Continuous Mortality Investigation
Institute and Faculty of Actuaries

High Age Mortality Working Party

Modelling population exposures at very high ages

Supplementary Technical Paper to Working Paper 100

June 2017

Contents

1. Introduction

This paper acts as a technical addendum to Working Paper 100 and sets out the CMI High Age Mortality Working Party's proposed framework for modelling population exposures at high ages.

Working Paper 100 includes a summary of this paper's findings, and no further summary is provided in this paper. In addition, we do not repeat the acknowledgements and the membership of the Working Party. There is some repetition of technical material in this paper, when read alongside Working Paper 100, in order to provide a complete narrative that can also be read independently.

This paper reviews the approach currently used by the Office for National Statistics (ONS) and its limitations, and then moves on to testing variants designed to address these. The structure of the paper is as follows:

- Section 2 sets our review in context and recaps the widely used Kannisto-Thatcher (KT) method for estimating high age population figures.
- Section 3 explains the current approach used by the ONS, alongside the approach used by the Human Mortality Database (HMD), and considers some key limitations.
- Section 4 describes variant approaches that could be adopted to address these limitations.
- Section 5 introduces our assessment of the variant approaches.
- Sections 6 – 10 summarise our analysis of the variant approaches under a range of diagnostics, examining both performance and impact relative to the current ONS methodology.
- Section 11 considers the constraints currently applied to the KT estimates and whether these are appropriate.
- Section 12 summarises the conclusions of our analysis, and sets out proposed modifications to the current methodology.

This paper has been informed by (and builds upon) previous research including the papers

- 'Phantoms never die: living with unreliable population data' (Cairns et al) and
- 'Accuracy of official high age population estimates in England & Wales: an evaluation' (ONS).

The ONS report provides an excellent review of the quality of high age population and deaths data for England & Wales, including known issues and suggestions for improving the estimation methods used. We are grateful to the ONS for inviting members of the Working Party to attend its workshop in early 2016 where initial results for the paper were shared and discussed. This work has helped to shape our analysis, which we hope builds on the foundations laid by the ONS review. We have addressed research areas suggested in their paper, including both the incorporation of a trend allowance in the projection of survivor ratios, and the investigation of alternative join ages, under the KT method.

A full list of the previous research and other material reviewed in this paper can be found in a References section at the end of Working Paper 100.

2. Background

2.1 Context

Estimating the population at high ages in England & Wales (E&W) is important as a prerequisite for:

- national population projections;
- estimates of current mortality rates by age (and hence life expectancy); and
- estimates of 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 mortality rates for portfolio curves at high ages (as described in Working Paper 100), potentially including the SAPS and Annuities base mortality tables at high ages where the underlying datsets lack credibility.

A brief history

Population estimates for E&W are currently published by the ONS on a yearly basis.

At ages up to 89, these are based on information collected as part of each 10-year census, and calculations to roll these forward to the middle of each calendar year in both census years and years in between each census.

Prior to 2007, the individual age estimates for ages 90+ were calculated for E&W by the ONS (and previously by the Government Actuary's Department) for use in producing the National Population Projections and compiling National Life Tables. These estimates were made available for research purposes but were not officially published. Since 2007 the ONS began to publish the 90+ estimates by individual age as experimental statistics, in response to increasing interest from stakeholders including actuaries. The estimates gained National Statistics status in 2011.

In recent years there has been ever greater focus on these high age estimates, due to:

- the increasing numbers of people living to very old ages (the proportion of the total population aged 90+ was 0.4% in 1991, 0.9% in 2015 and is projected to be 2.3% by 2039);
- the recent heavier (relative to preceding years) mortality in national mortality at high ages with higher-than-expected mortality rates observed in E&W in 2015 and 2016, for example; and
- questions raised over the accuracy of the figures. For example, the 2011 Census suggested there were around 31,400 fewer people aged 90 and over than had been estimated by rolling forward the 2001 Census data to 2011, and issues have been identified with the current estimation methods for particular birth cohorts[1].

The UK Statistics Authority has recently completed an assessment of the annual 'Estimates of the Very Old' data series, and the ONS has published a paper reviewing the methods used to obtain high age population estimates and the quality of the underlying input data.

In the sections that follow we set out our own consideration of the key issues surrounding high age population estimates in the E&W national dataset and how these might be addressed.

Mid-year estimates vs exposure estimates

The population estimates produced by the ONS for each calendar year are intended as estimates of the number of people alive *at the mid-point* of that year. In contrast, for the purpose of mortality analysis we are more interested in the population exposure *over* the calendar year – that is, the average number of people alive across the year. This is because it is a measure consistent with the number of deaths over the year, so the two can be validly compared to provide an estimate of mortality rates.

If the population size varies linearly over the year, then mid-year estimates and exposure estimates should be interchangeable. However, where the variation in population size is non-linear, this relationship will no longer hold.

The analysis and proposals throughout the rest of this paper are intended to estimate population exposures (*not* mid-year estimates):

- Most of our analysis concerns the underlying accuracy of the KT method, which is relevant to the estimation of both mid-year estimates and exposures.
- However, the element of our proposal which deals with exposure smoothing (Section 10) is fundamentally targeted at estimating exposure figures (not mid-year estimates) and encompasses the concept of convexity adjustment to allow for non-linearity, noted above.
- For the purpose of estimating mid-year population figures (in line with the ONS's current output) the convexity adjustment embedded in this element of our proposal would not necessarily be appropriate.

2.2 The Kannisto-Thatcher method

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 (for example, from census questionnaires rolled forward with allowance for births, deaths and migration) 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. A more thorough review of the accuracy of each type of data can be found in the ONS's December 2016 report 'Accuracy of official high-age population estimates, in England & Wales: an evaluation'.

Against this backdrop, the KT method was proposed by Thatcher et al in 2002. It is one of a wider class of survivor ratio methods that aims to construct high age population estimates from the (presumed accurate) deaths data at those ages.

The approach is as follows (we consider first a single cohort and then extend to cover a wider range of cohorts):

- First, identify those birth cohorts which are now fully extinct. This is usually taken to be the set of cohorts that would be aged over ω now, where ω is an assumed maximum age of survival. (The ONS takes $\omega = 120$.)
  - For these cohorts, we have complete information on all individuals that were alive in the past from the record of their subsequent deaths.
  - This means we can reconstruct the population aged x last birthday at the beginning of year t ($P_{x,t}$) by the recursive formula:

    $$P_{x,t} = P_{x+1,t+1} + D_{x,t}$$

    where $D_{x,t}$ is the number of deaths occurring during year t amongst those aged x last birthday at the start of the year.

  This approach is known as the method of extinct generations (or the method of extinct cohorts) and was first proposed by Vincent in 1951.
- Next, consider the remaining cohorts – those that are not known to be extinct by the current year (which we will denote T). Specifically, let us consider the oldest such cohort, with year of birth Y.

- This cohort will be very old (and hence almost extinct) by year T, but there may be a small number of survivors remaining.
- The starting point of the KT method (and other survivor ratio methods) is to estimate the number of such survivors by assuming that the ratio of current survivors to total deaths over the last k years is the same as it was for the previous cohort, with year of birth $Y-1$ (where both survivor and death figures are already known); this ratio is referred to as the survivor ratio. The formula to determine $P_{x,t}$ is therefore as follows:

$$P_{x,T} = S_{x,T} \times \sum_{i=1}^{k} D_{x-i,T-i}$$

where

$$S_{x,T} = S_{x,T-1} = \frac{P_{x,T-1}}{\sum_{i=1}^{k} D_{x-i,T-1-i}}$$

- Once this first estimate $P_{x,t}$ has been determined, the population estimates at previous ages along the cohort can be reconstructed by adding back in the recorded deaths as before. Then the completed figures for that cohort can be used to estimate the population of the next cohort (with year of birth $Y+1$) in the final year T, again using the survivor ratio method – and so on until all cohorts have been populated.
- The KT method introduces two refinements to this basic approach.
  - Firstly, the survivor ratios are based on an average across the last m cohorts (not just the most recent cohort), to reduce random volatility as the data becomes sparse.
  - Secondly, a global correction factor c is applied when projecting the survivor ratio from the previous cohort ($S_{x,T-1}$) to the current cohort ($S_{x,T}$), rather than simply assuming that the two are equal. The value of c is determined by constraining the total population in year T to its published value (which is assumed to be reliable in aggregate even though the individual population estimates by single year of age are not). The idea is that this accommodates a scenario in which mortality rates (and hence survivor ratios) are changing over time, for example due to an underlying tendency of mortality improvement between one cohort and the next.
- The adapted survivor ratio formula under the KT method is therefore

$$P_{x,T} = c \times S_{x,T} \times \sum_{i=1}^{k} D_{x-i,T-i}$$

where

$$S_{x,T} = \frac{\sum_{j=1}^{m} P_{x,T-j}}{\sum_{j=1}^{m} \sum_{i=1}^{k} D_{x-i,T-j-i}}$$

Join age

The KT method assumes that death is the only mode of change to population figures from one year to the next, and makes no allowance for other population flows such as migration. This may be reasonable at higher ages, but becomes less robust towards younger ages (where, in any case, the census-based estimates tend to be more reliable). As such, the KT method is typically used to construct population estimates down only to a certain age, below which the census-based estimates are adopted without adjustment. Correspondingly, the correction factor c in the method is determined by constraining to the published population total for ages greater than or equal to the join age.

Throughout the remainder of this paper we refer to this age as the join age for the KT method.

Parameter values

In their original description of the KT method, Thatcher et al proposed the values $(k, m) = (5,5)$ and a join age of 90, which are also the values adopted in the ONS's current implementation of the method. However, they noted that preliminary testing had produced 'no clear optimal combination' for k and m. Furthermore, the paper did not test alternative join ages (instead, it tested the approach of constraining to the 90+ population *total* against an alternative approach of constraining to the *single year of age* population figure at age 90, which it found to be considerably less robust).

We therefore review alternative parameterisations of the KT method amongst the variants tested in this paper.

3. Current approach and issues

3.1 ONS approach

The methodology currently used by the ONS to determine its high age population estimates for E&W is set out online and can be summarised as follows:

- For ages 89 and below, adopt the census-based figures by single year of age 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, apply the KT method with parameters $(k, m) = (5,5)$ and join age 90, using deaths data to estimate the single year of age population counts at ages 90+. Deaths data here means registration data to 1992, and occurrence data from 1993 onwards (except for the final year which is registration data).
  - Determine the correction factor c by constraining the total population estimate for ages 90+ in the final year to the census-based figure for that year.
  - Calculate single year of age estimates for ages 90 and over for previous years by adding back the deaths.
  - Finally, apply a (uniform) annual scaling adjustment to the resulting population estimates at ages 90+ *in each previous year*, to constrain the 90+ population total to the census-based figure for that year.
- The population estimates over age 104 are then grouped into a single 105+ estimate for publication.

Chart 3.1 illustrates how the ONS's current implementation of the KT method (as described in detail in Section 2.2) works.

- In the diagram, it is the gold area for which we need to estimate population exposures (starting with point B). The population estimates in the rest of the figure are known, or have already been determined.
- 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.
- It then applies the calculated survivor ratio to the single birth cohort containing point B, multiplying up by the known death count over the prior k ages in the cohort to give the population estimate for point B.

Chart 3.1: Current ONS implementation of KT method

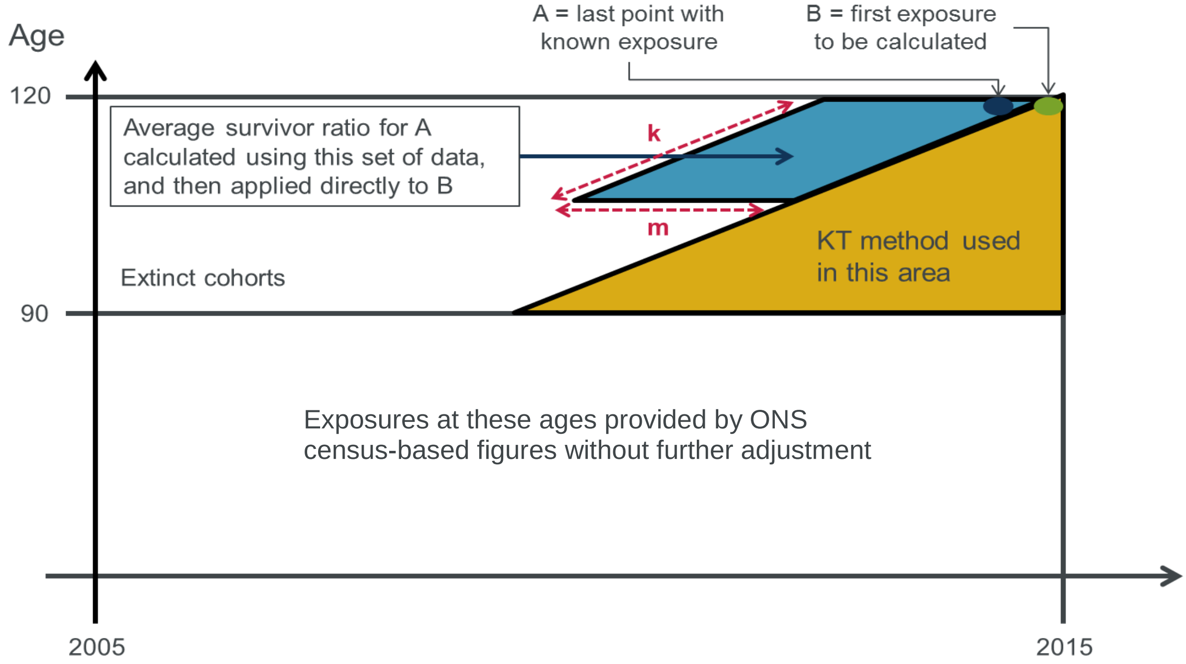

Subtleties

There are some subtleties in the detailed application of the KT method to the E&W population estimates of the ONS, in relation to the timing definitions of the data.

The KT method is designed to use:

- death counts during each year (based on age last birthday at the start of the year[1]); to construct
- population estimates at the start of each year (again, based on age last birthday at that point).

The deaths data available for E&W does cover deaths during each calendar year (as required) but the counts are by age at *date of death*, not age at the prior 1 January.

Furthermore, the population estimates prepared by the ONS are mid-year estimates (based on age at 1 July) rather than 1 January estimates.

As a consequence, the ONS first converts the raw death and population input statistics to the format needed for the KT method, using the approximations:

$$D_{x,t}^{KT\ input} = \frac{1}{2}\left(D_{x,t}^{raw} + D_{x+1,t}^{raw}\right)$$

$$P_{x,T}^{KT\ input} = \frac{1}{2}\left(P_{x,T-1}^{raw} + P_{x,T}^{raw}\right)$$

where:

- $D_{x,t}^{raw}$ is the number of deaths during year t of individuals aged x last birthday at date of death death; and
- $P_{x,t}^{raw}$ is the number of individuals alive at 1 July in year t who were aged x last birthday at that point in time.

[1] Note that in the context of the generic KT method, ‘year’ here does not necessarily mean a calendar year; it can be any period from one date to the day before that date in the following year. The ONS implementation of the KT method uses calendar year periods.

The KT method is then applied using the $D_{x,t}^{KT\ input}$ figures to produce 1 January population estimates up to and including 1 January T. These are constrained (using the KT correction factor c) so that the total 1 January T population estimate for ages 90+ matches the corresponding 90+ total from the converted population figures $P_{x,T}^{KT\ input}$ above. A final (extrapolated) estimate at 1 January T+1 is determined by ageing the population and subtracting the deaths over year T.

The official mid-year population estimates for year T and earlier years are then derived by interpolating the 1 January estimates, and scaling the resulting figures for each year so that the 90+ total matches the census-based total for that year.

The Human Mortality Database (HMD)

Before proceeding, it is worth briefly noting the approach to high age population estimates taken by the HMD, a collaborative international project that aims to provide researchers with online access to detailed and comparable national mortality data.

The HMD's methodology is described in full in its Methods Protocol, which is publicly available online. Like the ONS, it uses official national census-based estimates below a join age and applies the method of extinct generations (for extinct cohorts) and the KT method constrained to 90+ population totals (for non-extinct cohorts) above this age to produce population estimates by single year of age. An adjustment is made to the population estimates to derive exposures suitable for determining mortality rates.

However, the HMD approach differs in that it:

- adopts a more sophisticated Lexis adjustment approach to determining age at 1 January death counts from the raw age at deaths data, allowing approximately for the differing sizes of the two cohorts that contribute to the calendar year (rather than assuming a 50/50 Lexis split as in the ONS method);
- extends the death-based construction of population estimates back to a join age of 80 (not 90) for all cohorts that are extinct, and also for non-extinct cohorts that are aged 90 or older in the final year;
- constrains only the non-extinct cohorts in the KT method by reference to the 90+ population totals and only in the final year (there is no further constraining of the extinct or non-extinct cohorts in any earlier year).

This latter point means that, for all cohorts, the HMD population estimates at high ages are directly consistent with the raw deaths data, except potentially at the join age, which is not the case for the ONS estimates post-scaling. Our previous analysis in Working Paper 85 demonstrated this distinction. It does not, of course, mean that the HMD approach is necessarily more accurate for estimating the population of non-extinct cohorts (including, in particular, the population in the final year).

The diagram in Chart 3.2 is sourced from the HMD Methods Protocol and summarises the HMD methodology adopted in different regions of the dataset.

Chart 3.2: Methods used for population estimates by the HMD

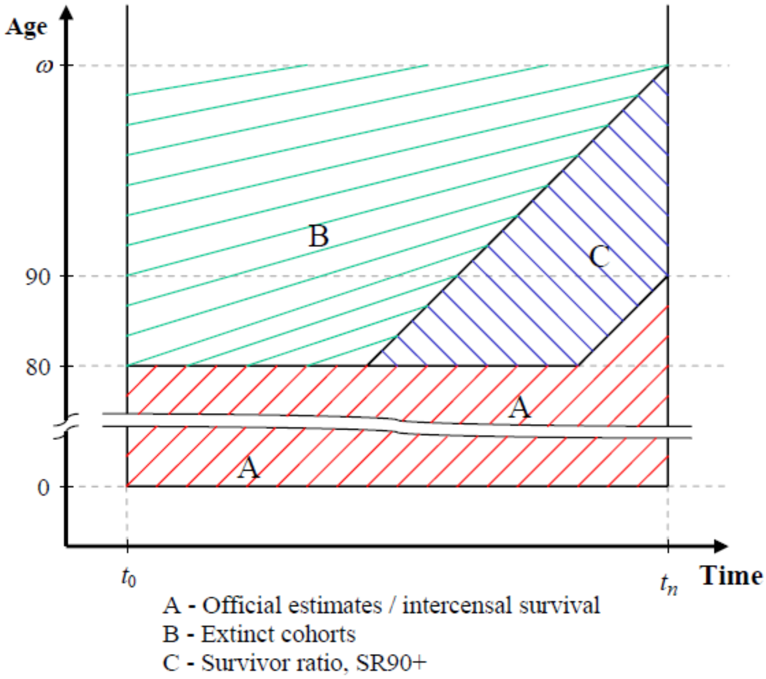

Source: HMD Methods Protocol

In Working Paper 85, we noted that we had not been able to replicate the HMD approach independently, so our interpretation of their approach should be considered provisional. Since the publication of that paper, we have now been able to replicate the HMD estimates to within a small margin. The main focus of our analysis is, in any case, on variants to the current ONS methodology (which we have also been able to replicate to a high degree of accuracy).

3.2 Accuracy of ONS population estimates

In December 2016 the ONS released its paper 'Accuracy of official high age population estimates in England and Wales: an evaluation'. Their findings are summarised below; we think these are consistent both with the findings of our earlier work in Working Paper 85 and the work that is presented in this paper.

Death reporting data

The Economic and Social Research Council (ESRC) Centre for Population Change at the University of Southampton funded the validation of deaths data recording for a sample of deaths in E&W of semi-super-centenarians (aged 105 to 109) during the period 2000 to 2014, in particular the matching of date of birth information given in the death certificate against birth records. 1,141 deaths in the sample were born in E&W or Scotland and, of these, 96% had a date of birth match between birth and death records, with a potential 1-2% further 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 also indicates a very high level of date quality validation. However, it should be noted that around 13% of male deaths and 8% of female deaths aged 105 and over during 2000 to 2014 were of people born outside the United Kingdom. These were not included in the sample as it can be very difficult to obtain birth certificates for those people.

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. However, there can be delays between death occurrence and death registration. To provide some comfort on this, 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 further the impact of both cross-border UK and international migration on population estimates. This is an important theme as our extinct generation mortality modelling in Working Paper 85 assumed that migration had a negligible impact. The work performed by the ONS concluded that:

- levels of cross-border UK migration are small for ages 90 and above (there is a net outflow from E&W equating to around 0.03% of the E&W population);
- cross UK migration from patient register (PR) 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;
- 0.2% of international migration is in respect of ages 90 and above – around 350 people per annum (0.07% of the age 90+ population).

The ONS has therefore concluded that migration has a minimal impact on the mid-year population estimates of lives aged 90 and above.

Kannisto-Thatcher modelling

The KT method relies on the two elements discussed, namely good quality deaths data and assuming negligible migration, to determine population estimates at very high ages.

In their paper, the ONS also considered the following areas of their implementation of the KT method:

- The KT method has also been performed 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 with population registers for each country for each individual age as at the start or end of each calendar year during this period.
  - The KT estimates fitted well to the population registers at all ages and times, although this fit deteriorates as the time to extinction increases (i.e. for younger ages and at recent calendar years), with KT estimates being lower than the population registers.
  - The fit is good on most ages up to 94 but deteriorates above this point, most likely because small deviations in small numbers of lives results in a proportionately larger difference.

- The KT method is used to estimate populations as at 1 January for each calendar year. Deaths data is provided by age at death in each calendar year, with an assumption being made on the uniformity of birthdays to determine the split by age at 1 January.
  - The ONS has compared this approach against the population registry as at 1 January for Finland.
  - Overall the difference between the census registers and the KT estimates is (relatively) minor, though there is a clear bias towards slightly overstated population estimates when an assumption of uniform birthdays is applied to the deaths data, for most of the period covered. The fit is improved when using the deaths data in the format required, with no need to split between birth cohorts.
- 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 synthetic populations.
  - A variety of different assumed mortality rate scenarios were modelled to assess the change in the correction factor.
  - It was found that the correction factor is close to 1 (i.e. no adjustment) when mortality is not changing over time, with bigger adjustments of up to 15% in either direction depending on the mortality scenario.
  - The ONS' implementation of the KT method therefore relies heavily on the total 90+ population being accurate.

Comparison of counts determined from administrative data

The ONS compared the KT estimates for E&W against Statistical Population Datasets which have been created by combining multiple administrative datasets: the National Health Service (NHS) Patient Register (PR), the Department for Work and Pensions (DWP) Customer Information Service (CIS) and the Higher Education Statistics Agency (HESA) databases.

For ages 90 and above, the administrative dataset returns population estimates for 2011 that are broadly similar to the unconstrained KT estimates (0.5% higher for males and 0.4% lower for females). However, a difference emerges over more recent years, with the administrative dataset estimates consistently higher than the unconstrained KT estimates for both males and females (for example, 6.7% higher for males and 3.3% higher for females in 2015).

Implications for our analysis

Members of the Working Party were invited by the ONS to attend its workshop in early 2016 where initial results for its paper were shared and discussed. This work has 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, relative to preceding years (notably 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 on the fit to data assuming a different join age to the currently assumed age 90. We have tested the impact of assuming a range of different join ages.

The ONS testing of deaths data has so far found it to be of high quality with only minor issues, so we have assumed in our work that the deaths data are accurate and complete, and reliable for use as inputs to the KT method. The ONS also reported that the impact of migration at older ages is minor, which is an important consideration for our testing of lower join ages.

3.3 Known issues

In this section, we set out potential concerns with the current ONS approach to estimating population exposures at high ages for E&W.

A number of recent papers have highlighted apparent anomalies with either published population estimates, or mortality rates in E&W. Deaths registration data is considered to be complete and not expected to contain material inaccuracies due to the method of registering deaths in E&W. The source of the anomalies is therefore most likely to lie within the exposures data.

- Inaccuracies in census-day population estimates:
  - The population figures derived from the census are estimates only. Despite ONS efforts to minimise errors, there is potential for mis-estimation at all ages.
- Population estimates at age 90 and above:
  - The ONS publishes estimates for the E&W population by single year of age up to age 89, with the population for ages 90 and above provided as a single figure. Estimates for the population aged 90 and above by single age are then determined using the KT survivor ratio method, constrained to the official mid-year population estimate of the total population aged 90 and over. This approach could lead to a discontinuity between the estimates for ages 89 and 90, and could also lead to inaccuracies in the estimates above age 90 if the KT method is not parameterised appropriately.
- Phantoms:
  - If a census estimate of the number of people at a given age is higher than the true figure then, even if the data on subsequent deaths and migration is accurate, these people would not be removed from the estimate. These phantoms would not die and so the exposure estimate would diverge further from the true figure over time (in percentage terms), until correction at the next census. This problem is particularly acute at older ages, where populations are low and mortality rates are high. A similar problem of negative phantoms holds if a census underestimates the true population.
- Potential for discontinuity at census dates:
  - Results from decennial censuses are a key input into national population estimates. There is potential for this to give rise to discontinuities if the approach taken is different for different censuses, or if the roll-forward approach used for inter-censal periods is inaccurate.
- Unusual patterns of births. The 1919 and 1920 cohorts are particularly affected by an unusual pattern of births following the return of troops at the end of World War I. Cairns et al noted two areas where this could cause issues:
  - The rolling-forward of census figures, from the 29 April 2001 census date to mid-2001, assumed an even spread of birthdays during each year at all ages. This is reasonable where birth rates were relatively stable over a year and from year to year. However, there are some birth cohorts for whom this is not a good approximation. This led to a material difference between the projected and actual populations that was revealed at the time of the 2011 Census and caused estimates to be revised. There is potential for similar problems to be present in the current data, though we understand the roll-forward method used after the 2011 Census may have been changed to allow more appropriately for the uneven distribution of births.
  - The use of mid-year population as an estimate of central exposed-to-risk seems reasonable, unless there is an irregular pattern of births in a calendar year. Where there is an irregular pattern of births, a convexity adjustment could be considered to the mid-year population estimate to determine an appropriate exposed-to-risk.
- Differences between ONS and HMD data:
  - Analysis by the Working Party in Working Paper 85 indicated that mortality rates calculated directly from death registrations for extinct cohorts are higher than mortality rates published by the ONS. Given that the same registered deaths are used in both approaches, the difference can be attributed to the differing population estimates. The differences are less pronounced when comparing extinct generation mortality with mortality data published by the HMD.
  - The difference between extinct generation and ONS approaches tends to increase during each decade. This is consistent with census inaccuracies underlying the ONS data increasing (in relative terms) as census data is projected over the course of each decade.

4. Variant approaches

4.1 Desirable features

Having reviewed the current approach to estimating high age population exposures, including the accuracy of the underlying data sources and known issues with the methodology, we now move on to considering potential variant approaches.

In this section, we set out key features which we suggest are desirable in the modelling approach. (We go on to develop some of these into specific diagnostic tests in Section 4.3).

- Consistency with official population estimates:
  - The method should produce results which line up closely with the official estimates at younger ages, with minimal inconsistences at the join to younger ages.
  - Where the official estimates are considered to be reliable at older ages, a method should line up as closely as possible to the high age totals within the official estimates (i.e. minimising the annual scaling adjustments that are applied each year to constrain to these totals).
- Predictive performance:
  - The method should perform well in predicting actual population figures, where the actual figures are known with relative certainty.
  - It should demonstrate stability in performing well across a range of different potential future scenarios, in relation to the population data and mortality levels/changes exhibited in the population.
  - This can be demonstrated with back-testing and synthetic data.
- Internal consistency:
  - Methods should produce results that contain no obvious anomalies or discontinuities either by age or calendar year, within or between birth cohorts.
  - The final population estimates should exhibit decrements consistent with the official deaths data.
  - Where testing demonstrates the official population estimates to be reliable at a total level (but not necessarily by individual year of age), the method for determining high age population estimates should produce a smooth transition to the official estimates at the join age.
- Ease of communication/implementation:
  - Each alternative method should be assessed not only by the appropriateness of the results it produces, but also by how easy it is to communicate/implement.
  - Where methods give similar results, simpler methods are preferred.

4.2 Variants considered

We have reviewed a range of variants on the current KT method adopted by the ONS (which was described in Section 3.1). These are summarised in Table 4.1, along with the broad motivation for including them in our analysis.

Table 4.1: Variants considered

Variant	Description	Motivation
Mortality trend	When projecting survivor ratios, allow explicitly for the average recent trend in survivor ratios over time to continue into the next year.	To better capture the local pattern of mortality improvements by age and time, leading to a more realistic projection.
Parameters k and m	Test 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).	Decreasing k and m means averaging over less data in the survivor ratio (and vica versa). We wish to investigate whether the improved resolution of doing so is offset by data noise or whether it leads to better performance.
Join age	Test the impact of reducing join age below 90. If join age is N this means that the KT method is extended down to age N (and constrained to the N+ population total).	We are concerned that the 90+ totals used to constrain the KT estimates may not be accurate. Reducing the join age and constraining to a larger population total (e.g. 85+ or 80+) may be more robust.
Lexis adjustments to deaths data	Adopt a more sophisticated approach to determining age at 1 January death counts from the age at death input data (using a Lexis triangle decomposition).	This approach is already adopted by some bodies (such as the HMD). It may lead to greater accuracy and improve the internal smoothness of the population exposures.
Exposure adjustments	Adjust the modelled population exposures for convexity and birth distribution (or apply pragmatic smoothing to similar effect).	Cairns et al identified anomalies in the current estimates arising from distributional effects. This variant aims to correct these.

Unconstrained variants

The ONS currently constrains the population estimates produced by the KT method to its official population total for ages 90+.We considered testing unconstrained variants of the method, in response to concerns over the integrity of the 90+ totals described previously. At this stage, however, we have limited the scope of our testing to the dimensions described above. Given that this includes:

- variants designed to capture the age shape of the population more effectively (e.g. mortality trend allowance and Lexis adjustments); and
- variants that constrain to alternative totals (for example, lower join ages permit more freedom in the high age population estimates)

we feel that this provides sufficient sensitivity for a preliminary analysis.

We note that a number of the variants produce correction factors close to 1 (and small annual scaling adjustments), which is similar to applying the KT method without constraint in any case. We expand on this in Section 11.

4.3 Diagnostics

In order to compare the performance of different variants, we have applied a set of numerical tests designed to draw out the desirable features listed in Section 4.1.

Table 4.2 summarises the diagnostics performed. Further details, including illustrative examples, can be found in Appendix 4 to this technical paper.

Table 4.2: Diagnostic tests

Test	Description	Motivation	Key consideration in assessment
Impact diagnostic	Assessment of: a) period life expectancy estimates by age (based on the raw average mortality rates implied for calendar years 2011-2015 inclusive) and b) population estimates by age (for calendar year 2015) relative to the current ONS approach.	Test whether a variant affects the population exposures to a material extent.	All variants
Final year balancing adjustment	This is the adjustment implied by the 'correction factor' c in the KT method, i.e. equal to $c-1$.	In general we prefer final year balancing adjustments close to zero, as these imply a lower discrepancy between the underlying KT method and the official population total to which it is constrained in the final year.	Trend, varying k and m and join age
Annual scaling adjustment	For each variant, and for each calendar year over the period 1972 to 2015, we have calculated the percentage adjustment required to scale the raw KT method exposures to match the official population total above the join age (in line with the constraint applied by the ONS). Our analysis shows the average magnitude of this adjustment across all years.	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.	Trend, join age
Cohort consistency	For a given age x (last birthday) and calendar year t, we define the cohort inconsistency as $$\left(P_{x-1,t-1} - P_{x,t} - D\right)/P_{x,t}$$ where • $P_{x,t}$ denotes the mid-year population estimate at age x (last birthday) for calendar year t, and • D is an estimate of the number of deaths between mid-year $t-1$ and mid-year t (who were aged x last birthday at mid-year t), given by $$D = \left(3D_{x-1,t-1} + D_{x,t-1} + D_{x-1,t} + 3D_{x,t}\right)/8$$	An internally consistent population should exhibit decrements from year to year which match the number of deaths reported. We prefer an average cohort inconsistency close to 0%.	Join age

Test	Description	Motivation	Key consideration in assessment
	where $D_{x,t}$ denotes the number of deaths aged x (last birthday) during calendar year t. We then calculate the average magnitude of this metric across all ages 71 to 104 inclusive, and all calendar years 1972 to 2015 inclusive.		
Smoothness of mortality across join age	We define the deviance in $\log m_{x,t}$ for age x and year t as $$\frac{1}{3}\sum_{j=x-2}^{x+2}\left(\log m_{j,t} - \log m'_{j,t}\right)^2$$ where $\log m'_{j,t}$ is the estimate of $\log m_{j,t}$ derived by fitting a straight line in $\log m_{j,t}$ to the 5 age points $j = x - 2$ to $j = x + 2$. We then calculate the average deviance in $\log m_{x,t}$ at the join age, across all calendar years 1984 to 2015 inclusive.	We expect mortality rates for each calendar year to vary smoothly by age – in fact, we expect $\log m_{x,t}$ to be locally linear in x (where $m_{x,t}$ is the crude central mortality rate for age x and year t). A particular concern is the smoothness in rates at the *join age*, where we transition from the ONS's census-based population estimates at younger ages to the (constrained) KT method at higher ages. An average deviance closer to 0% is preferred, as this implies a smoother transition of mortality rates across the join age.	Join age
Cairns Blake Dowds Kessler (CBDK) Diagnostics	**CBDK 1 – mortality rates by cohort** This is a plot of crude central mortality rates by age, for each of 5 adjacent birth cohorts. **CBDK 2 – concavity by cohort** This is a plot of the empirical concavity function $C(x,t)$ by time for a given birth cohort, where $$C(x,t) = \log m_{x,t} - \frac{1}{2}\left(\log m_{x-1,t} + \log m_{x+1,t}\right)$$ This function measures the log-linearity of mortality rates by age for a given calendar year t at age x.	The test is whether the pattern of mortality rates by age looks similar for these successive cohorts. We expect $\log m_{x,t}$ to be locally linear, so the test is whether the concavity function stays close to zero, without any systemic bias. We expect “convexity adjustment ratio - 1” to be close to zero.	Join age, lexis adjustment, exposure smoothing

Test	Description	Motivation	Key consideration in assessment
	We also calculate the convexity adjustment ratio (as described by Cairns et al). This 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. **CBDK 2 – concavity heatmap** A related diagnostic is the presentation of the concavity function $C(x,t)$ as a two-dimensional heatmap by age x and calendar year t. This should provide a more comprehensive picture of any structural concavity in the dataset.	The test is whether this heatmap simply shows random variation around a concavity value of zero, or whether it exhibits structural features indicative of anomalies in the data.	
Synthetic population modelling	One way to assess the accuracy of the KT method is to use synthetic populations constructed for the age ranges under consideration. For these artificial datasets, assumed mortality and migration rates can be input and death counts determined in the format required for the KT method. The resulting single year of age population estimates derived using the KT method can then be compared against the 'known' population figures to give an indication of the method's accuracy.	The key benefit of this approach is that it provides a means of testing the method's ability to estimate population figures correctly, for a range of hypothetical scenarios with potentially different features. A direct test of accuracy is, in contrast, not possible for the ONS dataset (because there is no reliable source of 'correct' figures against which to compare the KT estimates).	Trend, varying k and m

The synthetic population analysis in this paper focuses on six populations, constructed from a range of alternative mortality improvement structures. To assist the reader, we summarise these six scenarios in Table 4.3. To ensure the robustness of our conclusions we have deliberately covered a wide range of scenarios, including:

- simple improvement regimes (1 and 2);
- regimes with more complex features (3 and 4); and
- cases with mortality modelled on real – raw or smoothed – experience (5 and 6).

As the populations are (by definition) contrived, it is important to consider performance across the whole range of different scenarios when testing variants of the KT method. For example, populations 1 and 2 would be expected *a priori* to support the introduction of a survivor ratio trend because they represent artificially smooth mortality trend regimes themselves. It is of course still reasonable to check this, but the trend variant should only be preferred if it performs more robustly across the wider set of scenarios (and the actual ONS dataset).

Table 4.3: 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

5. Results format

We have analysed the performance of different variants under the full range of diagnostic tests set out in the previous section. In the sections that follow, we consider each element of the methodology in turn and draw out some key features of the diagnostics that shed light on the performance of the alternative approaches tested:

- mortality trend;
- parameters k and m;
- join age;
- Lexis adjustments to deaths data;
- exposure adjustments.

Where we refer to the base variant, we mean our implementation of the ONS's *current* approach to estimating high age population exposures (as outlined in Section 3.1).

When testing alternative methods on the ONS E&W data, our analysis presented in this section focuses on the male dataset, because

- there are lower data volumes towards high ages for males than females (which means that it is arguably more important for the KT method to perform reliably, and that the dataset itself serves as a more robust test of performance) and
- male pension and insurance liabilities are typically more prevalent (and hence material to practitioners) than female liabilities due to historic workforce practices in the UK.

We have also tested the sensitivity of our analysis to the female ONS dataset for E&W and we are satisfied that our conclusions remain robust. Although we have not shown the full analysis for females in this paper, we have included an overview of the key diagnostics and life expectancy impacts of our proposal in Section 12 (alongside the corresponding overview for males).

Note that, for analysis based on the ONS data, our modelling uses the death registrations data by single year of age up to 104 and then applies an exponential run-off pattern to the ONS's published death count for ages 105+ to produce deaths at individual ages 105 and over. If the ONS has used more granular 105+ data in their actual application of the KT method then our implementation of the method may be slightly different to theirs at the highest ages. Our approach should nonetheless be fit for purpose as a benchmark for comparing variant performance.

The ONS data used in our analysis is detailed in a References section at the end of Working Paper 100.

6. Mortality trend

Motivation

Chart 6.1 plots the actual progression of survivor ratios in the male ONS data by year (to be precise, the average year of the survivor figures from the five cohorts included in each ratio). The survivor ratios have been exhibiting an upward trend over time as mortality rates have decreased, whereas the KT method starts simply by projecting the last observed ratio into the final year (to the right of the dashed line). If prevailing mortality rates have changed then the survivor ratio based on the past data would need adjusting before it is applied. Between 1995 and the early-2010s, the survivor ratio has increased to reflect the period of sustained mortality improvements.

Chart 6.1: survivor ratio progression – current method for ONS males – ages 90 - 120

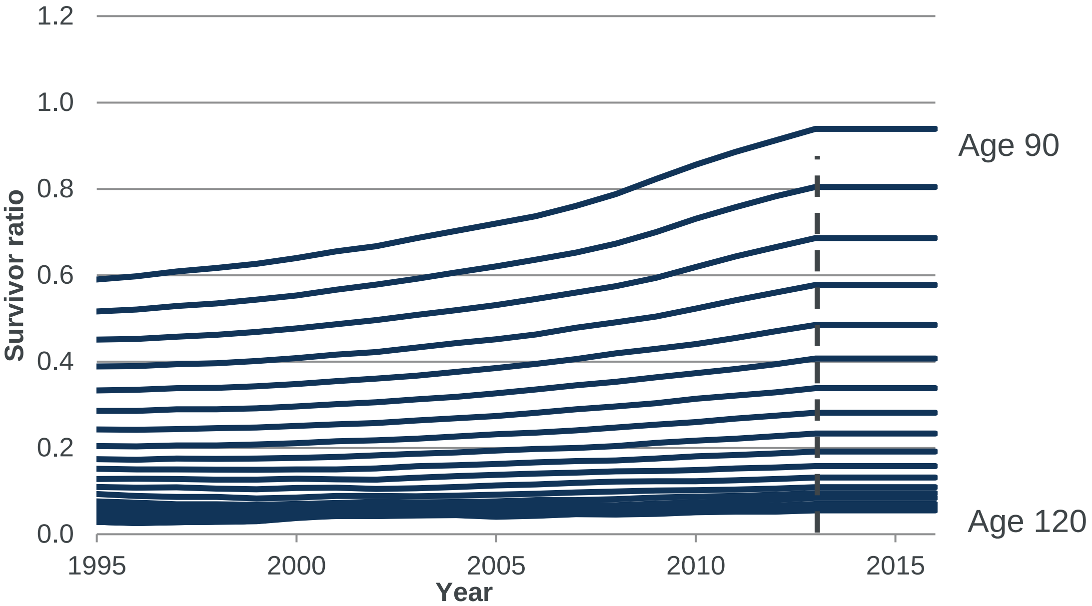

The current KT method permits one single overall adjustment by virtue of the global correction factor, c, described in Section 2.2. This adjustment implicitly captures the average rate of mortality improvement across all ages and years under the survivor ratio method. A trend of reducing mortality rates (i.e. mortality improvement) implies an increase in the survivor ratios over time and hence a correction factor greater than 1, which is what we see in practice. (Our implementation of the ONS's current approach for data up to 2015 leads to a correction factor of 1.083.)

In their December 2016 paper the ONS noted that there may be merit in adopting a more sophisticated approach, with an *explicit* allowance for survivor ratio trends in the KT method projection.

- This may be better able to capture local variations in mortality improvement by age and time, and so lead to a more robust projection across the age range.
- It may also reduce the need for a large global correction factor, which (though ensuring consistency with the official high age data in aggregate) risks producing artificial discontinuities around the join age, and inconsistencies from year to year.

National mortality experience in recent years has brought this area into sharp focus and underlined the need for a model which better reflects the structure of mortality trends. Improvements in mortality since 2011 (and most particularly over 2015) have diverged from the longer-term trend of the previous decade, and have exhibited quite different patterns for younger- and older-aged pensioners. A single global trend correction may be inadequate to account for such features.

We have therefore tested a version of the KT method which allows for a simple 5-year linear extrapolation of the trend in past survivor ratios by single year of age to estimate the survivor ratio for the final calendar year at each age. This is illustrated in Chart 6.2 – rather than use a single blue parallelogram (as illustrated previously in Chart 3.1) to determine the survivor ratio, we are now looking at a sequence of blue parallelograms through time and extrapolating the (linear) trend in survivor ratios to point B.

The choice of 5 years as the trend calibration period was motivated by a desire to project from just the most recent short-term trajectory, whilst ensuring a sufficient period of data to capture that trajectory. In preliminary analysis, we also tested a geometric approach to extrapolation (which broadly corresponds to a linear model for mortality improvement). However, this proved to be less robust at the highest ages because it was overly sensitive to survivor ratios close to zero.

Chart 6.2: KT method with trend

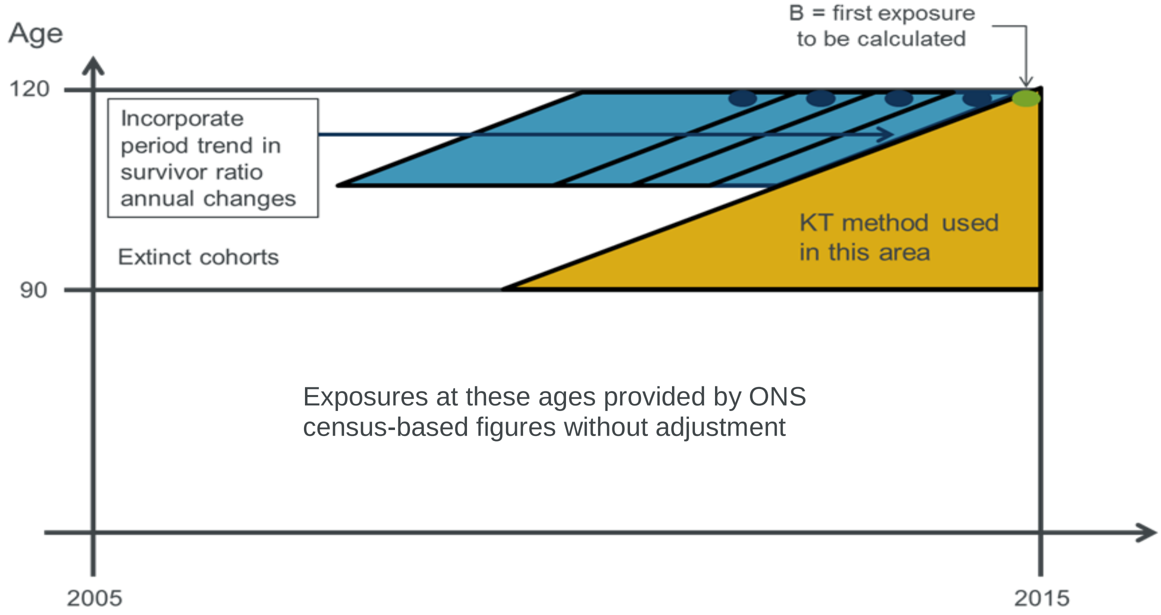

Results

As a first check, we tested both the current KT method and the trend variant on a very simple synthetic population, which assumed no changes at all in age-specific mortality rates from year to year.

- In this case, the age-specific population estimates produced by both versions of the KT method were identical to the 'true' figures from the synthetic populations for all ages and years. For the other synthetic populations tested, this was not the case.
- This suggests that the KT method produces the correct populations where there is no change in mortality rates over time but requires some adjustment where this is not the case.

We then compared performance on the range of synthetic populations from Section 4.3, before turning our attention to the ONS dataset itself. The full results across all six synthetic populations are included in the Excel Appendix for completeness – we present below just a subset of the output, focusing on Populations 4, 5 and 6, to illustrate some of the key features. (Unsurprisingly, the trend variant of the KT method exhibited clear outperformance for Populations 1 and 2, which are themselves modelled on a regime of uniform mortality trends – whilst this is comforting, we do not believe it merits further attention here.)

Synthetic populations – life expectancy and population estimates

In Chart 6.3, the tables on the left-hand side show the impact on life expectancy (estimated vs actual) across a range of alternative join age sensitivities. As discussed previously, the life expectancy calculations are based on survival truncated to age 106, to avoid excessive volatility from sparse data at the very highest ages.

The charts on the right-hand side show the underlying population exposures which are driving this impact, with calendar year 2015 as an example. The curves in this chart are plotted up to the point at which each population falls below 15 individuals, to avoid spurious rounding noise in the ratios shown at the very highest ages.

Chart 6.3: Error in estimated period life expectancy (LHS) and 2015 population exposures (RHS) by join age for synthetic populations, (k,m) = (5,5)

Population 4 Mortality reduction 2% p.a. (with cohort effects)

Trend	Actual	No trend	No trend	No trend	No trend	Trend	Trend	Trend	Trend
Join age		90	85	80	75	90	85	80	75
Parameter k		5	5	5	5	5	5	5	5
Parameter m		5	5	5	5	5	5	5	5
Life expectancy from age 65	21.70	-	-	-0.1%	-0.2%	-	-	-	-
Life expectancy from age 70	18.09	-	-	-0.1%	-0.2%	-	-	-0.1%	-0.1%
Life expectancy from age 75	14.73	-	-0.1%	-0.2%	-0.3%	-	-	-0.1%	-0.1%
Life expectancy from age 80	11.72	-	-0.1%	-0.2%	-0.9%	-	-0.1%	-0.1%	-0.2%
Life expectancy from age 85	9.07	-	-0.1%	-1.0%	-1.6%	-	-0.1%	-0.4%	-0.4%
Life expectancy from age 90	6.76	-	-0.9%	-1.7%	-2.3%	-	-0.6%	-0.8%	-0.9%
Life expectancy from age 95	4.65	-0.2%	-1.0%	-1.7%	-2.1%	-0.3%	-0.8%	-1.0%	-1.1%

No trend

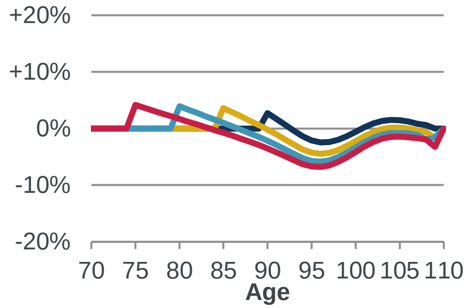

Trend

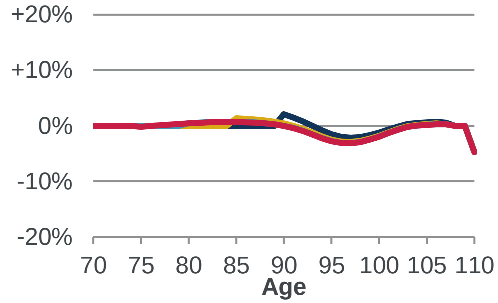

Population 5 Raw E&W

Trend	Actual	No trend	No trend	No trend	No trend	Trend	Trend	Trend	Trend
Join age		90	85	80	75	90	85	80	75
Parameter k		5	5	5	5	5	5	5	5
Parameter m		5	5	5	5	5	5	5	5
Life expectancy from age 65	19.31	-	-	-	+0.1%	-	-	+0.1%	+0.2%
Life expectancy from age 70	15.52	-	-	+0.1%	+0.2%	-	-	+0.1%	+0.2%
Life expectancy from age 75	12.15	-	-	+0.1%	+0.3%	-	-	+0.1%	+0.3%
Life expectancy from age 80	9.07	-	-	+0.2%	+0.9%	-	-	+0.2%	+1.1%
Life expectancy from age 85	6.45	-	-	+0.9%	+1.7%	-	-	+1.1%	+1.9%
Life expectancy from age 90	4.34	-0.1%	+0.1%	+0.9%	+1.7%	-0.1%	+0.3%	+1.2%	+1.9%
Life expectancy from age 95	2.83	-1.5%	-1.6%	-1.0%	-0.2%	-1.4%	-1.4%	-0.9%	-0.3%

No trend

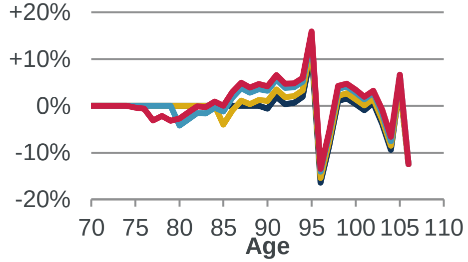

Trend

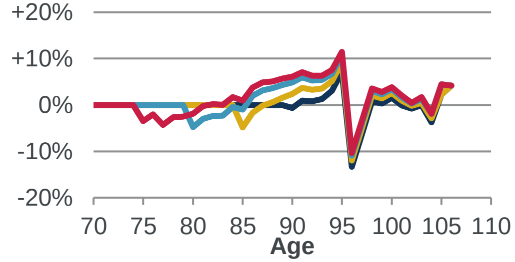

Population 6 Smoothed E&W

Trend	Actual	No trend	No trend	No trend	No trend	Trend	Trend	Trend	Trend
Join age		90	85	80	75	90	85	80	75
Parameter k		5	5	5	5	5	5	5	5
Parameter m		5	5	5	5	5	5	5	5
Life expectancy from age 65	19.45	-	-	+0.1%	+0.1%	-	-	-	+0.1%
Life expectancy from age 70	15.68	-	-	+0.1%	+0.2%	-	-	-	+0.1%
Life expectancy from age 75	12.28	-	-	+0.1%	+0.3%	-	-	-	+0.1%
Life expectancy from age 80	9.19	-	-	+0.2%	+0.8%	-	-	+0.1%	+0.4%
Life expectancy from age 85	6.56	-	-	+0.8%	+1.6%	-	-	+0.4%	+0.8%
Life expectancy from age 90	4.43	-	+0.3%	+1.0%	+1.7%	-	-	+0.3%	+0.7%
Life expectancy from age 95	2.87	-	+0.2%	+0.7%	+1.3%	-0.1%	-0.2%	-0.1%	+0.3%

No trend

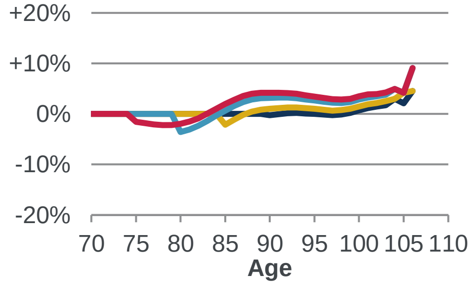

Trend

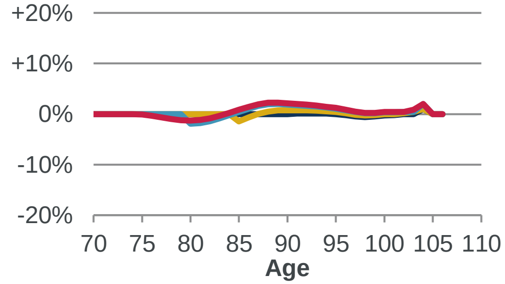

Join age 90 | Join age 85 | Join age 80 | Join age 75

In both cases, the range of alternative join ages is included purely to give a full picture of the robustness of the trend vs no trend variants under the application of different population constraints. It would not be appropriate to use these diagnostics to compare performance *between* different join age values. This is because, for the synthetic populations, we know that the high age figure to which we are constraining is 'correct'. Using a higher join age must therefore lead to a tighter fit, but tells us nothing about the most appropriate join age to use when constraining to the *ONS population data* (for which we don't know the correct population totals). We go on to consider join age more fully in Section 8.

Observations

For populations 4 and 6 the variant incorporating mortality trend delivers materially more accurate population exposures (and hence smaller errors in life expectancy) than the base variant in most scenarios. Furthermore, the accuracy of the exposure estimates is more robust to the choice of join age than the base variant, with a much tighter spread of figures in each case.

For population 5 the overall performance of the two variants is similar. It perhaps stands to reason that the trend variant may not outperform in cases where the improvement pattern exhibits more volatility from year to year, as this feature does not present a straightforward trend in improvement that can be easily extrapolated.

Furthermore, population 5 incorporates the same change in prevailing mortality improvement trends as the actual ONS data for E&W males (transitioning from relatively higher improvement rates in the decade up to 2010, to lower improvement rates over the past 5 years). It may be that this recent change in improvement rates is impairing the ability of a simple trend extrapolation to project survivor ratios effectively.

We should not disregard population 5 lightly, as it reflects the actual pattern of (noisy) mortality experience in ONS data, which is the form of data our method will ultimately need to deal with. Having said that, and bearing in mind the exceptional features of recent national mortality experience (with significantly lower improvements than the prevailing trend of the past decade), it may be that we are exposing the trend method to an *atypically* awkward region of data by focusing on the period ending in 2015. There is scope for further testing (against historic periods of actual mortality improvement data, and against reliable national population datasets from other countries) to establish whether the trend method outperforms in more typical circumstances.

In the meantime, it is important to note that the trend variant does not materially *underperform* the base variant in this case (or indeed for any of the six synthetic populations), which **indicates it is potentially more robust than the current approach, even under atypical scenarios**.

Synthetic populations – constraint adjustments

Chart 6.4 shows the range of final year balancing adjustments by join age for the trend variant compared with the base variant.

Chart 6.5 shows the average annual scaling adjustments for the two variants.

These are key measures of internal consistency. The charts demonstrate that in the majority of scenarios the trend variant of the KT method requires less adjustment to meet the total population constraints than its base counterpart. Again, there are some scenarios where the methods perform similarly (e.g. the relative performance differs by join age for population 5) but no cases, across these or the other three synthetic populations, where the trend variant of the KT method significantly underperforms the current approach.

This suggests that incorporating a trend allowance has improved the ability of the underlying method to correctly estimate the population exposures before any constraint is applied.

Chart 6.4: Final year balancing adjustment by join age for synthetic populations (k,m) = (5,5)

Population 4 – Mortality reduction 2% p.a. (with cohort effects)

No Trend

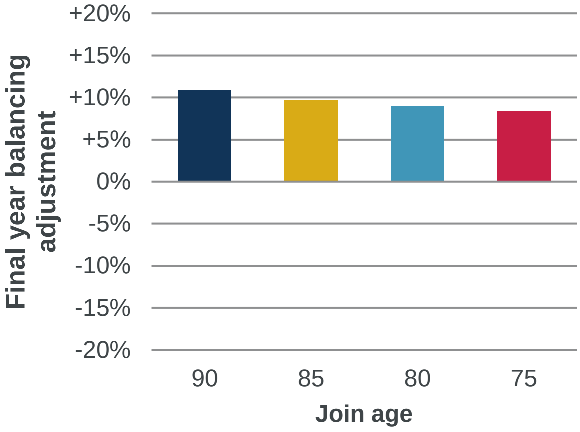

Trend

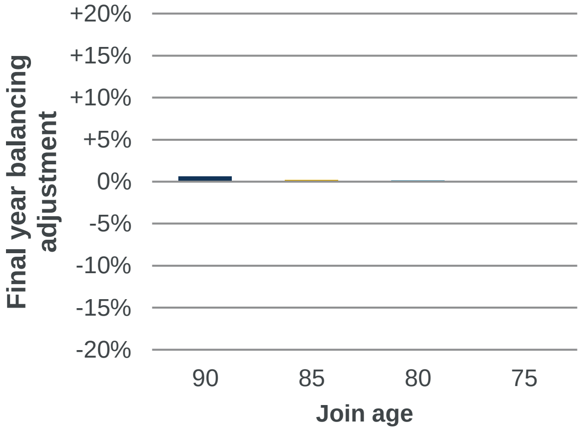

Population 5 – Raw E&W

No Trend

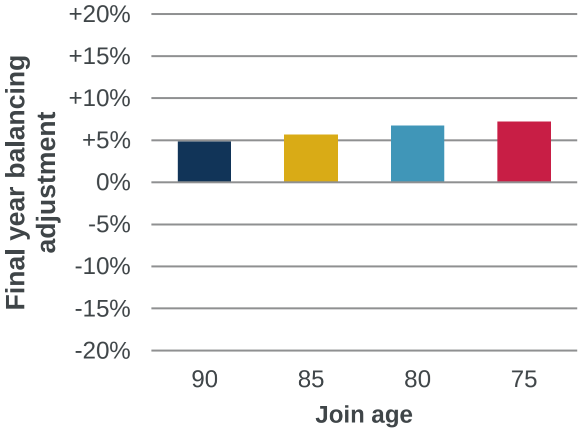

Trend

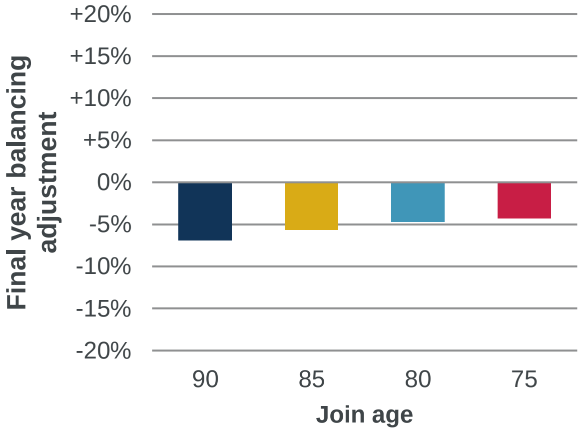

Population 6 – Smoothed E&W

No Trend

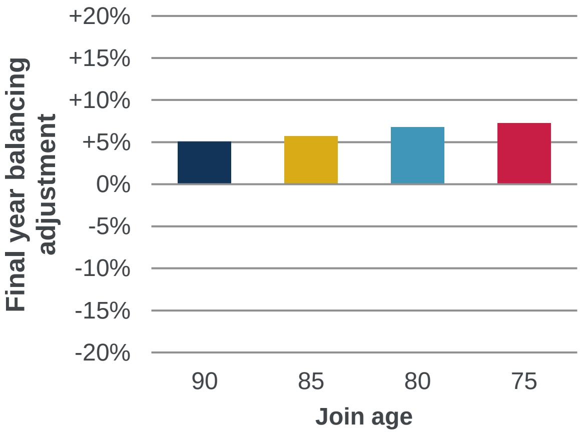

Trend

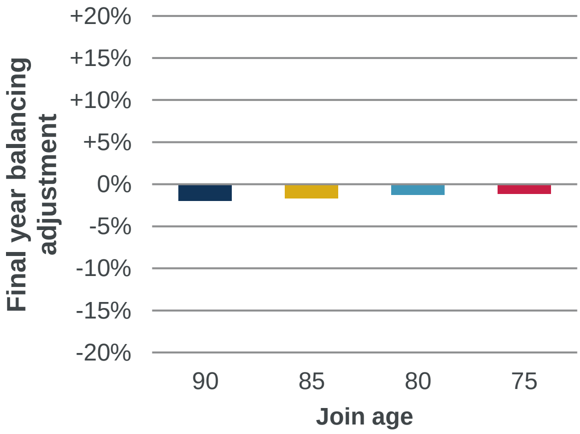

Chart 6.5: Average annual scaling adjustment by join age for synthetic populations (k,m) = (5,5)

Population 4 – Mortality reduction 2% p.a. (with cohort effects)

No Trend

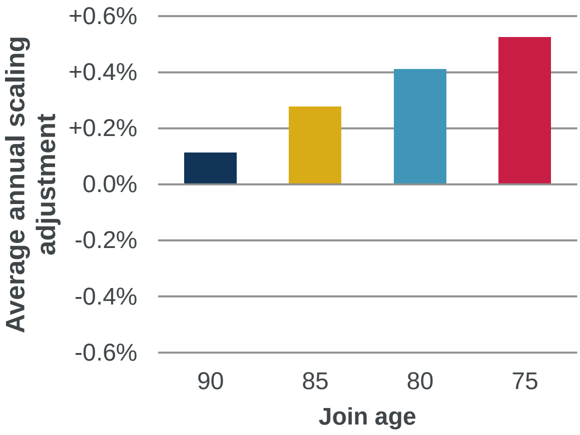

Trend

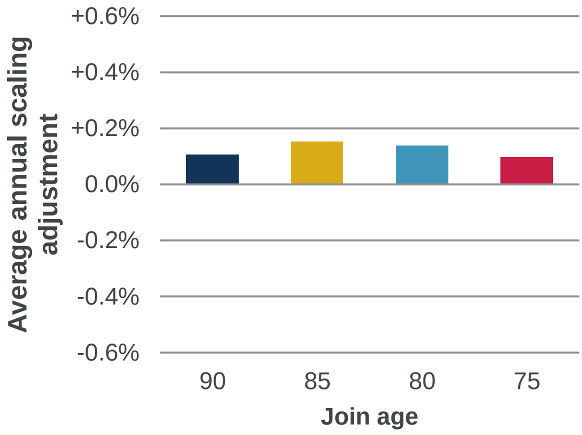

Population 5 – Raw E&W

No Trend

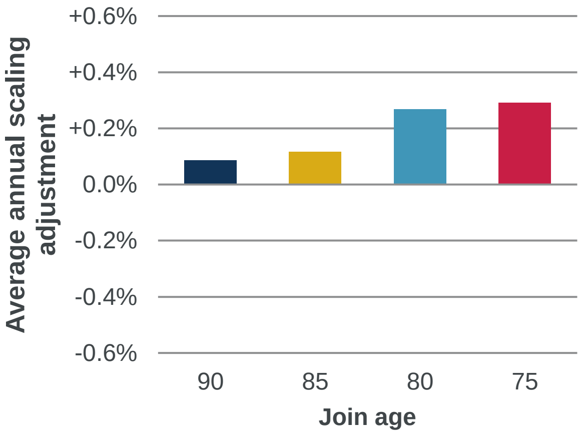

Trend

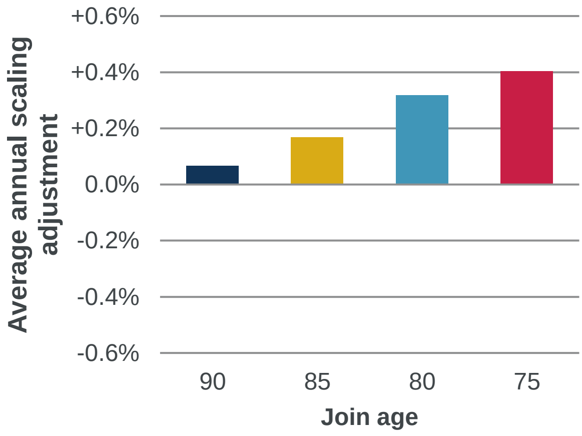

Population 6 – Smoothed E&W

No Trend

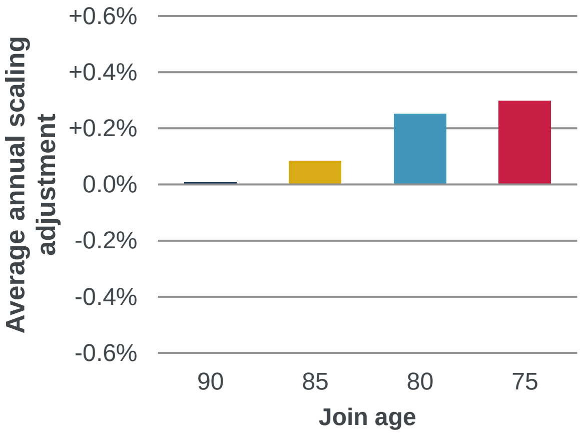

Trend

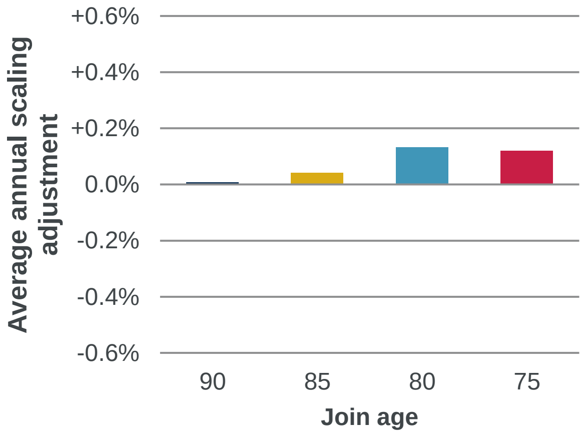

ONS male data – constraint adjustments

When applied to the ONS male data, the trend variant produces a tighter spread of final year balancing adjustments than the base variant, both:

- for calendar year 2015 (Chart 6.6), and
- when back-testing the KT method on periods ending in each calendar year from 1984 to 2015 inclusive (see the Excel Appendix).

It is interesting to note that in Chart 6.6, as well as being smaller in magnitude, the final year balancing adjustments for 2015 have also become negative under the trend variant of the KT method. This makes intuitive sense:

- The original balancing adjustments were positive because they were doing the job of capturing the average tendency to increasing survivor ratios over time (i.e. mortality improvement);
- Under the trend method this is no longer needed, because the survivor ratio projection captures recent past improvement trends directly. However, 2015 was a year in which mortality rates were generally higher than would have been expected based on the trend of recent prior years. As such, a small negative balancing adjustment is needed to constrain the projected population estimate above the join age to the ONS official total for that year.

Chart 6.6: Final year balancing adjustment in 2015 by join age for ONS male data

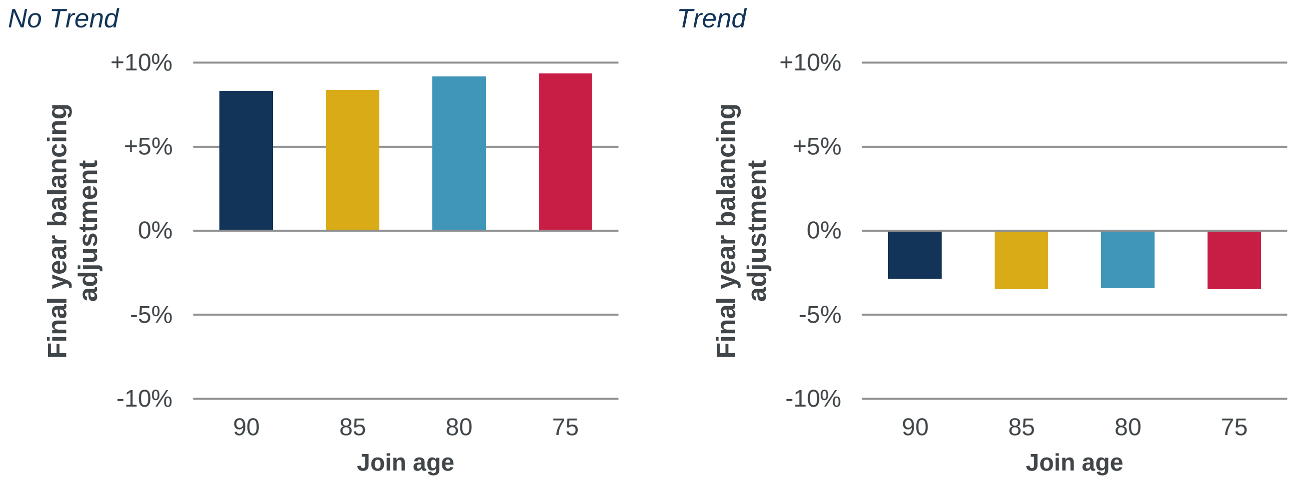

When applied to the ONS male data, the trend variant and base variant deliver similar performance under a number of our other diagnostics (e.g. average annual scaling adjustment, cohort inconsistency metric and smoothness of mortality across join age). Part of this may be due to the noisy pattern of improvements in the ONS data, similar to synthetic population 5, which is less conducive to the extrapolation of clear trend signals. However, it is important to note that incorporating trend does not appear to *weaken* performance under any of the metrics.

Conclusion and proposal

Allowing for mortality trends explicitly within the KT method does not lead to better performance in all scenarios. In particular, it may struggle to improve on the current approach for realistic data with noisy improvement structures.

It may be that a more sophisticated model of mortality improvements is better able to capture signals in the data and could improve the projection of survivor ratios. For example, one could attempt to include a cohort component to mortality improvements alongside the age-period trend approach we have tested.

Having said that, the simple trend variant does appear to function robustly across a wide range of scenarios (including for noisy data) and has the potential to deliver more accurate population estimates in certain circumstances.

We suggest that this provides sufficient motivation to consider incorporating the mortality trend allowance within the KT method. We estimate that the impact of doing so would be relatively modest, with a reduction in male period life expectancy of around 0.1% at age 90 and 0.4% at age 95 (based on the ungraduated average mortality rates for 2011-2015 inclusive).

7. Parameters k and m

Motivation

The current ONS method uses a value of 5 for both m (the number of past cohorts) and k (the number of prior ages along each cohort over which the deaths are summed) in calculating the historical survivor ratios.

This is in line with the 2002 paper by Thatcher et al, which considered alternative values m = 1 to 10, and k = 1 to 10, in each case testing the agreement of the KT estimates with official figures for countries with high quality population data. The paper showed that (k,m) = (5,5) appeared to perform well, but alternative parameterisations were possible with 'no clear optimal combination'.

The fundamental tension here is between:

- resolution (for which lower values of k and m are preferred, giving a more locally relevant survivor ratio to project into the next year); and
- stability (for which higher values of k and m are preferred, giving a larger region of data from which to calculate the survivor ratio).

There may also be an interaction with allowance for mortality trend in the KT method (which, in our suggested variant, involves the extrapolation of survivor ratios from 5 years' of past cohorts, and so may be better suited to lower values of m) and anomalies in the data (for example, different values of m and k change the region of data from which the survivor ratios are projected, and may place different weight on data with known anomalies such as the 1919/20 birth cohorts).

To help address these points we have tested a wide range of alternative m and k values, under both the current ONS method and the trend variant described above, in application to both the official ONS dataset and a range of synthetic datasets with different features.

Results

In testing alternative values of k and m, we again take the synthetic population scenarios as our starting point.

Synthetic populations – population estimates

Chart 7.1 shows the error in estimated population exposures for calendar year 2015 across a range of alternative (k,m) sensitivities, for the mortality trend variant of the KT method, for the synthetic populations 1 and 2 set out in Section 4.3.

The results in the left-hand charts suggest that a smaller survivor ratio window with higher resolution may deliver more accurate population estimates, although the very smallest values produce instability.

Reducing (k,m) all the way to (1,1) leads to unstable performance across the populations tested (for example, see left-hand charts for populations 1 and 2). Of the remaining values 2 to 5, it is (k,m) = (2,2) that appears to deliver the best performance in this set of charts, though the impact varies by age. For example, this parameterisation performs best in the left-hand charts for populations 1 and 2. Drilling into this further, we can test the impact of varying k (middle charts) or m (right-hand charts) independently from a baseline of (k,m) = (2,2).

The middle charts suggest that small values of k (e.g. 1) may lead to poorer performance, possibly due to the sparser data (and hence greater data noise) when calculating survivor ratios. This is seen particularly for populations 5 and 6. It further suggests that for values of k greater than 2 or 3 the estimates appear relatively insensitive to the precise figure chosen. Retaining the current value k = 5 appears to perform well (for example, it outperforms lower values of k in the middle chart for population 1).

The right-hand charts indicate the sensitivity of the population estimates to the value of m. There is no one value of m which performs best in all scenarios, but it appears that a value around 2 may be the most robust overall. In particular, m = 2 performs the best (amongst values 2 to 5) for populations 1 and 2. **Although not shown in this paper, our analysis of consistency metrics (the final year balancing adjustments and average annual scaling adjustments) also bears this out, suggesting that reducing the size of the survivor ratio window with a lower value of m may lead to a better fit.**

Chart 7.1(a): Error in estimated 2015 population exposures by (k,m) for synthetic populations (KT with trend, join age = 90)

Population 1 – Mortality reduction 2% p.a.

k = m

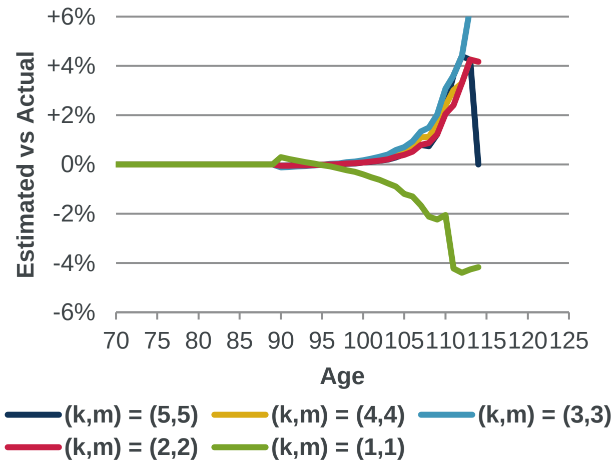

m = 2

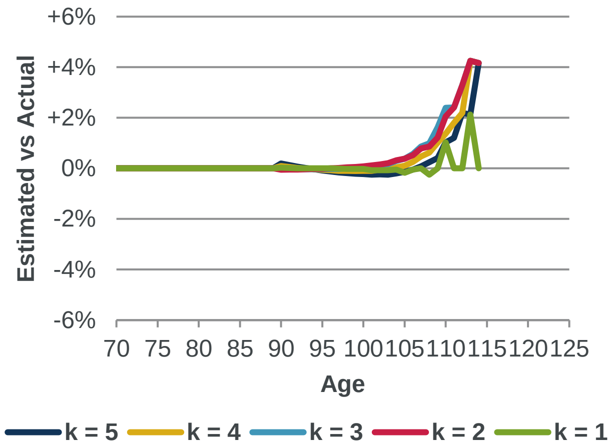

k = 2

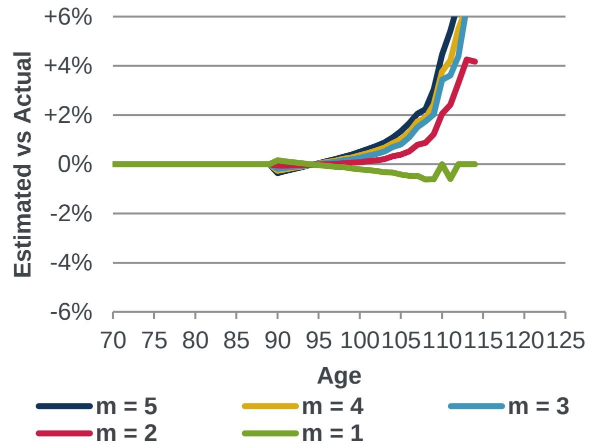

Population 2 – Mortality increase 2% p.a.

k = m

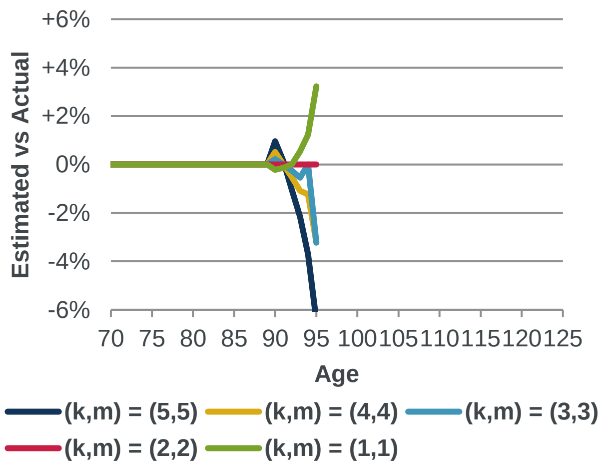

m = 2

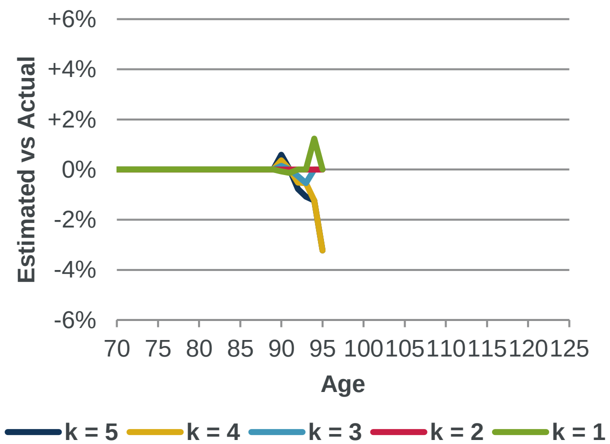

k = 2

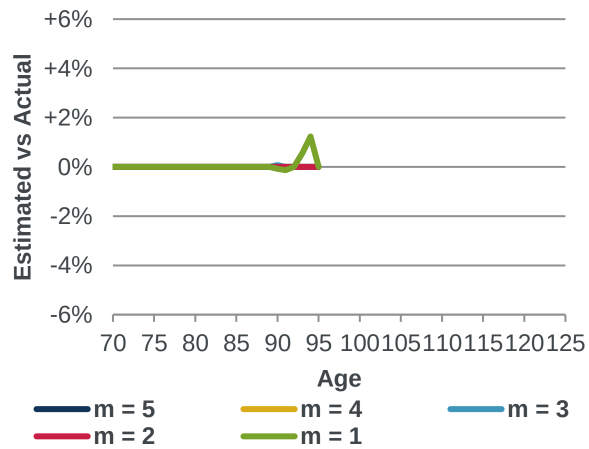

Chart 7.1(b): Error in estimated 2015 population exposures by (k,m) for synthetic populations (KT with trend, join age = 90)

Population 3 – 2015 shock

k = m

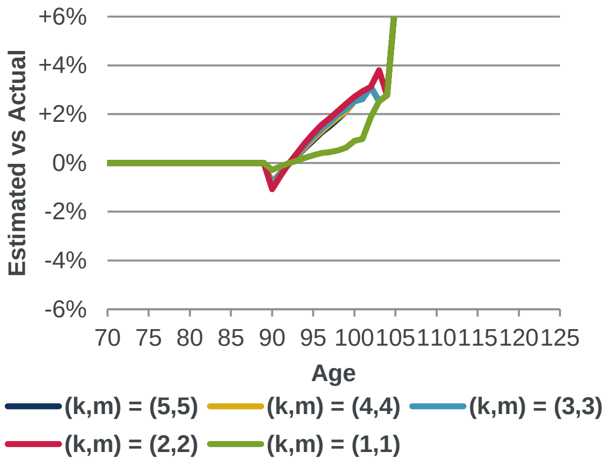

m = 2

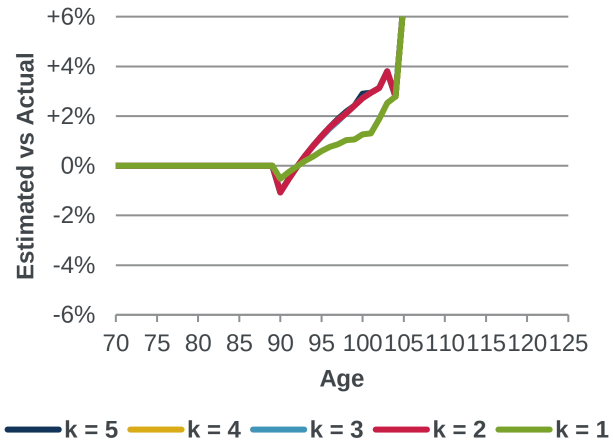

k = 2

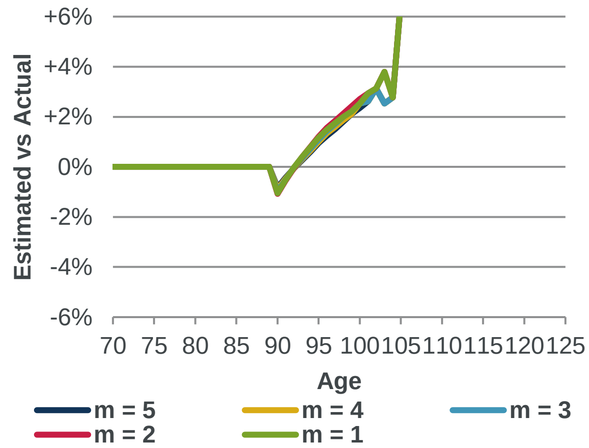

Population 4 – Mortality reduction 2% p.a. (with cohort effects)

k = m

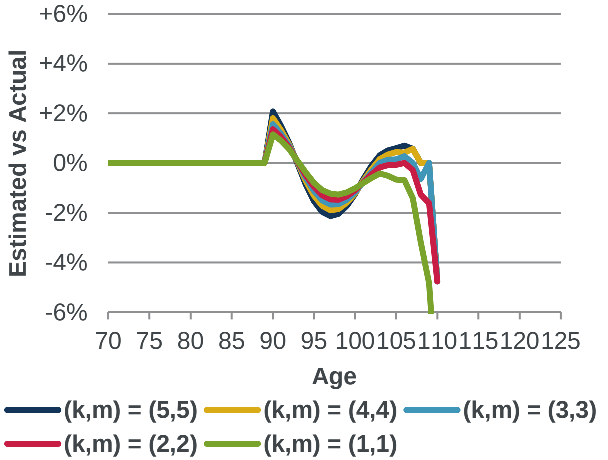

m = 2

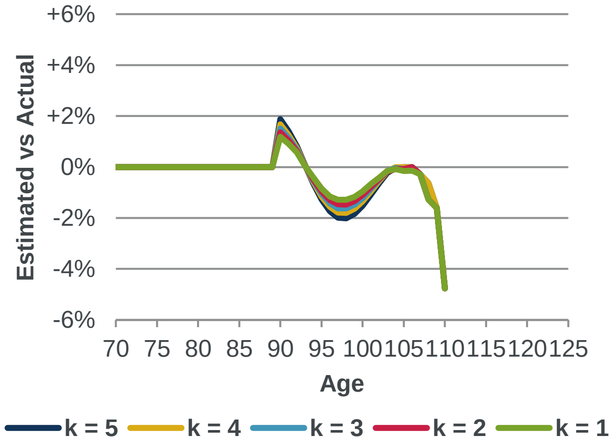

k = 2

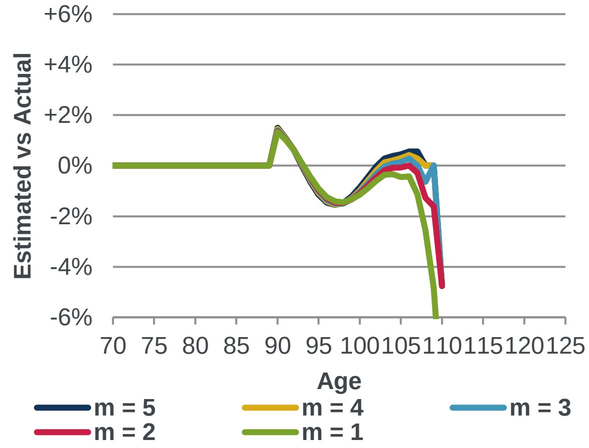

Chart 7.1(c): Error in estimated 2015 population exposures by (k,m) for synthetic populations (KT with trend, join age = 90)

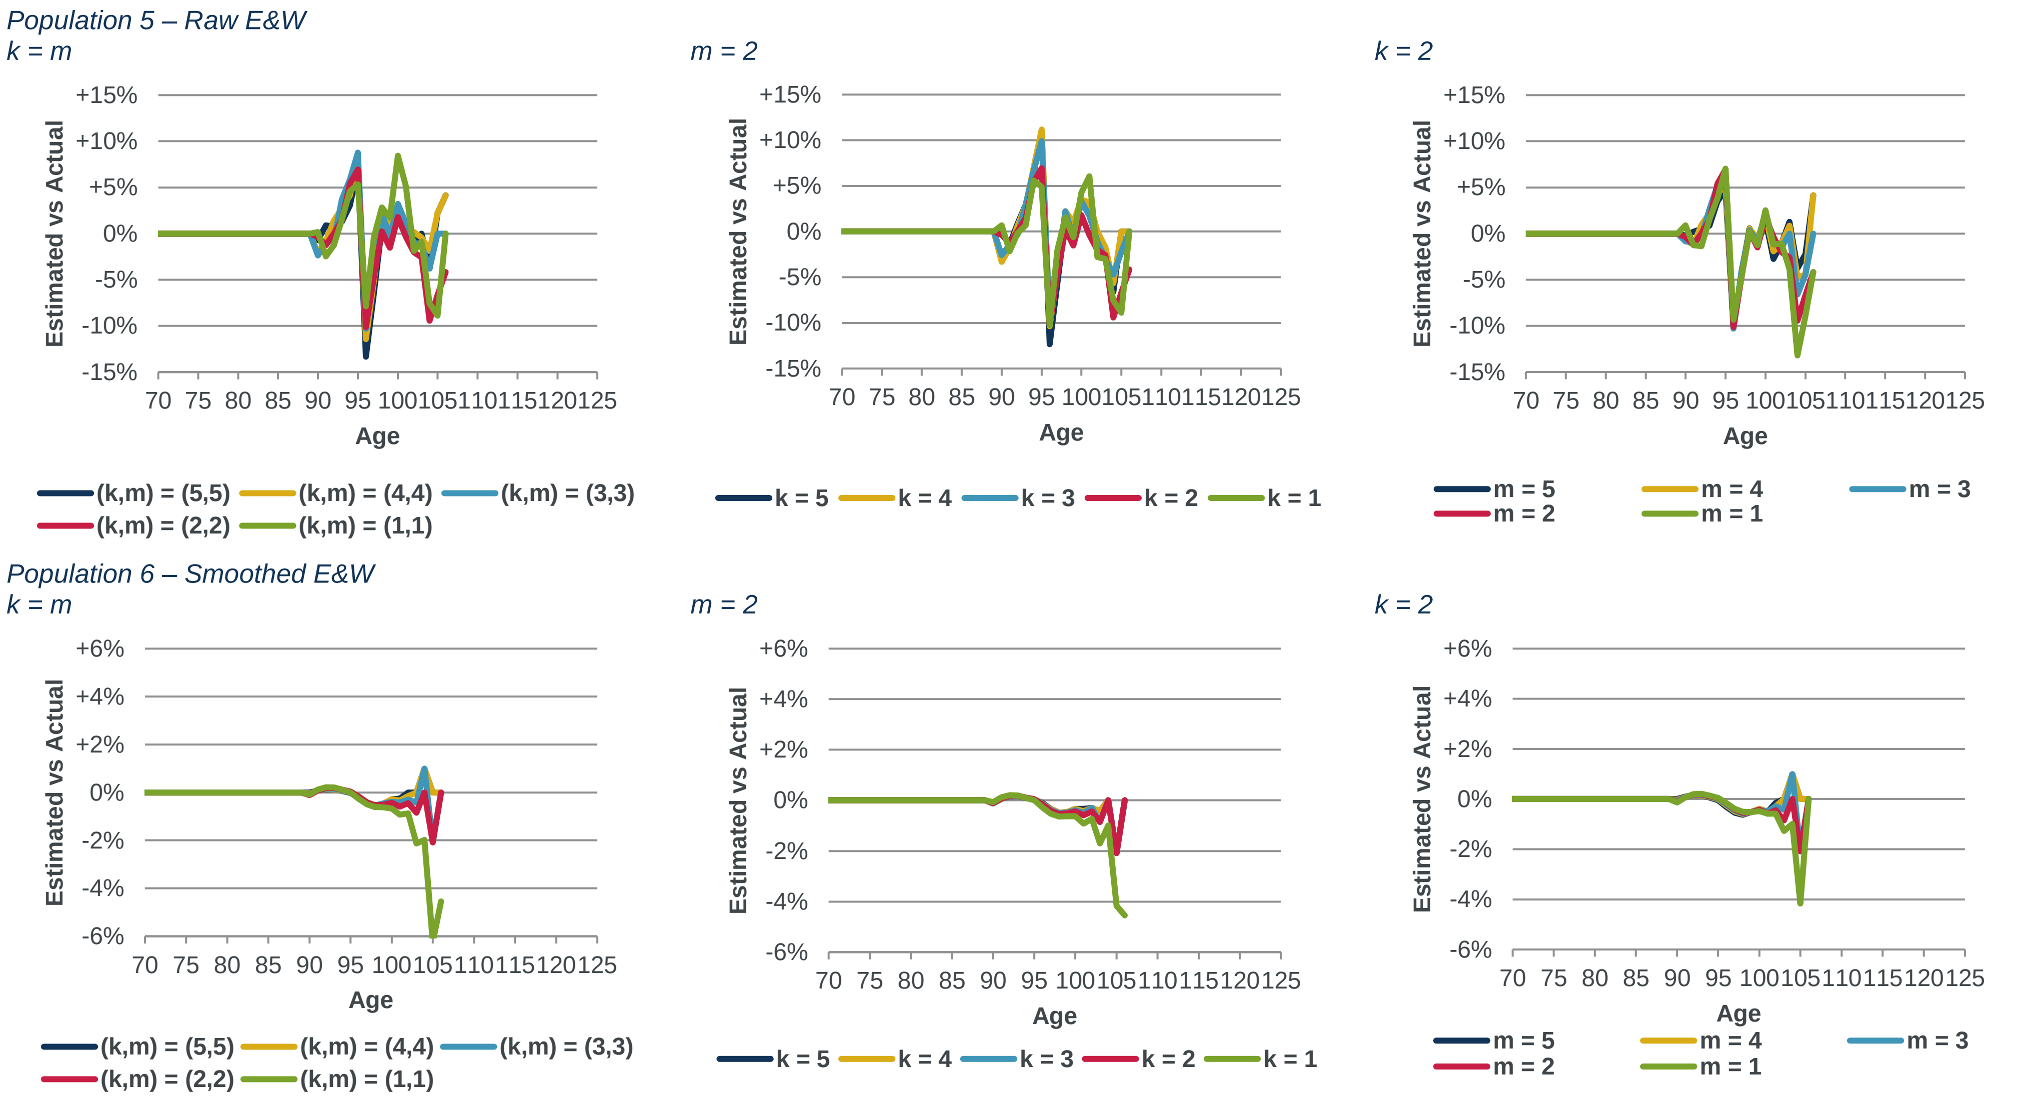

The synthetic population testing covers a broad range of scenarios intended to draw out projection robustness in different circumstances. Whilst it has provided *some* indication that reducing m to (say) 2 may lead to better performance in certain scenarios (and comfort that it does not materially underperform in other scenarios) the picture is far from conclusive. In particular, the scenarios in Chart 7.1 which provide strongest support for reducing m (populations 1 and 2) represent artificially smooth improvement regimes, and scenarios with more 'realistic' noisy patterns of improvements (e.g. population 5) are more ambivalent. (Indeed, the right-hand chart for population 5 suggests that reducing m below 4 or 5 may lead to *less* accurate high age population estimates – though of course this chart is based on k = 2 and so may be more prone to data instability for low values of m than when using k = 5.)

Given that there is already an established parameterisation (k,m) = (5,5) in the literature (albeit suggested in the context of the KT method *without* trend, and with the qualification that there is 'no clear optimal combination') we believe that further evidence is needed before a departure can be justified. With this in mind, we now turn to review the performance of different k and m values when applied to the actual ONS data for E&W.

ONS male data – final year balancing adjustments and population stability

Chart 7.2 shows the final year balancing adjustment in 2015 for a range of k and m values, again under the trend variant of the KT method.

Chart 7.2: Final year balancing adjustments in 2015 by (k,m) for ONS male data (KT with trend, join age = 90)

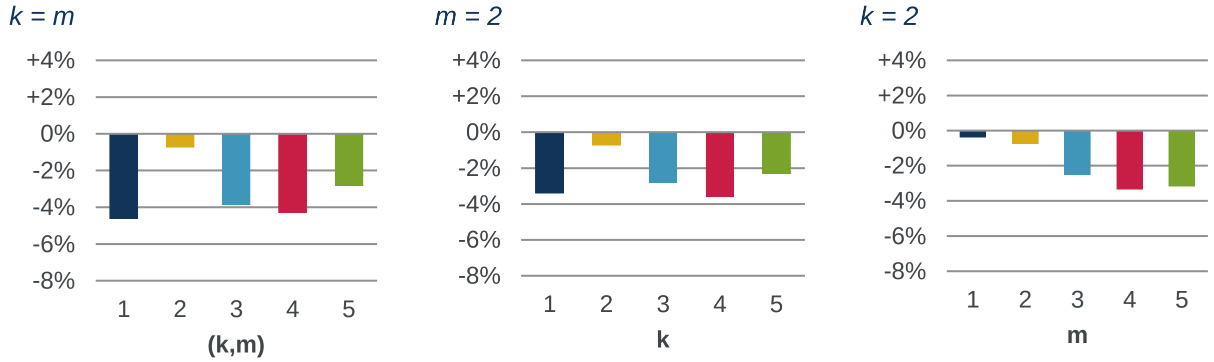

The charts again indicate that

- a reduction in the size of the survivor ratio window may deliver a better fit, although a very small window (k,m) = (1,1) can lead to deterioration (left-hand chart);
- in particular, lower values of m produce tighter final year balancing adjustments (right-hand chart);
- the case for reducing k is less clear-cut, and values of k that are too small (e.g. k = 1) may lead to instability.

The main concern with reducing either k or m is that it might undermine the stability of the KT method by increasing the level of random data noise captured in the survivor ratios. This is of particular relevance to the trend variant, which extrapolates the pattern of recent survivor ratios at each age from a smaller volume of local past data than the base variant (with no trend allowance) currently used by the ONS.

To provide some comfort on this, we have performed further analysis (contained in the Excel Appendix):

- We have assessed the impact on E&W male population estimates of adopting different k and m values, under the trend variant of the KT method, and have concluded that the pattern of population estimates by age under the trend variant is relatively robust to values of m and k, provided that the survivor ratio window is not too small.
- We have analysed the progression of survivor ratios in the male ONS data since the year 2000 over time. We found that reducing the size of the survivor ratio window from (k,m) = (5,5) to (k,m) = (5,2) does not materially destabilise the survivor ratio projection based on a 5-year trend period, but further reducing to (k,m) = (2,2) may lead to greater volatility, particularly at the highest ages.

- We have reviewed the impact on population estimates that would result from reducing m to 2. We observe that the variant with (k,m) = (5,2) delivers stable population estimates towards lower ages as the constraining total ranges from 90+ down to 75+.

Conclusion and proposal

Our testing suggests that reducing the number of cohorts m in the average survivor ratio of the KT method from 5 to (say) 2 may deliver a more robust projection. This makes intuitive sense in the context of the trend variant, which already looks back over the previous 5 cohorts to fit a trend to the survivor ratios. This approach may mean that including more than a couple of cohorts in the survivor ratios to extrapolate is both unnecessary from a data volume perspective, and unhelpful insofar as it may obscure actual trend patterns by averaging too heavily.

A possible further reason for the improved performance of m = 2 is simply that it shortens the survivor ratio window such that:

- the problem cohorts of 1919/20 are given less weight in the projection; and
- the period over which the trend in survivor ratios is determined falls predominantly after 2010 (when the relatively higher mortality improvement trend of the previous decade transitioned to the lower trend of the past 5 years) rather than straddling the two different improvement regimes either side of that year.

To the extent that this is a factor, it could be considered an accident of the particular period tested by our analysis rather than indicative of any general outperformance of lower m values. However, we suspect it is not the main driver of the improved performance, because the specifics of the historic improvements and the anomalous cohort features of 1919/20 which appear in the ONS data are not properties shared by the full range of synthetic population datasets we have tested (and these also provide support for a reduction in m below 5).

Based on this analysis we suggest that adopting parameters (k,m) = (5,2) alongside the trend variant of the KT method should be worthy of consideration.

Note that this suggestion is targeted principally at reliable estimation of the *E&W* population exposures, and has been tested in this light. The male and female ONS datasets and the synthetic population datasets used as the basis of our analysis in this paper all reflect E&W population data volumes. It may be that this choice of parameters performs less well for alternative (e.g. smaller) populations, if the data credibility leads to a different trade-off between stability and resolution.

We estimate that the impact of adopting both changes at once, though likely to be more robust across a range of scenarios, would actually produce an immaterial effect on period life expectancy compared with the current ONS method (based on the ungraduated average mortality rates for 2011-2015 inclusive). This is largely due to the tight constraint imposed by a join age of 90 under each method.

8. Join age

Motivation

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. Below age 90, the census-based population figures are adopted without adjustment.

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. The ONS provided a comprehensive overview in its own December 2016 paper, noting that:

- a range of high age data recording and processing issues had been identified in the 2011 Census, and previous censuses over the last 4 decades;
- the 2001 Census count for ages 90+ was 15% (males) / 8% (females) higher than expected based on independent research from the ONS Longitudinal Study (ONS-LS); and
- the original 90+ population estimate for 2011 (rolled forward from 2001 census data) was 7% overstated compared to the 2011 Census – an accumulating error of 0.7% per annum between successive decennial censuses.

Furthermore, 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', for which the 90+ population total is:

- potentially 0.3% overstated by recording discrepancies in year of birth (as compared with modal data values from a set of independent sources, including previous censuses and NHS Digital's MIDAS[2]);
- 2% higher than suggested by the sampling distribution of the ONS-LS dataset (though there were other possible explanations); and
- 3-5% higher than implied by administrative data from the Statistical Population Dataset version 1.0 (SPD V1.0)[3].

The 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.

Finally, our previous analysis in Working Paper 85 indicated that the 90+ population totals for past years of data (for which all cohorts are now extinct) were overstated relative to estimates implied by the deaths data.

All of this suggests that constraining to the 90+ census-based figures may be insufficiently robust in the determination of high age population exposures. The ONS has queried whether age 90 really is the 'optimum' age boundary for use in the KT method, or whether an alternative may be better suited.

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

Although a reduction in join age should lead to more robust constraints, there are two potential concerns worth noting here.

Firstly, the KT method assumes no migration, and migration rates tend to increase towards younger ages.

- This makes the KT methodology less robust, and could (in principle) offset the gain from constraining to a more reliable total.

[2] MIDAS is the 'Medical Research Information Service (MRIS) Integrated Database and Administrative System'. It was introduced during 2011/12 to improve automatic tracing of records to link and manage study cohorts for researchers.
[3] We understand that the SPDs are research outputs and subject to ongoing development. They are not official estimates (and so the 3-5% figure quoted here may be overstating any true error in the 90+ totals).

- In practice, however, the migration figures for E&W are immaterial well below age 90, so this is unlikely to cause a problem. For example, annual net migration at ages 80+ between 2012-2014 amounted to only around 0.1% of the total 80+ population.

Secondly, the weight placed on the survivor ratio assumptions within the KT method increases for cohorts which are furthest from extinction.

- This makes the estimates towards younger ages in the final years most susceptible to error, as demonstrated in the ONS's recent analysis of Swedish and Finnish data.
- Using a lower join age extends the scope of the KT method into younger ages, which is precisely where it is least robust.
- This should be ameliorated to some extent by the application of the total population constraint, but is nonetheless an important limitation to consider.

Results

For a specific join age, the synthetic population testing can show the potential for errors in the estimates where the actual population is known (for example, see Chart 6.3 in Section 6). But the synthetic population testing is agnostic in relation to making comparisons between join ages. This is because, for the synthetic populations, we know that the high age figure to which we are constraining is 'correct'. It is therefore necessarily the case that using a higher join age in applying this constraint must lead to a tighter fit of the population estimates to the actual synthetic population.

This may look like better performance in the diagnostics, but it tells us nothing about the most appropriate join age to use when constraining to the ONS population data (for which the principal motivation for analysing join age is concern over the accuracy of the official figures).

Therefore in this section we consider the diagnostic results as applied to the ONS data only.

ONS male data – consistency diagnostics

Chart 8.1 shows the final year balancing adjustment and average annual scaling adjustment under four different join ages (75, 80, 85 and 90), for the trend variant of the KT method using (k,m) = (5,2), applied to ONS male data. **The final year balancing adjustment is relatively insensitive to the choice of join age, but the average annual scaling adjustment reduces markedly for lower join ages, demonstrating a closer fit to the data.**

This second diagnostic includes historic years for which almost all high age cohorts are now *extinct*, which means that the lower average annual scaling adjustment for join ages 85, 80 and 75 implies closer agreement between the historic 85+, 80+ and 75+ totals and a reconstruction of the corresponding population figures based on the method of extinct generations.

Chart 8.2 expands on this, by plotting the annual scaling adjustment over time (right-hand chart) alongside the proportion of the population above the join age which is extinct by 2015 (left-hand chart), for a join age of 90 and 85 respectively.

The left-hand chart demonstrates that, for years prior to around 2005, almost the entire population above the join age comprises cohorts that are now extinct. This means that the KT estimate of the population for prior years is effectively based on a pure extinct generations construction.

The right-hand chart shows that, over the same historic period (prior to 2005), the size of the annual scaling adjustments is systemically lower for a join age of 85 than for a join age of 90.

This tells us that the official ONS population total for ages 85+ in historic years is consistently in better agreement with an extinct generations estimate than is the ONS population total for ages 90+. In other words, it suggests that the 85+ total may be more reliable.

As an aside, we note that the annual scaling adjustments in the right-hand chart reduce towards zero as one approaches the final year in the KT method implementation. This makes sense, because the final year balancing factor is already serving as a constraint to the official ONS population total in the final year (which means that no further scaling is needed in that year, and the degree of adjustment required for immediately prior years will also be small).

Chart 8.3 (left-hand chart) shows the average cohort inconsistency metric, for join ages 75, 80, 85 and 90. This indicates an improvement in cohort consistency with reducing join age. Chart 8.3 shows the average mortality deviance across the join age, for join ages 75, 80, 85 and 90. This suggests that the smoothest join to younger ages is for a join age around 80.

Chart 8.1: Constraint adjustments by join age for ONS male data (KT with trend, (k,m) = (5,2))

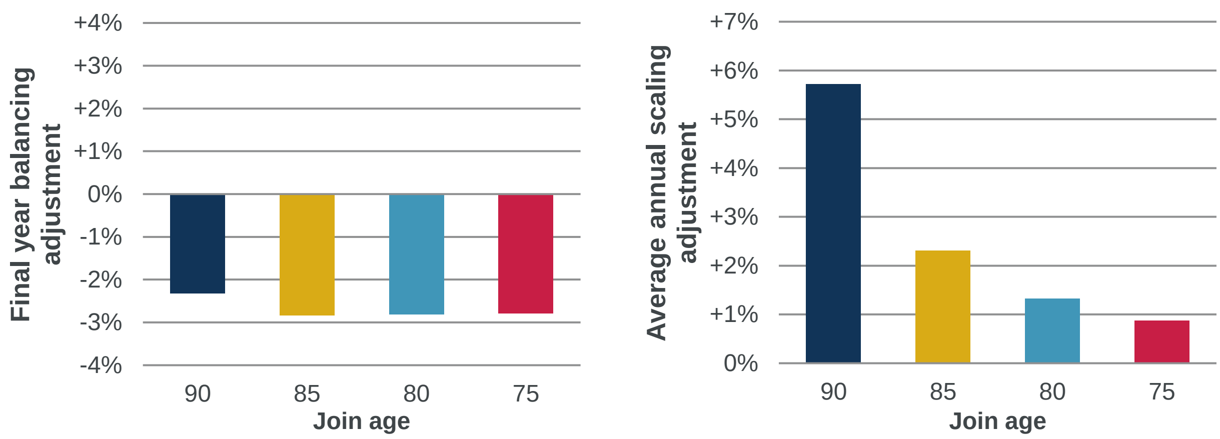

Chart 8.2: Annual scaling adjustments over time for ONS male data with join age = 85 and 90

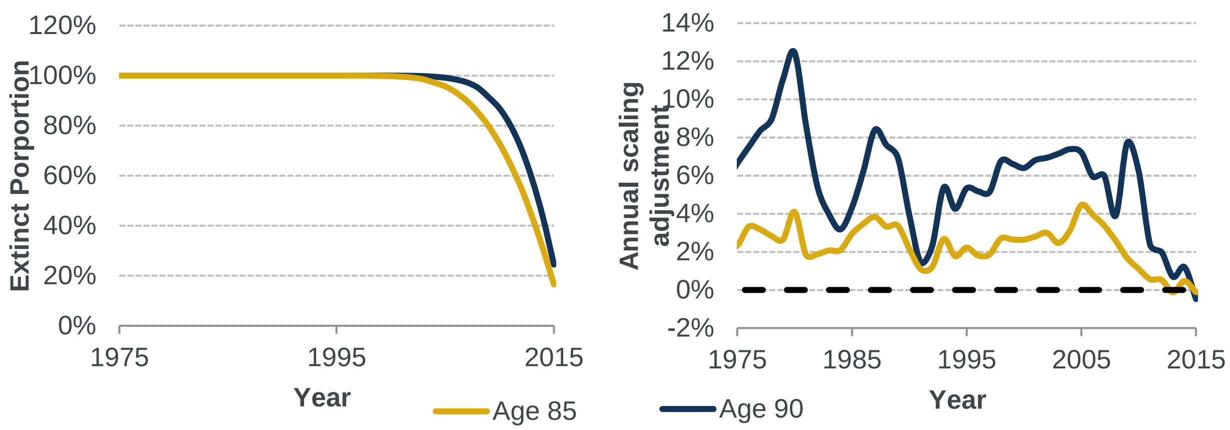

Chart 8.3: Average cohort inconsistency and average mortality deviance by join age for ONS male data (KT with trend, (k,m) = (5,2))

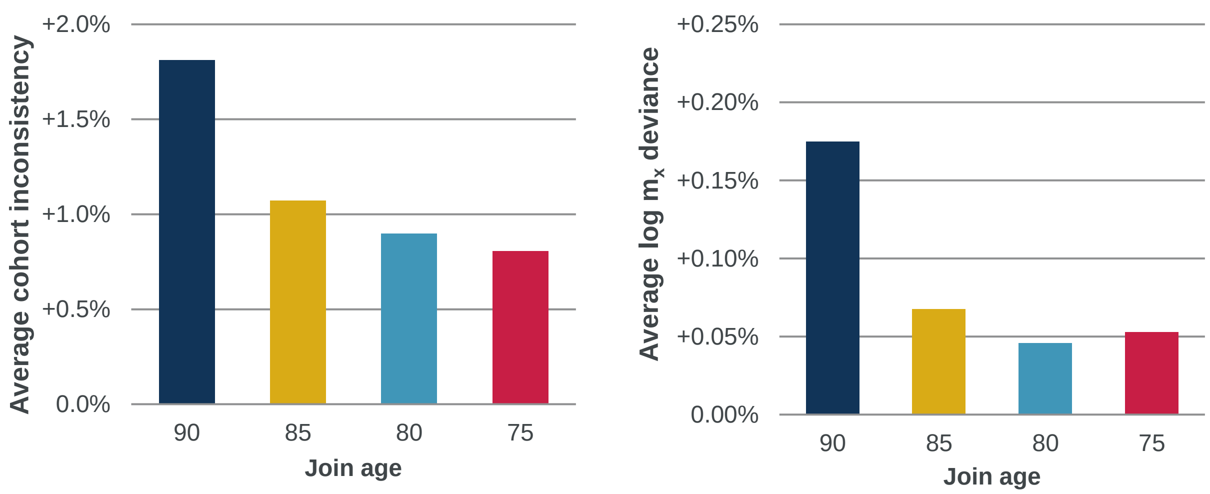

ONS male data – other considerations

Chart 8.4 shows the CBDK 2 concavity heatmap diagnostic for different join ages, under the trend variant of the KT method. It appears that some of the anomalies in the official population data are ameliorated if the join age is reduced, so extending the KT method down into the problem region.

- An example is the diagonal cohort effect in the top-left corner of the heatmaps (circled). This area comprises fully extinct cohorts, so the extension of the KT method into the region places no reliance on the survivor ratio projection in this case. The uncorrected high age anomalies in this region of the dataset are of course a known issue and one of the reasons that the CMI Mortality Projections Committee now excludes pre-1976 data in the calibration of the CMI Mortality Projections Model.
- Another example is the discontinuity at the join age which appears as a horizontal line when a join age of 90 is adopted (circled in the bottom-right heatmap). The join is visually smoother in the other heatmaps, i.e. when reducing join age below 90.

The impact of reducing the join age from 90 to (say) 80 would be two-fold:

- Firstly, the population estimates between 80 and 89 would be altered compared with the current approach. The extent of this is small in percentage terms, because:
  - At an aggregate level, the 80-89 population total makes up the vast majority of the 80+ population total, which means that the effect of constraining the 80+ population to its official figure will be similar to constraining the 80-89 population to its official figure.
  - The shape of the population by age appears relatively stable under the KT method (even in the final year 2015, for which the greatest dependence is placed on the projection of survivor ratios). We have of course only shown these population impacts for one year, and the picture could be different when testing the KT method on alternative end years.
- Secondly, the population estimates for ages 90 and above would be altered, because a different population constraint is being applied to the 90+ constraint under the current approach.
  - We estimate that the 2015 population estimate for ages 90+ would reduce by around 1.0%-1.5% if the join age were reduced from 90 to any of 75, 80 or 85 (assessed under the trend variant of the KT method with $k = 5$ and $m = 2$).
  - A reduction of this size in the 90+ population total seems reasonable given the evidence summarised earlier in this paper that the official 90+ total may currently be overstated by a similar order of magnitude.

Chart 8.4: CBDK 2 concavity heatmaps by join age (trend variant) for ONS male data

Join age 75

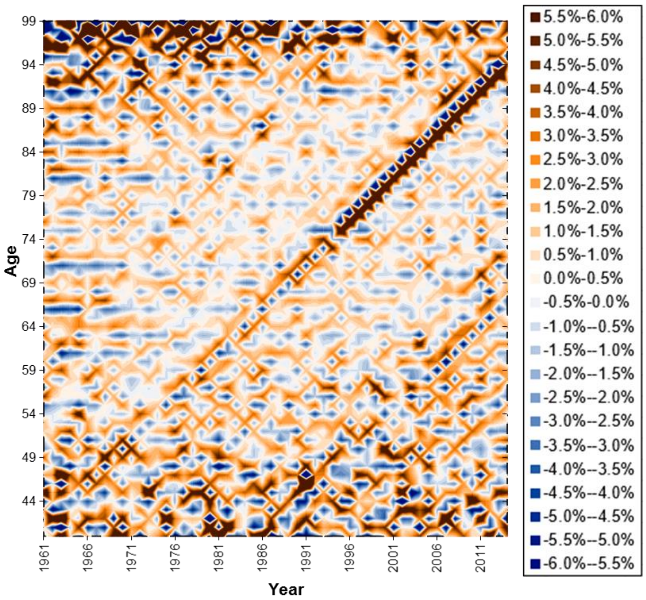

Join age 80

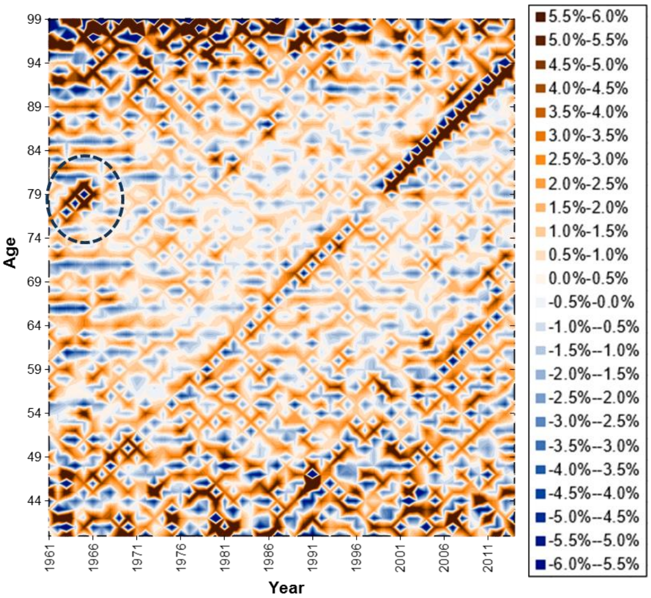

Join age 85

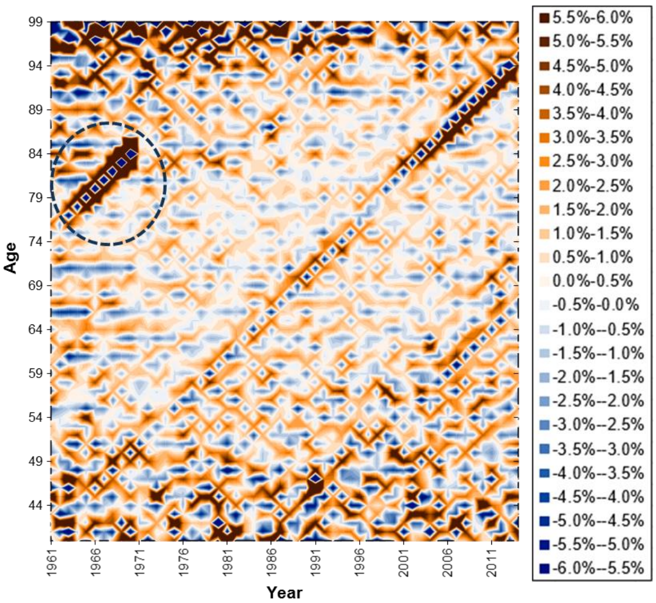

Join age 90

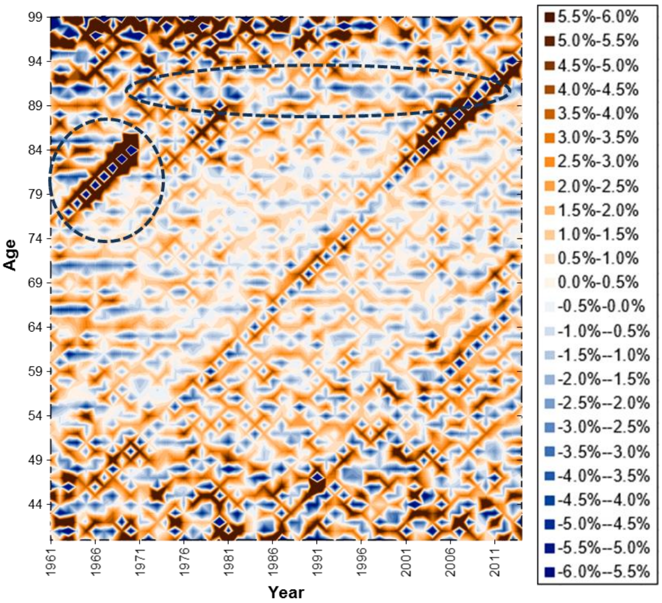

Conclusion and proposal

Our diagnostics suggest that reducing join age from 90 to 80 may deliver population estimates with improved internal consistency, and may ameliorate certain data anomalies.

We have prior reasons to doubt the reliability of the 90+ official population total to which the KT estimates are currently constrained, and in contrast our average annual scaling diagnostic suggests that the 80+ total population figure from historic years may be a good fit to the figures implied by extinct generation death records. Furthermore, we note that other bodies (such as the HMD) already extend the KT method down to lower join ages in some regions of the data.

Moving to a join age of 80 appears to reduce the 90+ population total by a margin that is reasonable given independent evidence, and the younger 80-89 population total appears to be robust under the trend variant of the KT method with (k,m) = (5,2) (not least because it is stabilised by the application of the 80+ population constraint). This latter point allays some concern that reducing the join age might extend the KT method into a younger age region where it is inherently less robust (due to greater distance from extinction and hence increased dependence on the survivor ratio projection). This would be the case under an unconstrained version of the KT method, but the large population size for ages 80-89 compared with 90+ provides stability when constraining the method.

Furthermore, there is strong evidence from recent ONS research that migration remains immaterial down to age 80, and so should not invalidate the assumptions of the KT method.

Notwithstanding the above, we are conscious that extending the KT method down as far as age 80 would place heavy reliance on the extrapolation of survivor ratios into younger ages. Regardless of the assumptions for migration, 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 our estimates.

From a pragmatic perspective, we note that the impact on high age population totals for 2015 appears to be broadly similar whether one adopts a join age of 75, 80 or 85 instead of 90. This may in itself reflect the increasing reliability of the official ONS population figures below age 85 (exhibiting less divergence from the KT estimates). In practice, it means that there may be little to gain from adopting a join age lower than 85, with the greater risks to the reliability of the KT method which that entails.

Bearing all this in mind, we propose a reduction in the join age from 90 to 85. It may improve the robustness of the KT method and avoid undue dependence on the 90+ population total currently used as a constraint.

The effect of making this change for calendar year 2015 (having already adopted the trend variant with k = 5 and m = 2) would be to increase implied mortality rates and so reduce life expectancy. We estimate that moving to a join age of 85 would reduce male period life expectancy by around 0.2% at age 85, 1.3% at age 90 and 1.2% at age 95 (based on the ungraduated average mortality rates for 2011-2015 inclusive).

9. Lexis adjustments to deaths data

Motivation

The KT method requires death counts over a 12-month observation period (typically a calendar year) by age at the start of that period, i.e. at 1 January. 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[4] in Chart 9.1.

Chart 9.1: Lexis diagram of deaths data timing

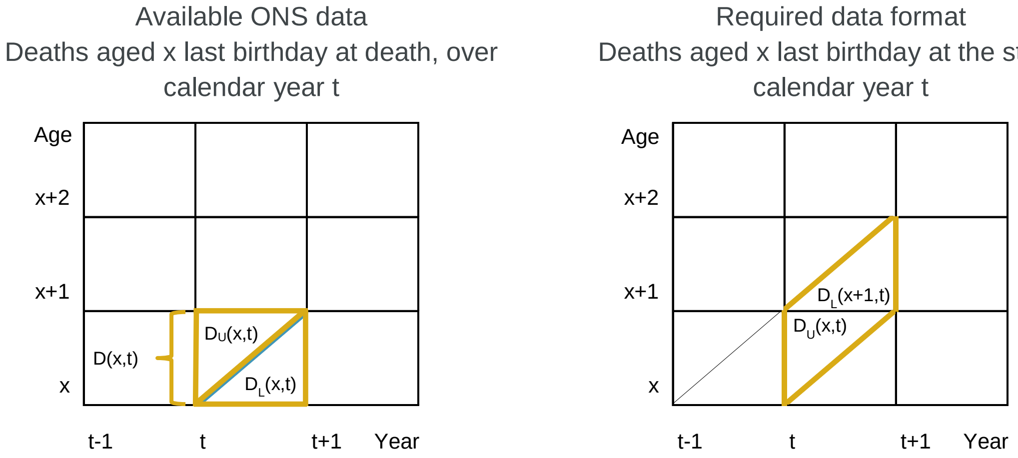

The deaths aged x at death in calendar year t, **D(x,t)**, can be split into two Lexis triangles (left-hand grid). The upper triangle, **D_U(x,t),** represents deaths of individuals aged x at the start of calendar year t, who died before reaching their birthday in the year of death. The lower triangle,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

This 50/50 assumption is plausible for younger ages where the mortality curve is relatively flat and shallow, but quickly becomes questionable at higher ages when the mortality curve gets steeper.

We have therefore tested a variant which aims to allocate deaths to the appropriate age triangles using a more accurate Lexis adjustment.

The following refinements have been included:

- 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.

[4] The Lexis diagram takes its name from the German statistician Wilhelm Lexis, who introduced it in his 'Introduction to the Theory of Population Statistics' in 1875 (though other authors were developing this idea around the same time).

- 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.

For each age x and calendar year t, we then estimate the proportion of deaths in the upper triangle by:

1. Constructing the monthly exposure run-off grid by age x-1 to x+1 and calendar year t to t+1;
2. Applying relevant monthly mortality rates to the monthly exposures above to construct the monthly deaths by age and calendar year;
3. Summing over deaths aged x in calendar year t to give us D(x,t);
4. Summing over the upper triangle in D(x,t), i.e. deaths aged x at the beginning of calendar year t, to give us $D_U(x,t)$;
5. Determining the proportion of deaths in the upper triangle as $D_U(x,t) / D(x,t)$.

We repeat the process to populate a table for all ages and calendar years.

The monthly birth pattern is taken from ONS data on live births by year and month. We have assumed ELT17 to be a suitable base mortality curve to give us the relative mortality differential by age, and CMI_2015 Core improvements to allow for changes in mortality rates over time. Furthermore, we have assumed net migration to have nil impact on monthly exposure patterns.

Results

Table 9.1 shows the impact on period life expectancy of adopting our more sophisticated approach to Lexis triangle decomposition of the deaths data (when converting to a 1 January age definition). A full grid of Lexis adjustments is shown in Appendix 3 for males (and provide in the Excel Appendix for both males and females). The impacts are expressed relative to our proposal so far, i.e. the trend variant of the KT method with (k,m) = (5,2) and join age 85. We have also included a comparison based on join age 80 for completeness – this shows similar impacts.

Introducing the Lexis adjustments causes a material reduction in life expectancy towards higher ages.

The principal reason for this is that the lower triangle of each death cell carries more exposure under the survival profile of the cell than the upper triangle (by virtue of being 'younger'). This means that at higher ages increasingly more weight is given to the lower triangle in the decomposition. This in turn reduces the 1 January death counts provided as inputs to the KT method, and hence reduces the constructed population estimates (leading to higher mortality rates and so lower life expectancies).

Table 9.1: Period life expectancy impact of refined Lexis approach (ONS males) on mortality rates for 2011-2015, KT with trend, (k, m) = (5,2)

Age	Join age 85	Join age 80
Under 80	nil	nil
85	−0.1%	−0.2%
90	−0.3%	−0.3%
95	−1.1%	−1.1%

In contrast the impact of the Lexis adjustments on the final year balancing adjustment and average annual scaling adjustment is small – see Table 9.2.

Table 9.2: Constraint adjustment impact of refined Lexis approach (ONS males) on KT with trend, (k,m) = (5,2)

	Join age 85		Join age 80	
	50/50 Lexis	Refined Lexis	50/50 Lexis	Refined Lexis
Final year balancing adjustment	−2.8%	−2.8%	−2.8%	−2.8%
Average annual scaling adjustment	+2.3%	+2.6%	+1.3%	+1.4%

We have also reviewed the impact of the Lexis adjustments on the internal consistency of the population distributions, as exhibited in the CBDK diagnostic plots.

Chart 9.2 shows one example, comparing the pattern of mortality rates by cohort for the 50/50 and refined Lexis approaches applied to birth cohorts between 1917 and 1921 (where distributional anomalies have been identified in the population). In each case the KT method adopts our proposed changes so far, i.e. introduction of the trend allowance, (k,m) = (5,2) and join age 85.

Chart 9.2: CBDK 1 – mortality rates by cohort (1917 to 1921 cohorts) for ONS male data

With 50/50 Lexis decomposition

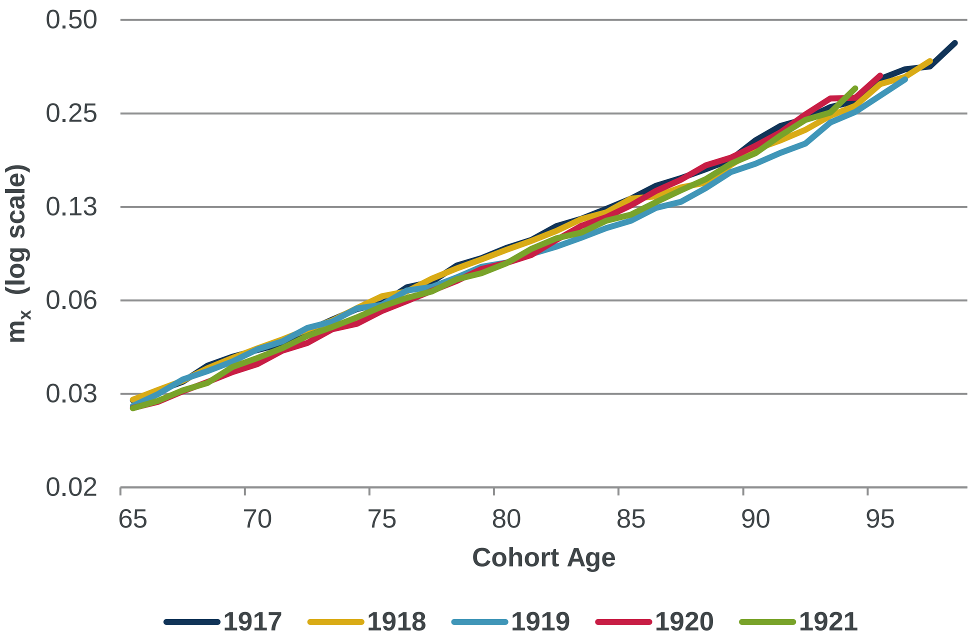

With refined Lexis adjustments

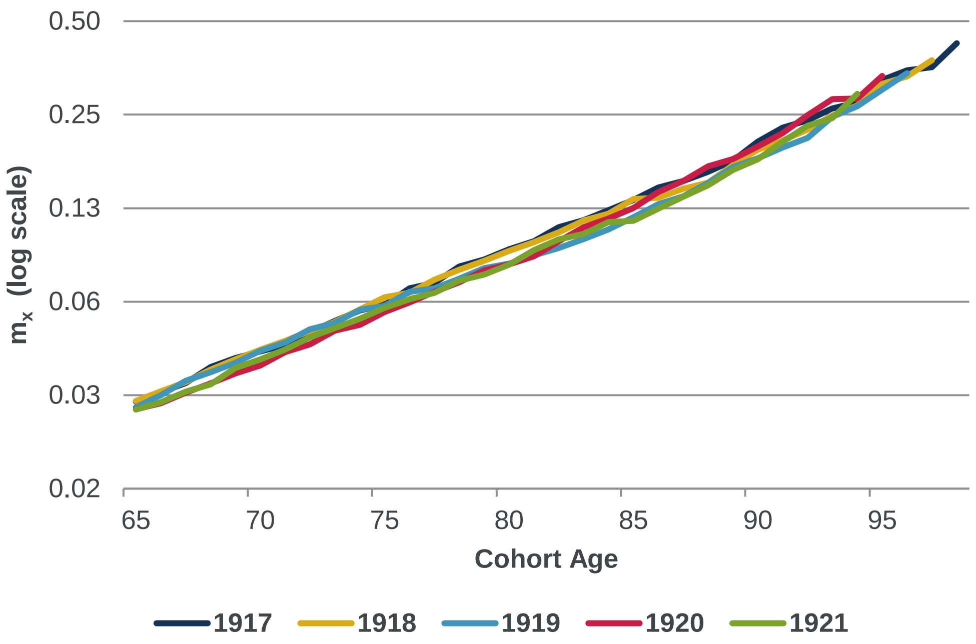

The second chart (with refined Lexis adjustments) shows a more log-linear cohort mortality curve. In particular, the light blue line representing the 1919 cohort is less divergent from its neighbours towards the highest ages.

Conclusion and proposal

Chart 9.2 shows an improvement in consistency between the mortality curves for adjacent cohorts in the high age region. Our overall impression is that incorporating the refined Lexis adjustments has scope to produce more coherent mortality estimates, with a material impact on life expectancies at older ages.

Furthermore, the principle of applying Lexis adjustments is theoretically more accurate – and a well-established practice already in use by other bodies such as the HMD. Indeed, the HMD describe it as 'one of the most important steps in computing the mortality rates and life tables'.

The ONS has previously considered this area (based on Finnish data) and found that the distinction in age definitions for input deaths data to the KT method 'does have some effect on the quality of the resulting KT estimates but this is relatively minor.' However, this assessment appears to be based only on aggregate 90+ population totals, and the conclusion may have been different if analysing the impact by single year of age. We note that the direction of movement in the 90+ totals from the ONS's analysis is consistent with our own findings in this paper.

Based on our analysis we suggest that there is merit in making allowance for the more sophisticated Lexis adjustments within the KT method.

Note that the precise impact of adopting a Lexis adjustment approach will vary depending on the specific distributional adjustments made (which are complicated and incorporate a range of assumptions, for example the base mortality curve used to run down exposures for survival). We believe that the set of adjustments used in this paper is reasonable, but it is by no means the only set that could be adopted.

10. Exposure adjustments

Motivation

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 (see Appendix 4) they also proposed specific adjustments that could be made to resolve them, including

- adjustments to allow for the actual distribution of births in the 2001 Census to mid-year roll-forward (which affects population estimates for 1992 to 2010 due to the method used by the ONS to back-fill population estimates between census years) and
- Convexity Adjustment Ratios (CARs) to estimate the exposed-to-risk over the course of a year from the mid-year population figure (rather than simply assuming that the two are equal).

They also conducted an overarching Bayesian analysis, which identified adjustments to make the exposures data more internally consistent. The smoothing adjustments derived from this Bayesian approach capture both effects above as well as other, unknown, effects (for example, the impact of inconsistencies that could arise around the join age due to transitioning between the KT method and census-based figures at that age). There may of course be further anomalies not captured by this framework.

We are grateful to Andrew Cairns for sharing information on these adjustments with us, and for running some of our test scenarios through the Bayesian framework.

Conscious that ease of communication and implementation is a desired feature of our modelling (see Section 4.1), we have also considered an alternative (pragmatic) approach to smoothing anomalies in the final exposure estimates.

Under this approach, we follow the exposure adjustment method set out in detail in CMI Working Paper 91. To summarise, for a given age x and year t, we

- work with crude $log\, m_{x,t}$ (where $m_{x,t}$ is the central mortality rate for age x and year t);
- smooth by fitting a straight line to the data for ages $x - 2$ to $x + 2$;
- if the crude and smoothed rates for age x are far apart, we adjust the exposure to match the smoothed mortality rate;
- we assess 'far apart' using deviance residuals.

Chart 10.1 depicts this approach as applied to an illustrative exposure estimate for age 91 in 2011. In this example, the exposure would be adjusted for age 91 (and for age 92).

Note that the nature of this adjustment is a little different to the other KT method variants tested in this paper: The variants considered so far serve only to change the age distribution of the population above the join age (rather than the total population, which is still being constrained back to the official total at the end of the estimation process).

In contrast, this exposure smoothing step applies after the constraint to official data and so has the potential to introduce an overall departure in the population total above (and indeed below) the join age from the official figures, if this serves to produce a smoother progression of mortality rates.

Chart 10.1: The CMI's exposure smoothing method (applied to age 91 in 2011)

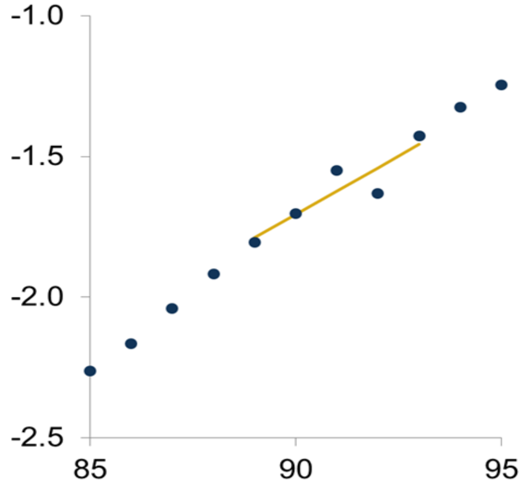

Results, conclusion and proposal

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 relate only to the conversion of the death input data and do not tackle distributional effects in the population when subsequently interpolating the 1 January population estimates from the KT method to mid-year estimates, and
- some of the issues reside below the join age (so are unaffected by the KT method in any case).

In the course of our investigation we liaised with Andrew Cairns to understand the impact of the Bayesian adjustment framework proposed in Cairns et al in addressing the residual anomalies. From testing the framework on a number of the scenarios in this paper, we conclude that it does appear to successfully resolve the majority of issues (leading to much smoother plots under the suite of CBDK diagnostics). However, the result of this sophisticated adjustment framework was relatively immaterial in terms of overall life expectancy figures and broadly indistinguishable from the simpler and more accessible approach to exposure smoothing set out in Working Paper 91.

There is an argument for adopting no exposure smoothing at all, based on overall materiality. However, we are concerned that this might overlook isolated cases – such as the 1919/20 birth cohort – where the materiality could be much higher, and hence risk making the ultimate set of population exposures less robust for wide-ranging use. Given that the CMI's smoothing adjustments are straightforward to implement, it feels safer to protect against this risk by incorporating them.

As a result, we prefer to adopt the CMI's exposure smoothing method as a pragmatic alternative which is easier to communicate and implement. In the figures that follow (Charts 10.2 and 10.3) we illustrate the effectiveness of this method in ameliorating the data anomalies identified by the CBDK diagnostic tests. Please refer to Appendix 4 for full details of these tests.

We estimate that incorporating the CMI exposure smoothing would increase male period life expectancy by around 0.1% at age 85, 0.4% at age 90 and 0.3% at age 95 compared to our proposed variant up to this point (based on the ungraduated average mortality rates for 2011-2015 inclusive).

Chart 10.2: CBDK 2 – concavity by cohort (1920 cohort) for ONS male data

Without CMI exposure smoothing (trend, k = 5, m = 2, join age = 85, Lexis adjustments)

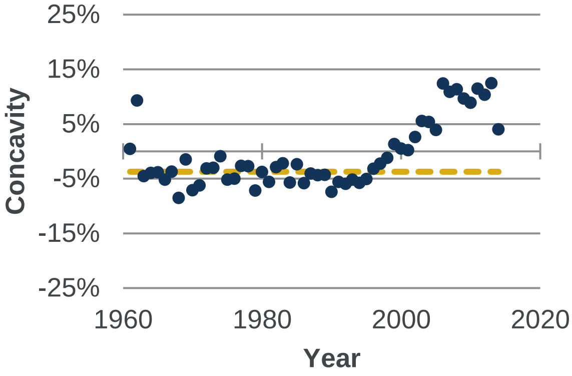

With CMI exposure smoothing from 1975 (trend, k = 5, m = 2, join age = 85, Lexis adjustments)

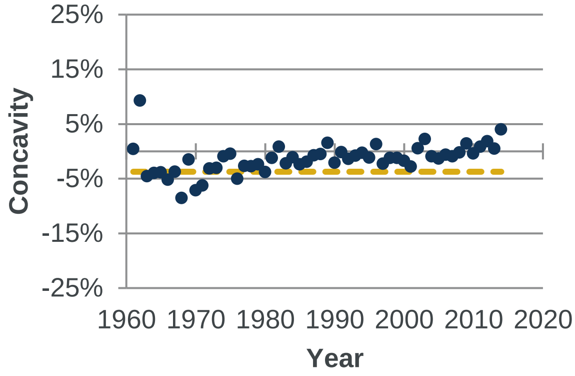

Chart 10.3: CBDK 2 – concavity heatmap for ONS male data

Without CMI exposure smoothing (trend, k = 5, m = 2, join age = 85, Lexis adjustments)

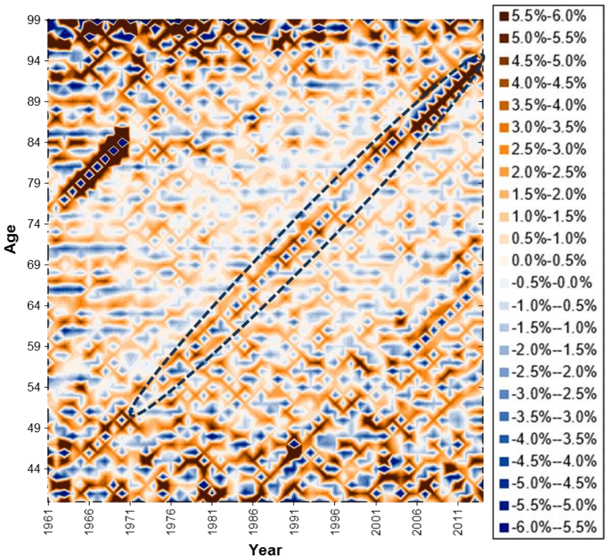

With CMI exposure smoothing from 1975 (trend, k = 5, m = 2, join age = 85, Lexis adjustments)

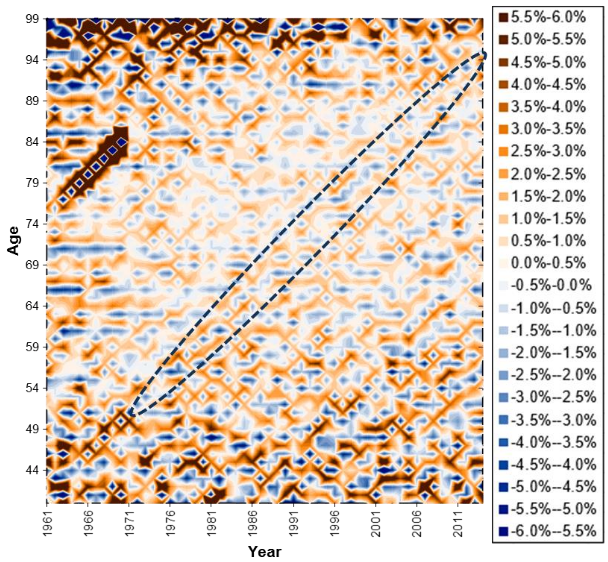

11. Constraining the KT estimates

Are the constraints accurate?

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. Below age 90, the census-based population figures are adopted without adjustment.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+.

There is an accumulation of evidence suggesting that the 90+ population figures to which the KT estimates are constrained may not be accurate. As summarised in Section 3, and our own work presented in Working Paper 85, a consistent picture is painted of official estimates of the total 90+ population being higher than estimates reconstructed from subsequent death records for extinct cohorts. Our analysis from Section 8 suggests that this overstatement may typically be of the order 2-8% (see Chart 11.1). It also indicates that the estimates for historic years show internal discontinuities around age 90, where the ONS transitions from census-based figures to the constrained KT method (recall Chart 8.4).

This issue is not confined to E&W. In his 2004 paper Kirill Andreev observes: 'Evidence available from existing demographic data indicates that size of a population at very old ages is rather vulnerable to age misreporting errors. Commonly, data errors act in such a way that population at the highest ages is inflated.' He cites the US as a particular example of this, with 'strong evidence of age overstatement in the oldest segment of the population in the 1980 Census.'

Chart 11.1: Official 90+ ONS population estimate vs population estimate reconstructed from deaths data, for cohorts expected to be extinct (or almost extinct) by 2015

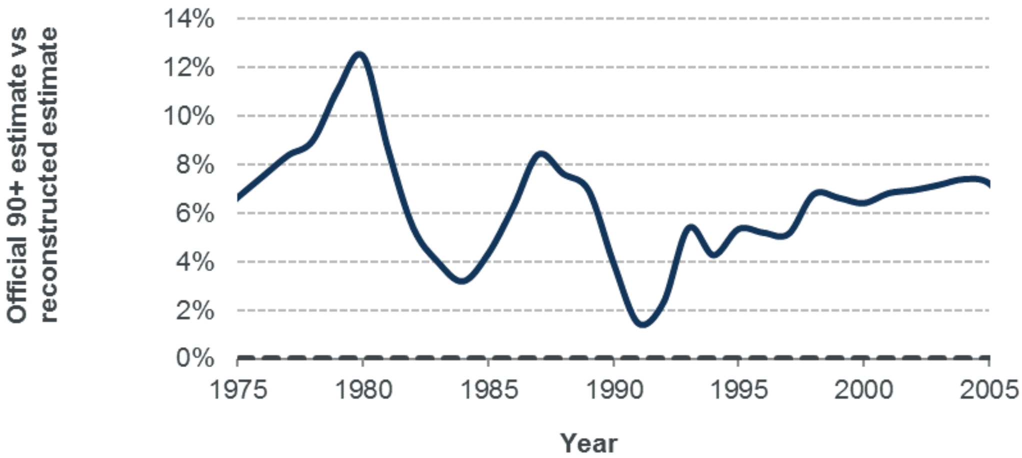

All of this suggests that constraining to the 90+ census-based figures may be insufficiently robust in the determination of high age population exposures.

We have considered whether to constrain at all, or whether it is more accurate simply to adopt the pure KT estimates at high ages without adjustment. **For estimating current or recent population figures (where the cohorts in question are not extinct) it is widely accepted that use of unconstrained KT estimates in the ONS's current format would not be robust. The principal reason for this is that the projection of survivor ratios in the method makes no allowance for mortality improvements over time.** This means that the method tends to underestimate population size at the highest ages, and there is independent evidence of this for both E&W and a range of other countries. Our proposals in this Working Paper paper aim to refine the KT method with direct incorporation of a mortality trend allowance, which should make the underlying projection more robust. Nonetheless, an attempt to apply the method with no constraint places significant reliance on the way in which survivor ratios are extrapolated – which is both relatively simplistic and dependent on noisy data.

The case for applying the KT method without constraint is different (and probably stronger) in relation to historic years, for which all individuals in the high age population are since believed to have died. In this case, the KT method amounts to a reconstruction of the former population based purely on subsequent death records. There is no reliance on the projection of survivor ratios in this scenario, and if the death records themselves are accurate then the unconstrained KT estimates should theoretically also be accurate. We note that the HMD effectively takes this approach in relation to past years – adopting the KT method at high ages but without scaling the resulting estimates to match official totals. We have some sympathy with this argument, and we would expect population figures reconstructed from the method of extinct generations to come close to the true population for historic years. **However, the adoption of unconstrained KT estimates in this form still presents a number of theoretical and practical issues:**

- **the KT estimates are reliant on distributional assumptions around the timing of deaths during each year, the absence of migration, and of course the quality of the death records themselves;**
- **a judgement is required on how historic the data needs to be before all high age cohorts can now be assumed fully extinct,** and on whether and how to adjust KT estimates for years which are not sufficiently historic;
- the KT method cannot be extended to arbitrarily low ages (in particular, because the assumption of nil migration is likely to become invalid). This means that there is an inevitable transition from official census-based figures at younger ages to KT-based estimates at higher ages. **There is a potential for discontinuity at this join age, and for inconsistency versus the total census-based population figures (i.e. for all ages) if the high age KT estimates are not constrained**.

The last point is important. Many of the concerns with the quality of census-based population totals at the highest ages relate to misrecording of information (such as date of birth), which serve to attribute individuals to incorrect age categories. There is no reason why this in itself should distort the overall population count for *all ages combined*, which we expect to be relatively accurate (especially in percentage terms, as the high age issues are diluted by the much larger cohorts alive at younger ages).

Bearing these points in mind, our preferred starting point is to retain the principle of constraining the KT estimates (for both recent and historic years) – but to investigate whether there are more accurate population totals to use for this purpose than the ONS's current 90+ figure.

In broad terms, we expect that towards lower ages, the ONS's official census-based population figures should become proportionately more accurate (as the population counts become larger and less susceptible to some of the specific census recording issues identified by the ONS), and the KT estimates should become proportionately less accurate (as one moves further into non-extinct cohorts, and migration starts to have an effect). There is likely to be a cross-over point at which the official figures become more reliable than the KT estimates.

This motivates consideration of constraining to a population total that extends into younger ages – for example, ages 85+ or 80+ instead of 90+. Such an approach would place less reliance on the official ONS total for ages 90+ (with which we and others have concerns over accuracy), reducing the constraint on the KT method at the highest ages whilst still ensuring consistency with the ONS figures published for the population as a whole.

Our analysis from Section 8 suggests that reducing the join age to 85 after allowing for mortality trend improves performance across a range of diagnostics, including agreement with extinct cohort estimates (Chart 8.2) and the internal consistency of mortality rates across the join age (Chart 8.3).

12. Conclusions

12.1 Proposal

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 ratio from 5 to 2;
- Extend the high age population method down to a lower join age than 90 to avoid placing undue reliance on the underlying census-based estimates in this age range – we propose age 85;
- 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 calendar year) 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.

12.2 Impact diagnostics - summary

If 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 (potentially with modest increases in life expectancy at some ages).

Table 12.1 summarises the impact at ages 85, 90 and 95 based on

- the average (ungraduated) mortality rates by age over calendar years 2011-2015, and
- the graduated mortality rates for the same period.

We show ungraduated life expectancy impact as a diagnostic of the underlying changes to mortality rates by age, before any smoothing. The graduated figures are provided as a guide to the potential impact in practice. For each measure, the impact on life expectancy at ages 80 and below is very minor (0.1% or less).

Note that the impacts shown are purely an illustration based on the particular period 2011-2015. We have seen different impacts for mortality rates based on alternative periods of data (for example, Appendix 1 covers a sensitivity for the period 2006-2010).

Table 12.1: Impact of proposal on period life expectancy from ages 85, 90 and 95, ONS E&W 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 figures in Table 12.1 are all measures of *base mortality* impact, i.e. the effect of our proposal on current/recent mortality rates. We have also reviewed the impact on projections of *future mortality improvement*, by calibrating the CMI's latest Mortality Projections Model, CMI_2016, to ONS E&W data prepared using the current vs proposed variants of the KT method, and comparing the resulting cohort life expectancies. Our proposal has a more modest impact on these cohort life expectancies, and in the opposite direction to the impact on period life expectancies, with increases of between 0.2% and 0.4% observed at ages 85, 90 and 95.

12.3 Impact diagnostics – male base mortality

Chart 12.1 illustrates the step-wise impact on male period life expectancy at age 95 from each element of our proposal, and Chart 12.2 shows the underlying population estimates (for 2015). We have chosen age 95 for the illustration in Chart 12.1 simply because it is sufficiently high to demonstrate clear sizeable effects from the impact of our proposal. The effect on life expectancy at younger ages is smaller, and covered later (in Chart 12.9) for completeness.

The two main drivers of the change for males are the reduction in join age from 90 to 85, and the use of refined Lexis adjustments in converting the deaths data.

The reduction in join age means that the age distribution of over-85s is left free to match the shape implied by the KT method (with our proposed modifications), rather than the over- and under-90 totals each being fixed to the official ONS estimates.

- The result is an increase in the estimated population at ages 85-89 and a corresponding decrease in the estimated population at ages 90+.
- The impact on ages 90+ is greater in percentage terms due to the lower populations at those ages.
- This increases the implied mortality rates at higher ages, so reducing life expectancy.

A key reason for this varying impact by age is that past survivor ratios for the elderly have exhibited a consistently higher upward trend at younger ages than at older ages, a feature which is better captured by our proposed method than by the traditional version of the KT method (as currently used by the ONS).

The use of Lexis adjustments is consistent with the approach taken by other bodies such as the Human Mortality Database (HMD), who describe it as 'one of the most important steps in computing the mortality rates and life tables'. (Although operating on similar principles, the precise adjustments used by bodies like the HMD are likely to differ from those we have determined, due to differences in the detail of our approaches.) It leads to a more accurate estimate of deaths aged at 1 January for the KT method, and tends to produce lower population estimates at very high ages due to allowance for the survival decrement within each single year of age. Again, this increases the implied mortality rates at higher ages, and reduces estimates of life expectancy as a consequence.

Chart 12.1: Cumulative impact on period life expectancy from age 95 (2011-2015) – proposal vs current method for ONS males

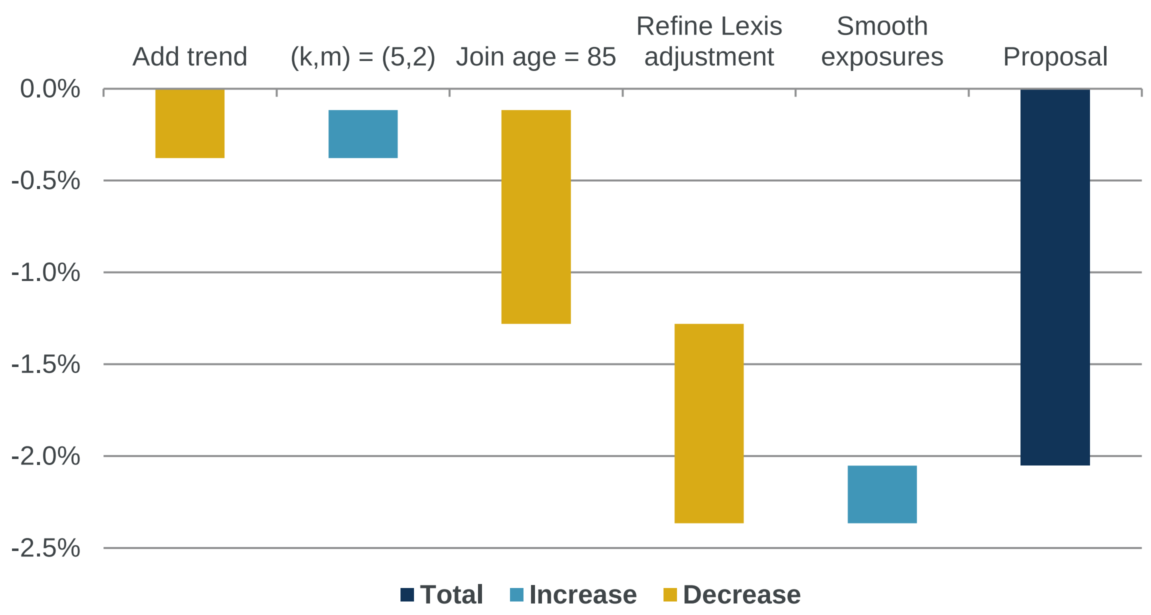

Chart 12.2: Cumulative impact on 2015 population estimates – proposal vs current method for ONS males

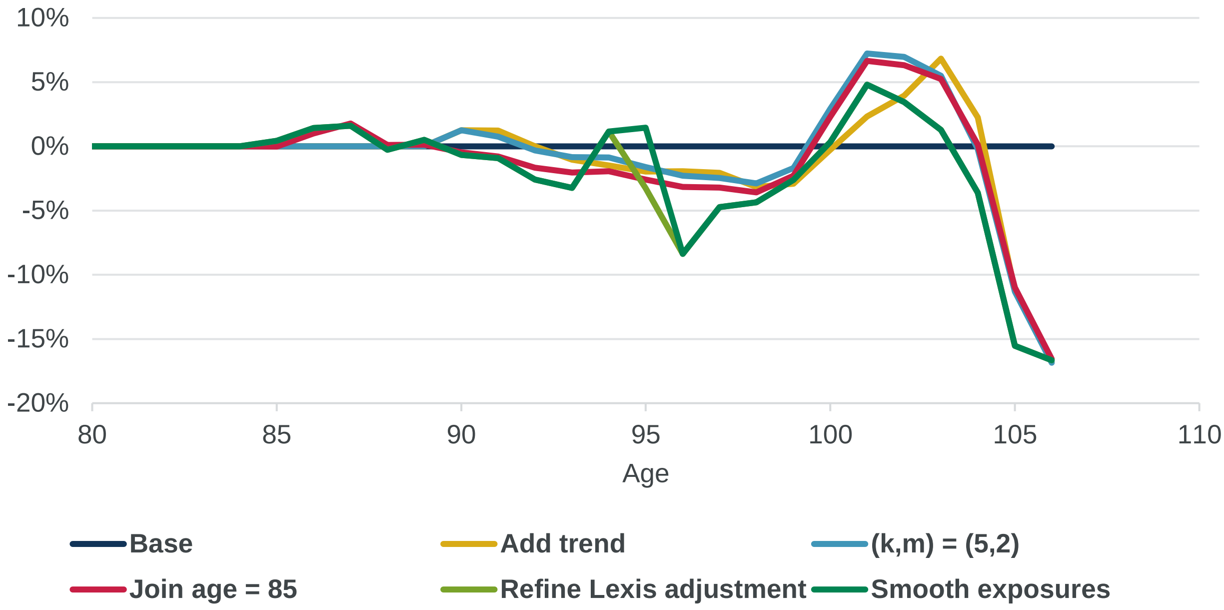

Commentary

Chart 12.2 shows that the 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. This is the underlying reason for the higher mortality rates (and hence lower life expectancies) at high ages which emerge in Chart 12.1.

(The hump exhibited between ages 100 and 105 is an incidental feature driven by a particular aspect of the trend implementation, which we explain further in Appendix 2. It is not in fact material to life expectancy, even at age 95.)

The main age shape to the population changes has been introduced by altering the survivor ratio projection in the KT method (to incorporate an allowance for mortality trend, with m reduced from 5 to 2). This change in shape is not material with join age 90, but impacts the life expectancies once the join age has been reduced to 85 (since this weakens the constraint placed on the 90+ population total, allowing it to reflect the survivor ratio projection under the KT method).

We have already noted that this reduction in 90+ population estimates seems reasonable against the backdrop of accumulating evidence that the official ONS 90+ figures may be overstated.

But it is worth taking a final step back to consider:

- why the incorporation of a mortality trend allowance in the survivor ratios is producing this effect; and
- whether it makes sense.

Chart 12.3 plots the progression of survivor ratios in the male ONS data since the year 2000, alongside the projected survivor ratio for the final year (to the right-hand side of the dashed line), for:

- the base version of the KT method currently used by the ONS (left-hand chart); and
- our proposed variant with mortality trend allowance and m = 2 (right-hand chart).

The past survivor ratios are plotted against the average year of the survivor figures from the m cohorts included in the ratio. (This is then consistent with the effective year of the *projected* survivor ratio to the right of the dashed line, which represents a single cohort.)

Chart 12.3: survivor ratio progression – proposal vs current method for ONS males

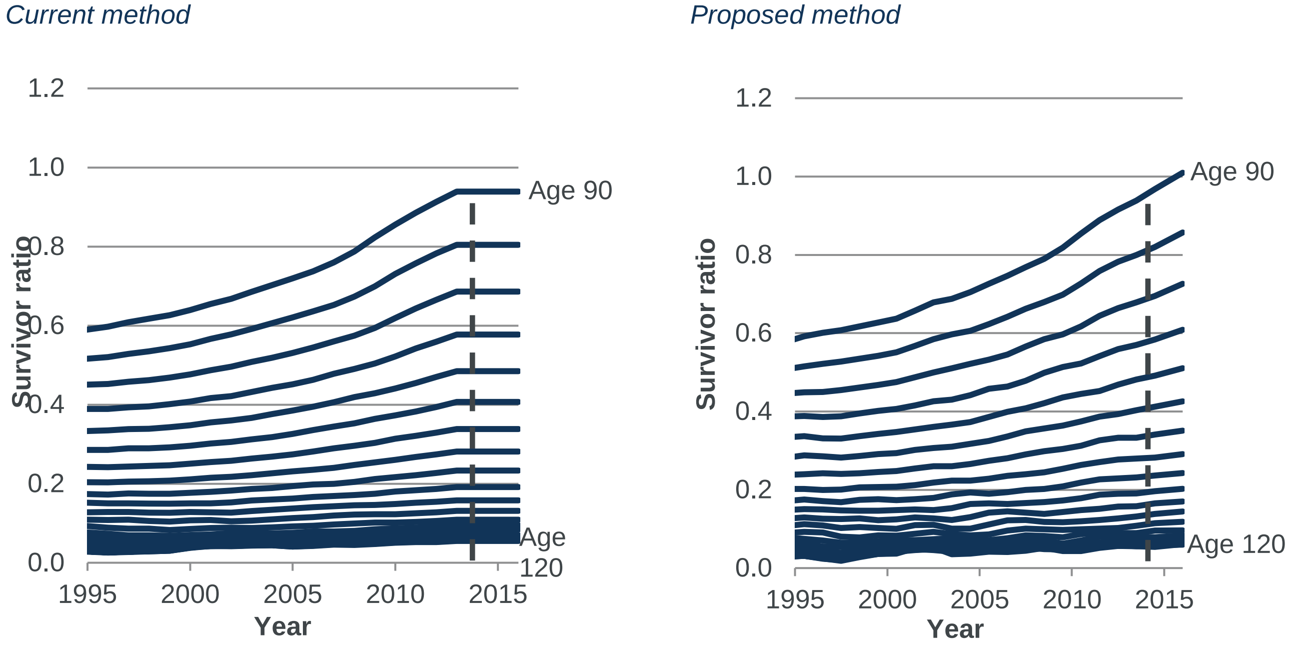

There are two features to draw out from these charts:

- there is a clear upward trend in survivor ratios over time; and
- the upward trend is greatest at younger ages, but tapers down towards zero at older ages.

This is indicative of a regime of progressive mortality improvement, with the rate of improvement reducing towards the highest ages – a message consistent with previous and current modelling of mortality improvements by the CMI.

Survivor ratios for the final year under the current method (which are projected with no allowance for the trend) are likely to be understated as a result. The KT method applies a global correction factor (to constrain to the high age population total in the final year), but this implies a common improvement rate in the survivor ratios across all ages. The net effect is that the final year survivor ratios at younger ages are still likely to be understated, whereas those at older ages are likely to be overstated post-correction.

In contrast, adopting an explicit trend allowance in the projection leads to an age-dependent pattern of improvements which appears to better reflect the recent trends in survivor ratios by age. Taking this approach leads to higher survivor ratios (and hence population estimates) at younger ages, and lower survivor ratios at older ages, than the current method adopted by the ONS. We note that this finding is consistent with previous research by Kirill Andreev, which also proposed an age-dependent mortality trend allowance in the construction of high age population estimates (albeit in a slightly different form to our proposal).

12.4 Impact diagnostics – female base mortality

Chart 12.4 illustrates the step-wise impact on female period life expectancy at age 95 from each element of our proposal, and Chart 12.5 shows the underlying population estimates (for 2015). As for males, the choice of age 95 for Chart 12.4 is purely illustrative – we have also covered the impact on life expectancy across a range of younger ages later (in Chart 12.10) for completeness.

For females, the reduction in join age has a similar impact on the general shape of the age distribution as for males, insofar as it generally increases the estimated number of individuals towards younger ages (and decreases the number towards older ages) in the 85+ population.

However, the crossing point occurs at a later age than for males (which makes sense given the typically longer life expectancy of females and hence shift in distribution of surviving individuals towards higher ages). This means that the reduction in life expectancy does not emerge until later ages in Table 12.1.

Chart 12.4: Cumulative impact on period life expectancy from age 95 (2011-2015) – proposal vs current method for ONS females

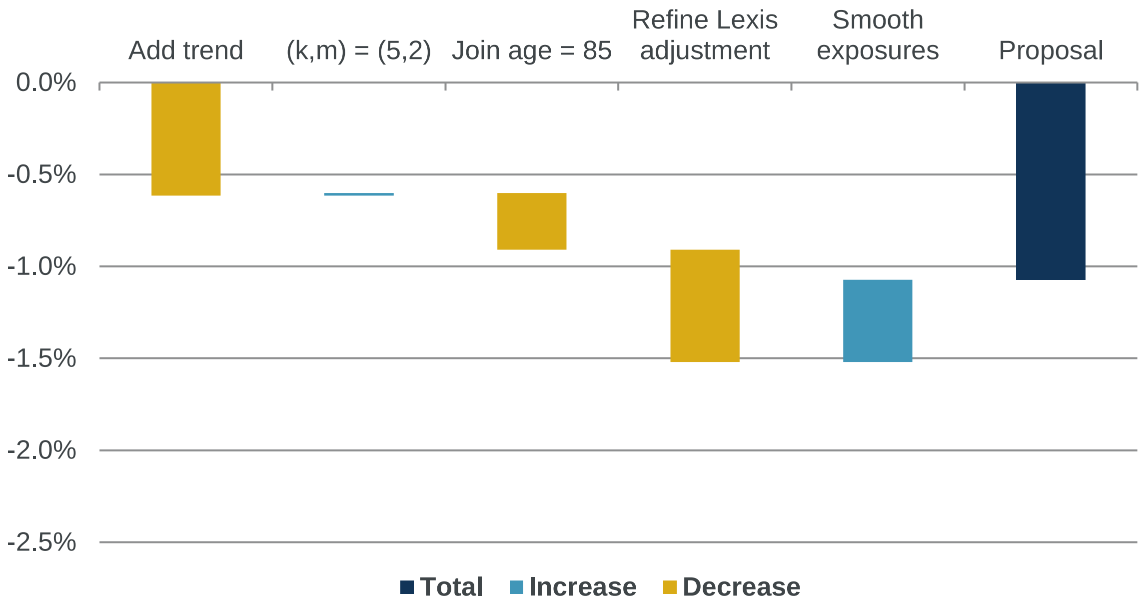

Chart 12.5: Cumulative impact on 2015 population estimates – proposal vs current method for ONS females

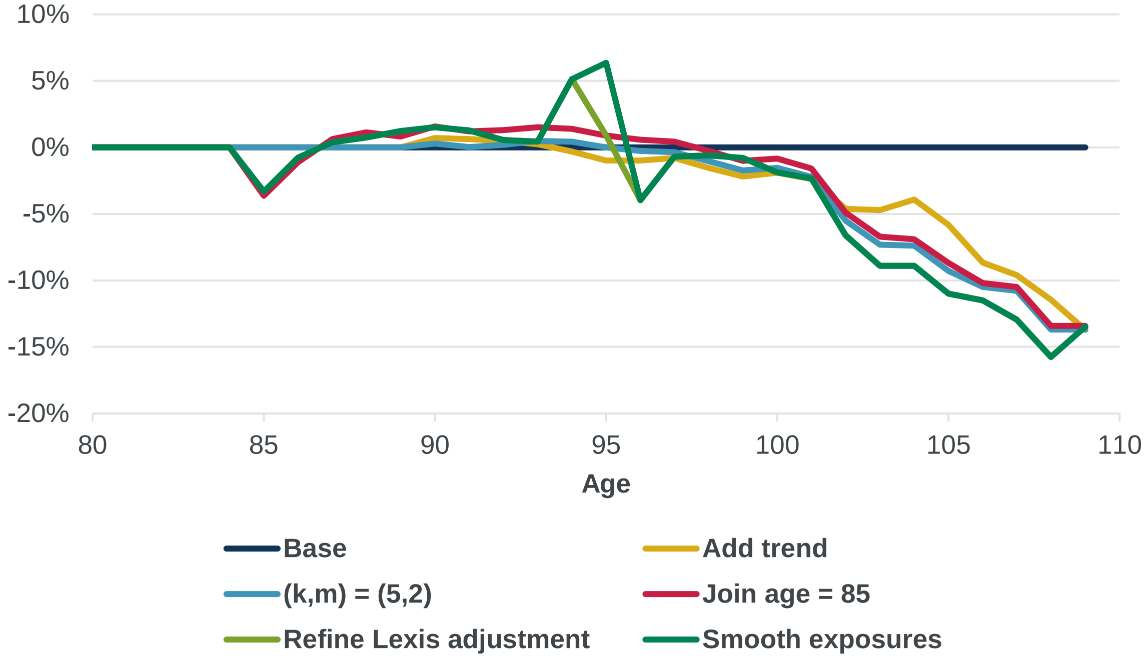

12.5 Impact diagnostics – mortality improvement projections

Sections 12.3 and 12.4 analysed the impact of our proposal on base mortality estimation, in the form of current population exposures and period life expectancies. This is likely to be a key area of interest for those graduating base mortality curves with reference to the national population, for example where using the national estimates to close off portfolio curves at high ages (as described in Working Paper 100).

However, the ONS E&W population datasets are also used as the basis for projecting national mortality *improvement* trends – not least, by the CMI's own Mortality Projections Committee when calibrating its projections model. It is therefore important to understand the effect of our proposed changes, not just on recent base mortality, but also on the projection of mortality trends based on historic data.

Population exposures

A sensible first step is to review the impact on population exposures by age, as in Charts 12.2 and 12.5 but extending back over prior years rather than showing 2015 alone. This requires a two-dimensional heatmap, which we have provided in Chart 12.6.

Chart 12.6: Impact on 2015 population estimates – proposal vs current method for ONS E&W data

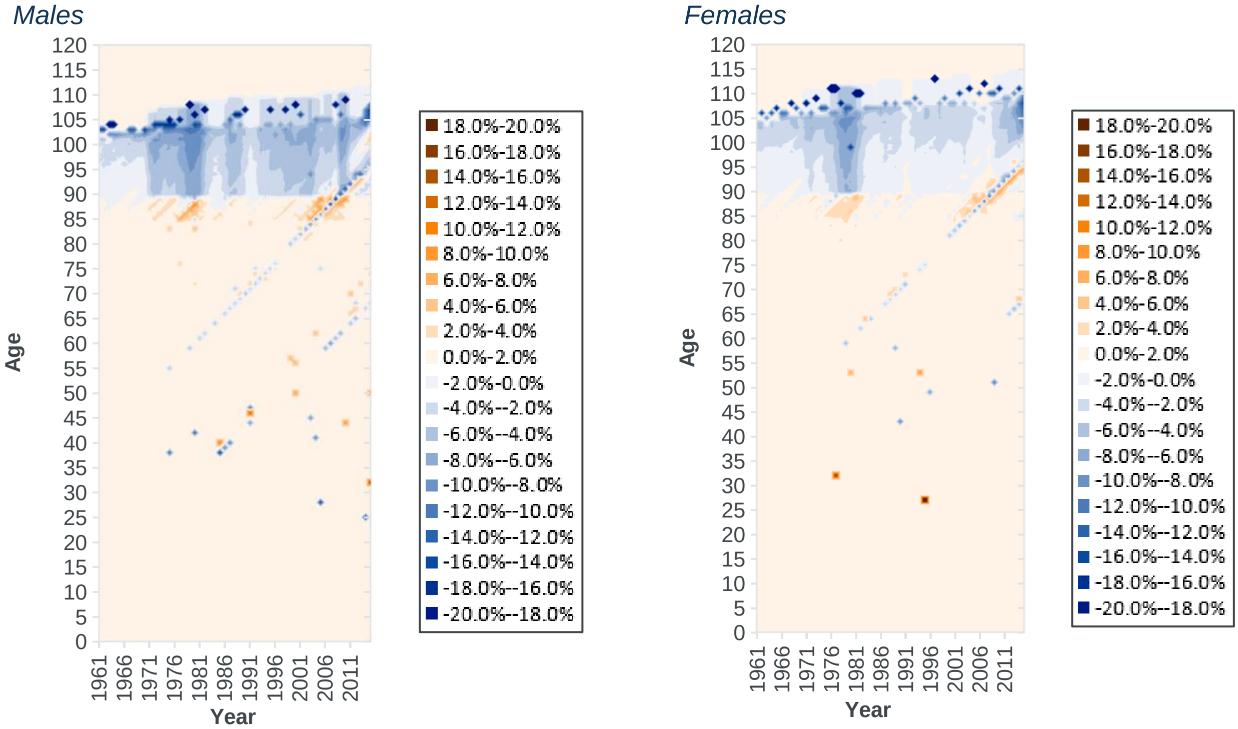

For both males and females, we can see a consistent pattern of reduced population exposures at very high ages (balanced by modest population increases between ages 85 and 90), not just for 2015 but across most of the historic period also. The stability of this pattern is comforting, and indicates that the results we have seen for 2011-2015 are not merely an accident of that particular period or of statistical noise in the data. It is also consistent with the premise that the official ONS 90+ population figures may be systemically overstated relative to an extinct cohorts reconstruction from the death records. Weakening the constraint to these 90+ totals (by reducing the join age to 85) would be expected to result in reductions to the 90+ population estimates historically, which is exactly what we see above.

The heatmaps also show a persistent adjustment to population estimates for the 1919/20 birth cohorts, appearing as a diagonal line across both the higher and younger age regions of each chart. This makes sense given what we know about the distributional anomalies associated with the 1919/20 birth cohorts, and the Lexis and exposure smoothing elements of our proposal which aim to address them.

Mortality improvement projections

Having determined the effect on historic population exposures, we can move on to measure the resulting impact on projections of mortality improvement rates. One way of assessing this is to consider the impact on the latest version of the CMI Mortality Projections Model[15] (the CMI Model), which 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 have applied the CMI Model to the ONS E&W death and population data up to calendar year 2015, taking the current method of estimating the high age population exposures as our starting point. Based on the Core parameters of the model, and an assumed long-term rate of mortality improvement of 1.5% p.a., this produces a grid of assumed mortality improvements for each sex. We have then calculated the cohort life expectancies that result from applying these improvements to a fixed initial mortality curve for each sex as at 1 January 2007 (we have used the published S2PMA and S2PFA tables for this purpose). The cohort life expectancies themselves are measured as at 1 January 2017 and calculated in line with the sample life expectancy methodology presented within the CMI Model itself.

Finally, we have repeated the process based on re-calibrating the CMI Model to ONS E&W data with our proposed modifications to the KT method introduced one by one. Charts 12.7 and 12.8 show the cumulative percentage impact on cohort life expectancy from age 95 of introducing these changes, for males and females respectively. By comparing cohort life expectancies using a fixed base mortality curve as at 1 January 2007, this chart effectively isolates the impact of our proposal on the mortality improvement projections alone.

The net effect on cohort life expectancy at ages 85 and 90 is similar, and covered later (in Charts 12.11 and 12.12) for completeness.

Chart 12.7: Cumulative impact on cohort life expectancy from age 95 as at 1 January 2017 – proposal vs current method for ONS males

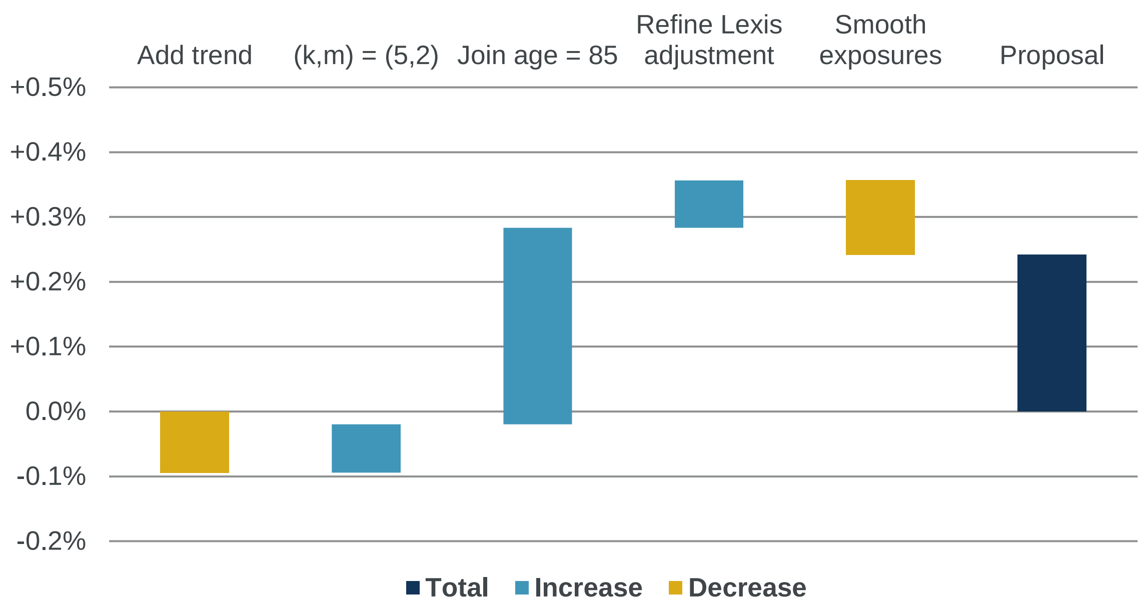

Chart 12.8: Cumulative impact on cohort life expectancy from age 95 as at 1 January 2017 – proposal vs current method for ONS females

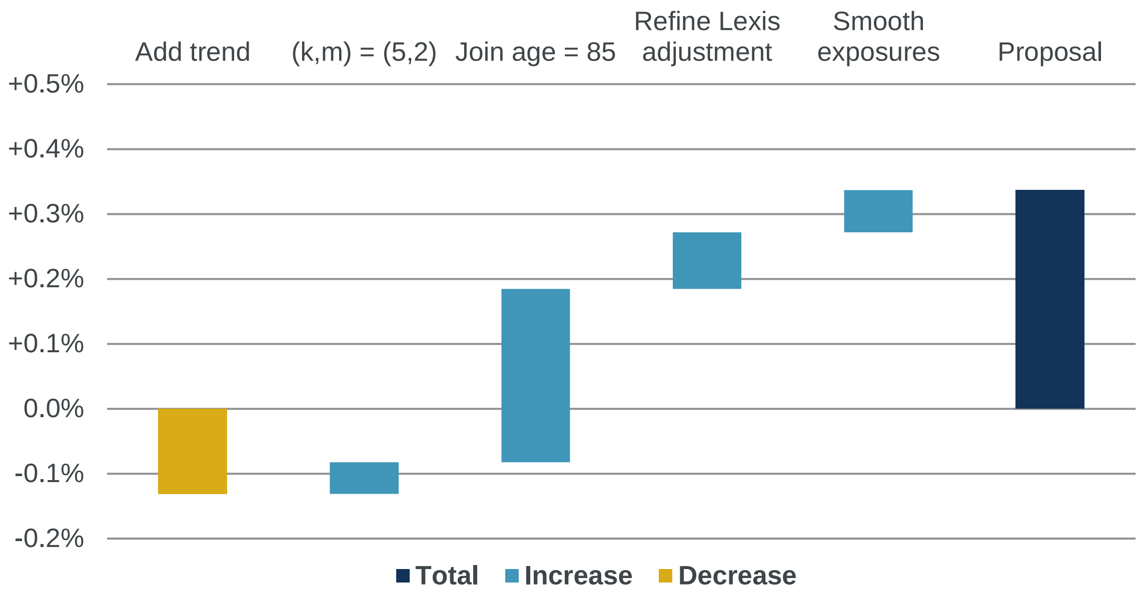

Commentary

Charts 12.7 and 12.8 each show small increases in cohort life expectancy as a result of our proposals, with the reduction in join age a key driver.

It makes sense that the overall size of impact here is smaller than the impact on period life expectancy shown in Charts 12.1 and 12.4:

- The main feature of our proposed changes is to reduce population exposures at ages 90+, which increases mortality rates in that age range and so reduces period life expectancy.
- However, we saw from Chart 12.6 that this effect was fairly consistent across past years – so although high age mortality rates have increased in each year, the rate of relative *improvement* in those rates from year to year has not been heavily impacted.
- Of course, the smoothing inherent in the CMI Model may also be damping the final impact.

It also makes sense that the reduction in join age is a key driver.

- The earlier elements of our proposal (adding a trend allowance and reducing m) affect the survivor ratio part of the KT method, which feeds into the population estimates for recent years but has no impact on historic years for which all the cohorts are now extinct. This may be why there is a modest reduction in cohort life expectancy from introducing trend (since the increase in base mortality which we have discussed previously is introduced in the most recent years, but *not* in historic years).
- In contrast, reducing join age to 85 allows the historic population exposures to depart from their ONS 90+ population constraints, which we have seen are consistently higher than exposures implied by the unconstrained method of extinct generations.
- Reducing the historic population exposures in this way means that the mortality rates for past years now also increase, uplifting the implied improvement rate and hence the cohort life expectancy figures.

That the *net* impact of our proposals should be a modest increase (rather than decrease) in cohort life expectancy is perhaps less obvious, as it depends on the relative impact of the changes on different years historically and the way these are captured by the CMI Model's age-period and cohort improvement components. We do, however, note that an increase in projected improvement rates is at least consistent with the further analysis in Appendix 1, which indicates that (for males) our proposals may lead to a greater increase in high age mortality over the period 2006-2010 than the period 2011-2015.

12.6 Performance diagnostics

For completeness, we have summarised the step-wise impact of our proposal (for each of males and females) on:

- a wider range of period and cohort life expectancies (Charts 12.9 to 12.12) and
- a number of other key performance diagnostics (Charts 12.13 to 12.16).

These charts are based on the ONS data for E&W.

We have also provided, in Appendix 1, a brief analysis of the sensitivity of our diagnostics to a different choice of data period. The results in that appendix are based on applying the KT method to data ending in calendar year 2010 (rather than 2015) and review the impact of our proposed changes on average mortality rates over calendar years 2006-2010 (rather than 2011-2015). Our proposal appears to be robust under this sensitivity, insofar as it exhibits similar outperformance over the current method across our main performance diagnostics. The actual impact on mortality rates does, however, appear to be more pronounced for the 2006-2010 period than the 2011-2015 period for males, with greater reductions in period life expectancy resulting. This may be a feature of the higher prevailing mortality improvements prior to 2011 (such that our introduction of a trend allowance in the KT method has more impact over this earlier period than over 2011-2015).

Chart 12.9: Cumulative impact on period life expectancy (2011-2015) – proposal vs current method for ONS males

Trend	Base	No trend	Trend	Trend	Trend	Trend	Trend
Lexis		50/50 Lexis	50/50 Lexis	50/50 Lexis	50/50 Lexis	Refined Lexis	Refined Lexis
Smoothing		No smoothing	No smoothing	No smoothing	No smoothing	No smoothing	CMI smoothing
Join age		90	90	90	85	85	85
Parameter k		5	5	5	5	5	5
Parameter m		5	5	2	2	2	2
Life expectancy from age 65	18.52	-	-	-	-	-	-
Life expectancy from age 70	14.72	-	-	-	-	-	-
Life expectancy from age 75	11.31	-	-	-	-0.1%	-0.1%	-
Life expectancy from age 80	8.29	-	-	-	-0.1%	-0.1%	-0.1%
Life expectancy from age 85	5.85	-	-	-	-0.2%	-0.2%	-0.1%
Life expectancy from age 90	4.04	-	-0.1%	-	-1.4%	-1.6%	-1.3%
Life expectancy from age 95	2.79	-	-0.4%	-0.1%	-1.3%	-2.4%	-2.1%

Chart 12.10: Cumulative impact on period life expectancy (2011-2015) – proposal vs current method for ONS females

Trend	Base	No trend	Trend	Trend	Trend	Trend	Trend
Lexis		50/50 Lexis	50/50 Lexis	50/50 Lexis	50/50 Lexis	Refined Lexis	Refined Lexis
Smoothing		No smoothing	No smoothing	No smoothing	No smoothing	No smoothing	CMI smoothing
Join age		90	90	90	85	85	85
Parameter k		5	5	5	5	5	5
Parameter m		5	5	2	2	2	2
Life expectancy from age 65	20.99	-	-	-	-	-	-
Life expectancy from age 70	16.89	-	-	-	-	-	-
Life expectancy from age 75	13.09	-	-	-	-	-	-
Life expectancy from age 80	9.67	-	-	-	-	-	-
Life expectancy from age 85	6.83	-	-	-	-	-0.1%	+0.1%
Life expectancy from age 90	4.66	-	-0.1%	-0.1%	-0.2%	-0.3%	-
Life expectancy from age 95	3.18	-	-0.6%	-0.6%	-0.9%	-1.5%	-1.1%

Chart 12.11: Cumulative impact on cohort life expectancy as at 1 January 2017 – proposal vs current method for ONS males

Trend Lexis Smoothing	Base	No trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend Refined Lexis No smoothing	Trend Refined Lexis CMI smoothing
Join age		90	90	90	85	85	85
Parameter k		5	5	5	5	5	5
Parameter m		5	5	2	2	2	2
Life expectancy from age 85	6.92	-	-	-	+0.1%	+0.1%	+0.3%
Life expectancy from age 90	4.46	-	-	-	-	+0.1%	+0.2%
Life expectancy from age 95	2.89	-	-0.1%	-	+0.3%	+0.4%	+0.2%

Chart 12.12: Cumulative impact on cohort life expectancy as at 1 January 2017 – proposal vs current method for ONS females

Trend Lexis Smoothing	Base	No trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend Refined Lexis No smoothing	Trend Refined Lexis CMI smoothing
Join age		90	90	90	85	85	85
Parameter k		5	5	5	5	5	5
Parameter m		5	5	2	2	2	2
Life expectancy from age 85	7.69	-	-	-	+0.1%	+0.1%	+0.2%
Life expectancy from age 90	5.01	-	-	-	+0.3%	+0.3%	+0.4%
Life expectancy from age 95	3.27	-	-0.1%	-0.1%	+0.2%	+0.3%	+0.3%

Chart 12.13: Cumulative impact on final year balancing adjustment in 2015 – proposal vs current method

Males

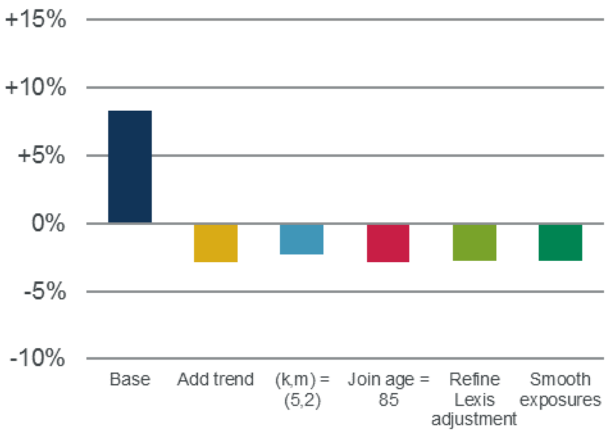

Females

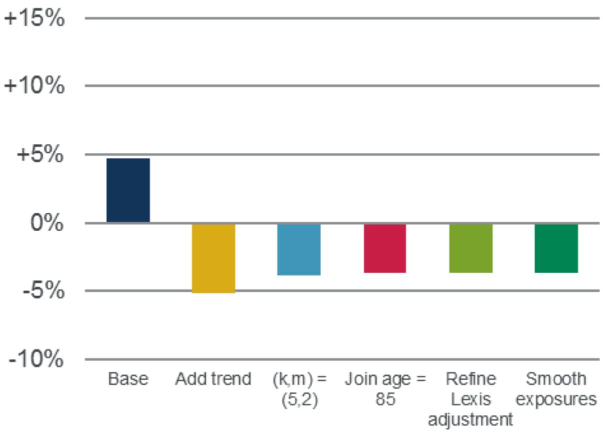

Chart 12.14: Cumulative impact on average annual scaling adjustment – proposal vs current method

Males

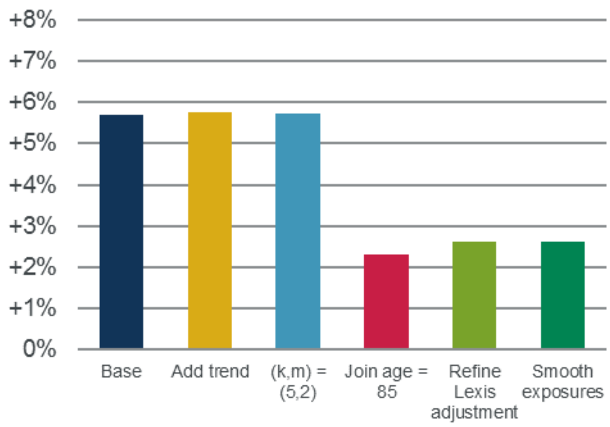

Females

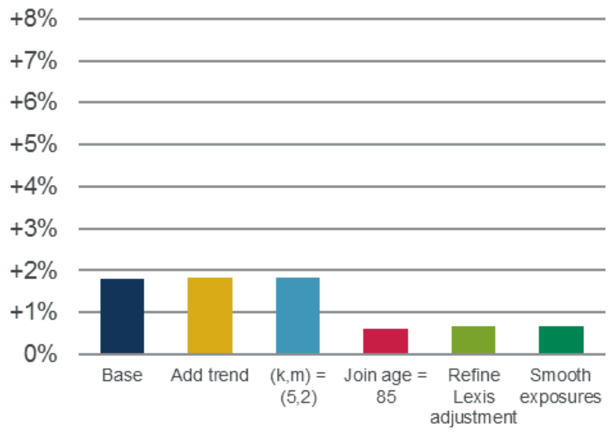

Chart 12.15: Cumulative impact on average cohort inconsistency – proposal vs current method

Males

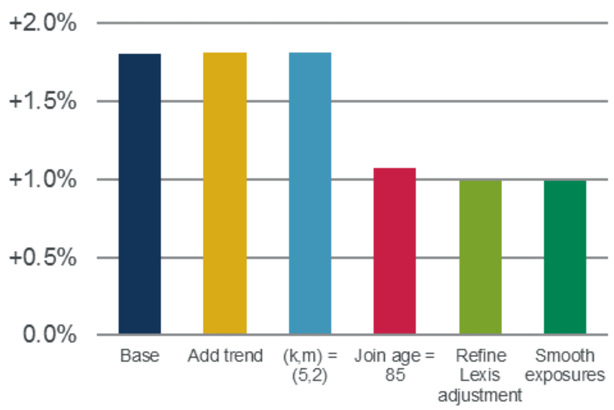

Females

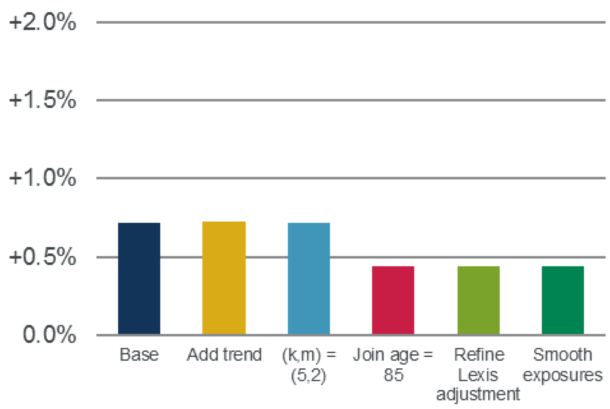

Chart 12.16: Cumulative impact on average mortality deviance across join age – proposal vs current method

Males

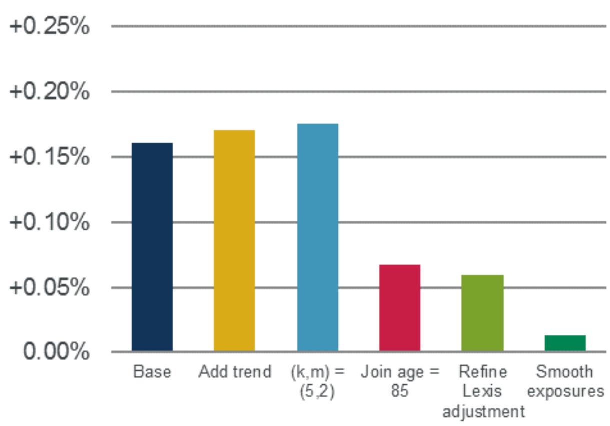

Females

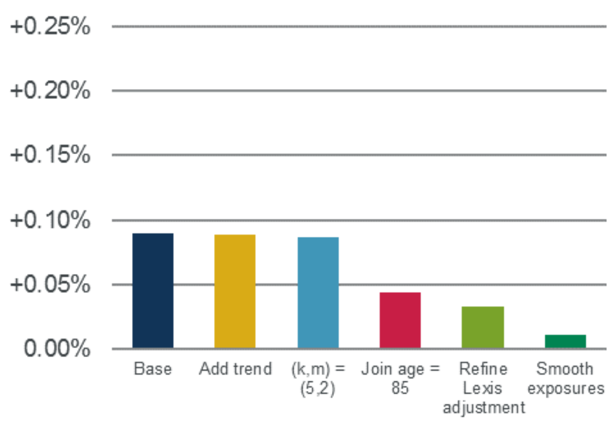

12.7 Further research

The Working Party hopes that this paper and the proposals it contains form a useful contribution to the body of research on high age mortality modelling. Nonetheless, there are several avenues where additional investigation may prove valuable. The Working Party does not presently intend to explore these areas further itself but offers a summary below for interested parties.

As noted in the main body of our analysis, the proposed allowance for mortality improvement trend in the KT method performs very well for simple patterns of mortality improvement but appears less effective for noisy improvement patterns akin to real data. Our testing suggests it fares no worse here than the current approach (i.e. with no trend allowance) but a more predictive model of improvements may be superior in accurately projecting the survivor ratios. There is scope for further research here, in testing the performance of different improvement models within the KT method (for example, models which attribute a cohort component to improvements). Equally, it may be that the recent history of changing improvement trends in the ONS male dataset for E&W is *atypical* of longer run patterns and, by focusing on this period, we may be underestimating the performance of the trend variant in more normal circumstances. Further comfort could be obtained in this area by testing the variant on alternative datasets and/or periods of time.

More generally, we have attempted to test the predictive power and robustness of different variants by applying them to a range of synthetic population datasets with different properties. We have endeavoured to include a range of different population structures (including populations evolved under realistic/noisy patterns of mortality improvement) but there is still scope for further back-testing on real datasets. In particular, it would be informative to review the performance of our proposed (and other) variants in application to

- historic ONS data for which all cohorts are now extinct, and
- high quality datasets from other countries where the population and death counts are believed to be reliable at high ages (e.g. Sweden and Finland).

The advantage of both these data sources is that they provide more trustworthy figures against which the population estimates under a given method can be compared for accuracy (in contrast to the ONS data for non-extinct cohorts, which is itself subject to reliability concerns).

By a similar token, our suggestion to reduce the size of the survivor ratio window (changing m from 5 to 2) is targeted principally at reliable estimation of the E&W population exposures, and has been tested in this light. The male and female ONS datasets and the synthetic population datasets used as the basis of our analysis in this paper all reflect E&W population data volumes. It may be that this choice of parameters performs less well for alternative (e.g. smaller) populations, if the data credibility leads to a different trade-off between stability and resolution. Analysis of alternative datasets may provide further comfort in this regard.

Appendix 1 Period sensitivity

The main analysis in this paper has focused on applying the KT method to ONS E&W data up to calendar year 2015 inclusive, and measuring the effect of our proposals on base mortality over calendar years 2011-2015. In this appendix we repeat a number of our key diagnostics for an alternative period instead, applying the KT method to data up to 2010 and measuring the impact over 2006-2010. The motivation for this is two-fold:

- 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).

By truncating the data to 2010, we aim to produce a meaningful sensitivity for which our impact measures do not overlap with the original period chosen (2011-2015). We are also conscious that prevailing national mortality trends up to 2010 were quite different to the lower mortality improvement regime of more recent years, which provides a useful change in scenario to help ensure our conclusions are reliable.

Performance diagnostics

Charts A1.1 to A1.4 show the step-wise performance of our proposed changes to the KT method when applied to the alternative data period.

The improved performance under our proposals which was exhibited in the main analysis (in particular, Section 12.6) is again evident here. There is a general trend towards better outcomes (i.e. smaller values of each diagnostic metric) as the successive steps of our proposal are introduced, with certain elements having a notable impact on particular diagnostics. For example:

- introducing the trend allowance and reducing m improves the final year balancing adjustments;
- reducing the join age to 85 drives the main reduction in the annual average scaling adjustment and average cohort inconsistency metrics;
- smoothing the final exposures noticeably reduces the average mortality deviance across the join age.

We note in passing that, for the period ending in 2010, introducing a trend allowance to the survivor ratio projection has not changed the sign of the final year balancing adjustments. This make sense given our earlier suggestion that the move to negative adjustments in the main analysis had been related to the off-trend heavy experience of 2015 (recall Section 6).

Taken together, the performance diagnostics for this alternative data period provide comfort that the changes we are proposing to the KT method in this paper are robust.

Chart A1.1: Cumulative impact on final year balancing adjustment in 2010 – proposal vs current method (applied to data up to 2010)

Males

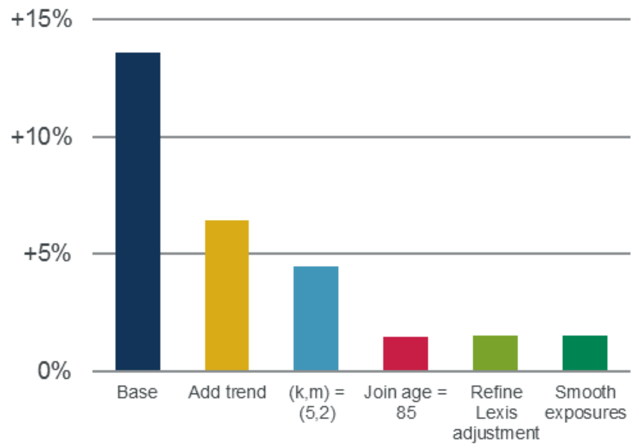

Females

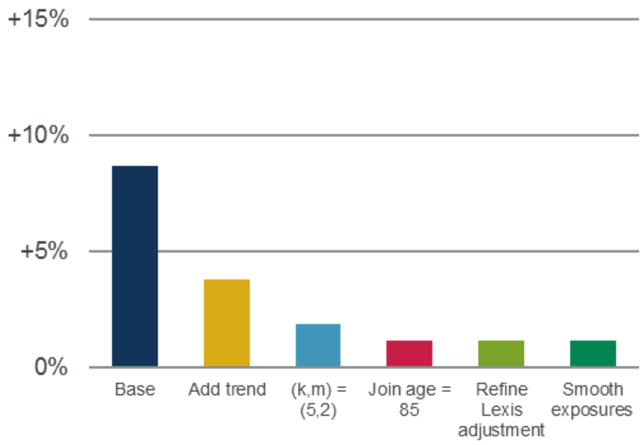

Chart A1.2: Cumulative impact on average annual scaling adjustment – proposal vs current method (applied to data up to 2010)

Males

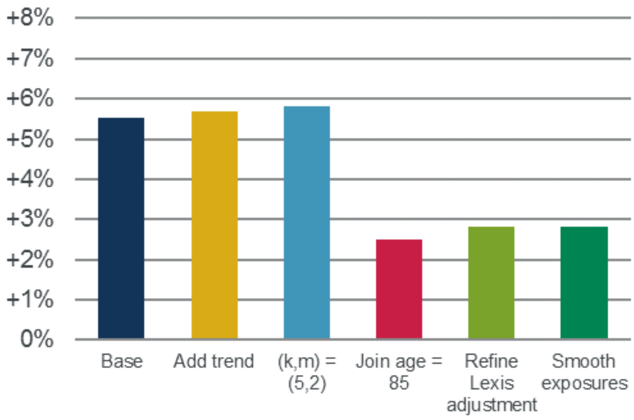

Females

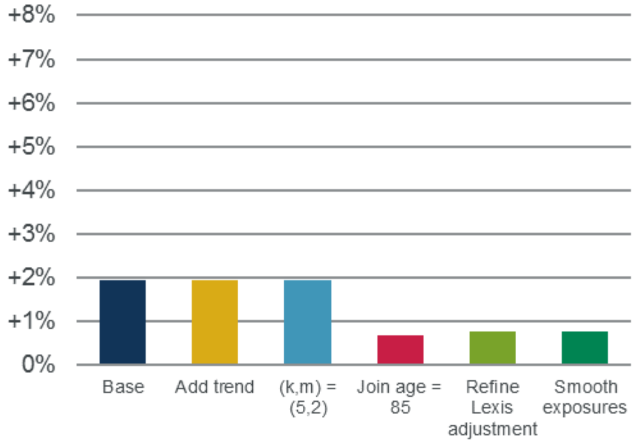

Chart A1.3: Cumulative impact on average cohort inconsistency – proposal vs current method (applied to data up to 2010)

Males

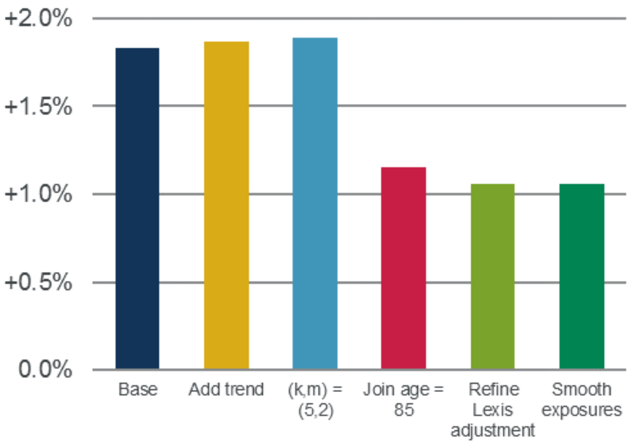

Females

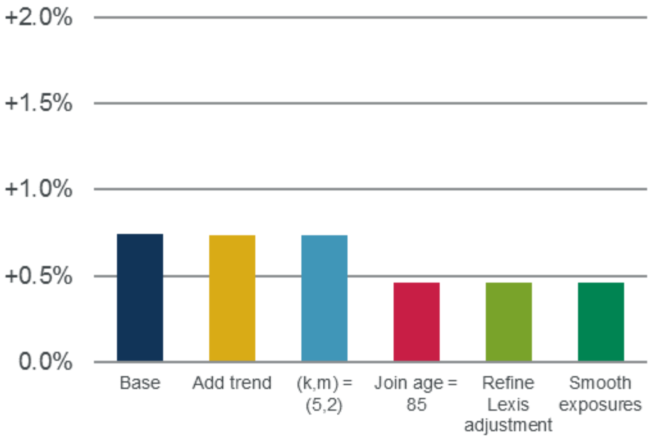

Chart A1.4: Cumulative impact on average mortality deviance across join age – proposal vs current method (applied to data up to 2010)

Males

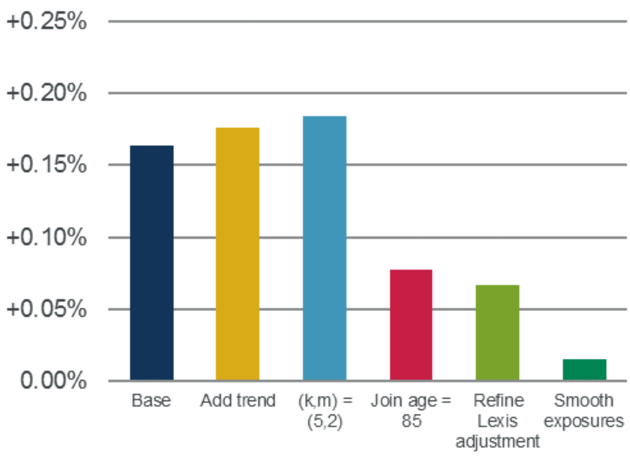

Females

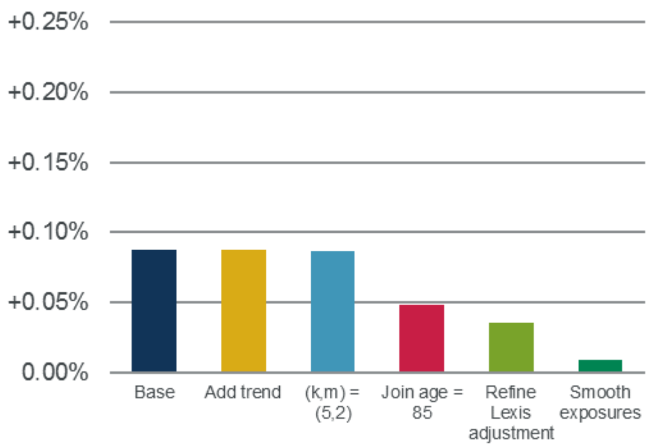

Impact diagnostics

Charts A1.5 and A1.6 show the step-wise impact on period life expectancy from each element of our proposal, when applied to the data up to 2010. For ease of comparison we have also included a final column showing the corresponding total impacts from our original analysis, i.e. based on the data up to 2015. (This final column is taken from Charts 12.9 and 12.10 respectively.)

Chart A1.5: Cumulative impact on period life expectancy (2006-2010) – proposal vs current method for ONS males (applied to data up to 2010)

		Impact on 2006-2010 period life expectancy						Impact on 2011-2015 period life expectancy
Trend	**Base**	No trend	Trend	Trend	Trend	Trend	Trend	Trend
Lexis		50/50 Lexis	50/50 Lexis	50/50 Lexis	50/50 Lexis	Refined Lexis	Refined Lexis	Refined Lexis
Smoothing		No smoothing	No smoothing	No smoothing	No smoothing	No smoothing	CMI smoothing	CMI smoothing
Join age		90	90	90	85	85	85	85
Parameter k		5	5	5	5	5	5	5
Parameter m		5	5	2	2	2	2	2
Life expectancy from age 65	17.68	-	-	-	-0.1%	-0.1%	-0.1%	-
Life expectancy from age 70	14.02	-	-	-	-0.1%	-0.1%	-0.1%	-
Life expectancy from age 75	10.71	-	-	-	-0.2%	-0.2%	-0.1%	
Life expectancy from age 80	7.89	-	-	-0.1%	-0.3%	-0.3%	-0.2%	-0.1%
Life expectancy from age 85	5.65	-	-0.1%	-0.2%	-0.6%	-0.7%	-0.5%	-0.1%
Life expectancy from age 90	4.02	-	-0.3%	-0.5%	-3.5%	-4.0%	-3.7%	-1.3%
Life expectancy from age 95	2.81	-	-2.5%	-4.2%	-5.9%	-6.9%	-6.9%	-2.1%

Chart A1.6: Cumulative impact on period life expectancy (2006-2010) – proposal vs current method for ONS females (applied to data up to 2010)

		Impact on 2006-2010 period life expectancy						Impact on 2011-2015 period life expectancy
Trend	**Base**	No trend	Trend	Trend	Trend	Trend	Trend	Trend
Lexis		50/50 Lexis	50/50 Lexis	50/50 Lexis	50/50 Lexis	Refined Lexis	Refined Lexis	Refined Lexis
Smoothing		No smoothing	No smoothing	No smoothing	No smoothing	No smoothing	CMI smoothing	CMI smoothing
Join age		90	90	90	85	85	85	85
Parameter k		5	5	5	5	5	5	5
Parameter m		5	5	2	2	2	2	2
Life expectancy from age 65	20.36	-	-	-	-	-	-	-
Life expectancy from age 70	16.33	-	-	-	-	-	-	-
Life expectancy from age 75	12.61	-	-	-	-	-	+0.1%	-
Life expectancy from age 80	9.34	-	-	-	-	-	+0.1%	-
Life expectancy from age 85	6.65	-	-	-	-	-0.1%	+0.2%	+0.1%
Life expectancy from age 90	4.57	-	-	-0.1%	-	-0.3%	-0.2%	-
Life expectancy from age 95	3.14	-	-0.7%	-0.8%	-0.3%	-0.9%	-0.9%	-1.1%

As for the 2011-2015 period in our main analysis, the charts show a reduction in period life expectancy at high ages as a result of our proposals. For females, this is of a similar magnitude to the reduction in the main analysis. For males, however, the size of impact is somewhat greater for the 2006-2010 period than the 2011-2015 period examined in the main body of this paper.

We considered whether this might be due to the timing of the decennial censuses underlying the official ONS population totals which are used to constrain the high age totals. We have already seen that there is growing evidence the official 90+ totals may be overstated, and that the principal impact of our proposals is to reduce population exposure estimates at the highest ages by departing from the constraint to these 90+ totals. The lag between the 2006-2010 period and its preceding census (in 2001) is greater than the lag between the 2011-2015 period and the corresponding census in 2011. So it might be reasonable to hypothesise that the official 90+ totals between 2006 and 2010 are less accurate than those between 2011 and 2015, with greater scope for overstatement and hence potentially greater life expectancy reductions once our proposals are introduced. In practice, however, this does not appear to be the case. We see similarly large life expectancy reductions to the 2006-2010 period when applying the KT method up to 2001 only. And of course the official population totals between 2001 and 2011 were restated by the ONS following the 2011 Census, so it is not clear that an accumulating roll-forward error should be expected as one advances through that period in any case.

On closer inspection, it appears that there may be another explanation for the larger impacts on male life expectancy observed over 2006-2010 compared with 2011-2015:

- As discussed in Section 12.3, the trend element of our proposal allows the projection of survivor ratios to reflect the age dependence of prevailing mortality improvements. This typically leads to a shallower increase in the projected survivor ratios (and hence lower population estimates) at the highest ages, where mortality improvements tend to be smaller.
- Over the years leading up to 2015, national mortality improvements were relatively low across *all* ages, not just the highest ages. This means that the introduction of localised mortality trends under our proposal produced only a modest age differential in the projection of survivor ratios – see Chart A1.7(a).
- However, the average level of improvements at younger ages had been much higher over the period to 2010 (and still tapered towards zero at the highest ages – Chart A1.7(b)). This meant that the age differential in improvement trends was greater, and the effect of introducing the survivor ratio trend allowance correspondingly larger.

Chart A1.7: Final survivor ratio projection for ONS males

Data up to 2015 | *Data up to 2010*

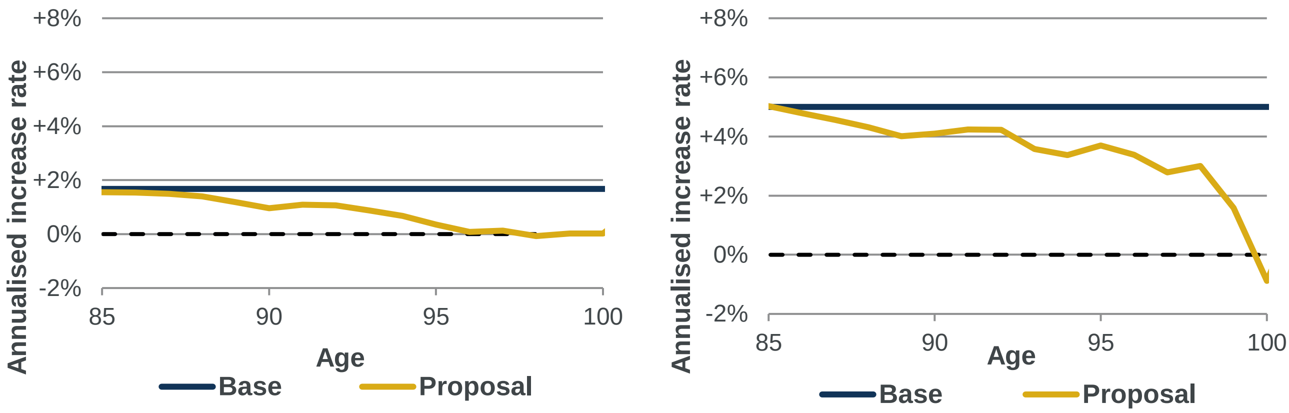

In contrast, the age shape of survivor ratio improvements for females was relatively flat up to age 100 over both periods shown – see Chart A1.8. There was a *slight* drop-off in the projection towards higher ages, for both periods, but this was not as pronounced as for males. Against this backdrop, it makes sense that the impact of our proposal in moving to a trend projection method is smaller for females than for males, and more consistent between the two periods analysed.

Chart A1.8: Final survivor ratio projection for ONS females

Data up to 2015 | *Data up to 2010*

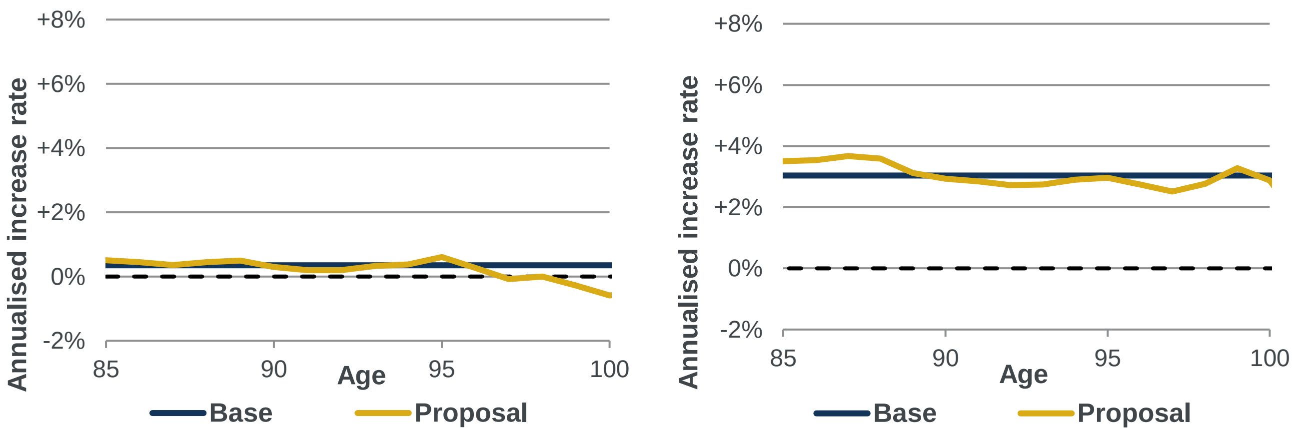

Appendix 2 Survivor ratio application

In Section 12.3 we noted a hump between ages 100 and 105 in Chart 12.2, which showed the impact of our proposals on the estimated 2015 population exposures for ONS E&W males. We have repeated the information as Chart A2.1 for ease of reference.

Chart A2.1: Cumulative impact on 2015 population estimates – proposal vs current method for ONS males

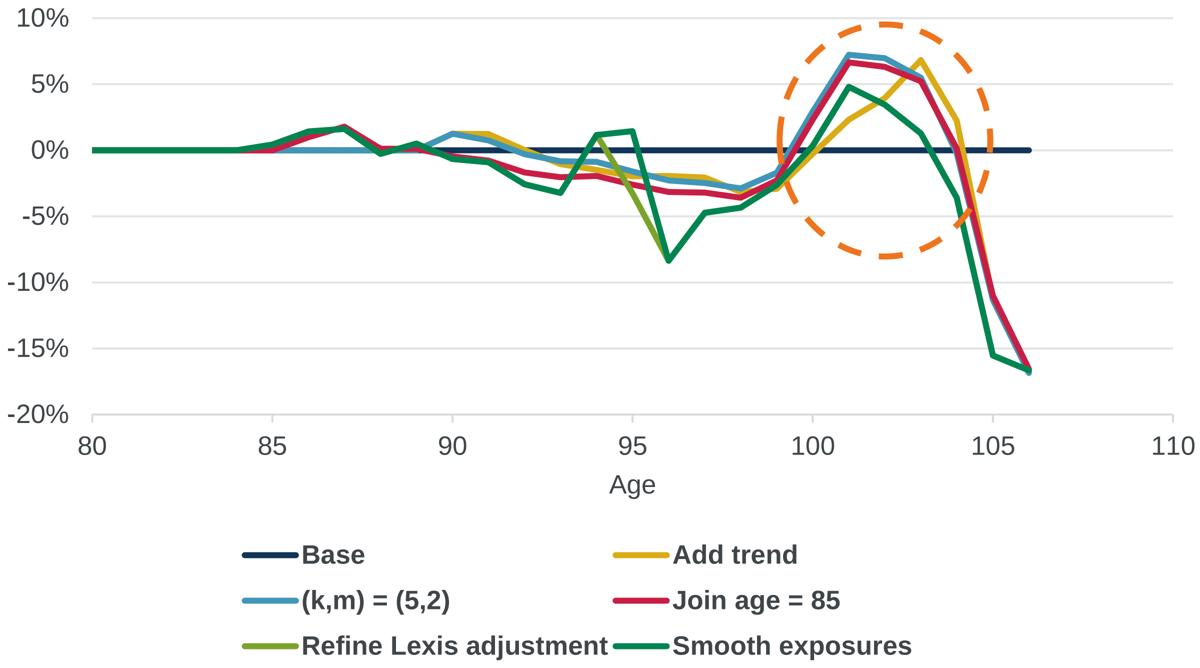

This hump is not apparent for females in 2015, or indeed for either sex in our period sensitivity of Appendix 1 (which applied the KT method to data up to 2010 only). It also does not accord with an intuitive interpretation of the trend allowance we have introduced in our proposals – if anything, we expect this to *reduce* exposures at the highest ages (due to the lower rate of survivor ratio increases typically observed at those ages).
On closer investigation it became apparent that there is another, quite simple, explanation for this shape:

- The population figures to which the chart relates are the mid-2015 estimates, based on an average of the 1 January 2015 and 1 January 2016 estimates produced directly by the KT method.
- The current version of the KT method used by the ONS
  - estimates the 1 January 2015 population using survivor ratios, and then
  - deducts deaths during 2015 to obtain an estimate of the 1 January 2016 population.

  Under this approach, the atypically high death counts for males in 2015 lead to *lower* mid-2015 population estimates than would otherwise be the case.
- In contrast, our trend variant of the KT method
  - estimates the *1 January 2016* population using survivor ratios, and then
  - *adds back in the* 2015 deaths to obtain an estimate of the 1 January 2015 population (in the same way that the KT method generally extends back to prior years).

  Under this approach, the atypically high death counts for males in 2015 lead to *higher* mid-2015 population estimates than would otherwise be the case.

This explains why a humped effect emerges at higher ages for males (where the population has become small enough that the different approach to the 2015 deaths has a material impact on the estimates). It also explains why there is no such effect observed when applying the KT method up to 2010.

Chart A2.2 shows how the impact of our proposals on the 2015 population estimates would look if applying the current version of the KT method with the same survivor ratio timing as our trend variant. You can see that the hump anomaly between ages 100 and 105 has disappeared.

Chart A2.2: Cumulative impact on 2015 population estimates – proposal vs current method for ONS males (restated)

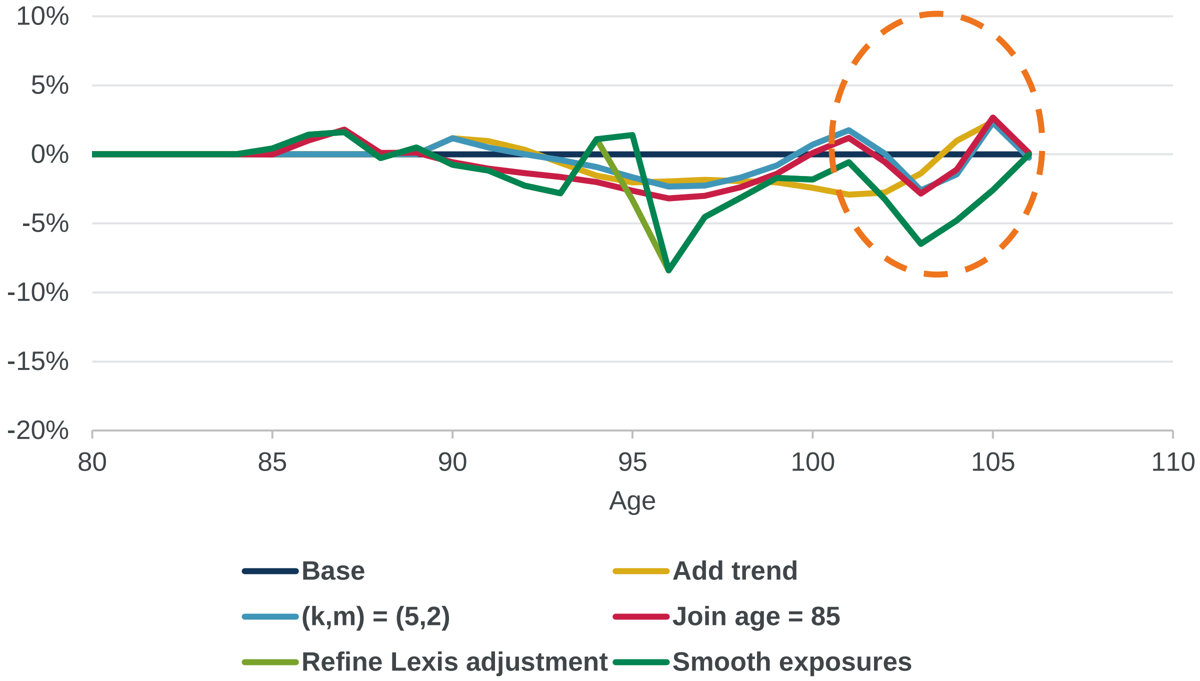

The difference in the application timing of the survivor ratios under our trend variant compared with the current ONS implementation of the KT method is nothing to do, conceptually, with the trend allowance itself. It is an incidental refinement to the implementation, introduced alongside the other elements of our proposal.

We do think that the modification we have introduced makes sense, insofar as it means the survivor ratios are applied to the latest available deaths data rather than the penultimate year of data. However, it is perhaps comforting to note that the effect of this change – although clearly visible in the pattern of population impacts by age – is focused towards such high ages that it is ultimately immaterial to life expectancy. This can be seen in Chart A2.3, which compares:

(a) the actual period life expectancy impacts under our proposal (from Chart 12.9) with

(b) the corresponding impacts which would be obtained if restating the current version of the KT method to use the same survivor ratio timing as our trend variant.

Chart A2.3: Cumulative impact on period life expectancy (2011-2015) – proposal vs current method for ONS males

(a) Original analysis (from Chart 12.9)

Trend Lexis Smoothing	No trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend Refined Lexis No smoothing	Trend Refined Lexis CMI smoothing
Join age	90	90	90	85	85	85
Parameter k	5	5	5	5	5	5
Parameter m	5	5	2	2	2	2
Life expectancy from age 65	-	-	-	-	-	-
Life expectancy from age 70	-	-	-	-	-	-
Life expectancy from age 75	-	-	-	-0.1%	-0.1%	-
Life expectancy from age 80	-	-	-	-0.1%	-0.1%	-0.1%
Life expectancy from age 85	-	-	-	-0.2%	-0.2%	-0.1%
Life expectancy from age 90	-	-0.1%	-	-1.4%	-1.6%	-1.3%
Life expectancy from age 95	-	-0.4%	-0.1%	-1.3%	-2.4%	-2.1%

(b) Restated analysis (for comparison)

Trend Lexis Smoothing	No trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend 50/50 Lexis No smoothing	Trend Refined Lexis No smoothing	Trend Refined Lexis CMI smoothing
Join age	90	90	90	85	85	85
Parameter k	5	5	5	5	5	5
Parameter m	5	5	2	2	2	2
Life expectancy from age 65	-	-	-	-	-	-
Life expectancy from age 70	-	-	-	-	-	-
Life expectancy from age 75	-	-	-	-0.1%	-0.1%	-
Life expectancy from age 80	-	-	-	-0.1%	-0.1%	-0.1%
Life expectancy from age 85	-	-	-	-0.2%	-0.2%	-0.1%
Life expectancy from age 90	-	-0.1%	-0.1%	-1.4%	-1.7%	-1.3%
Life expectancy from age 95	-	-0.6%	-0.3%	-1.5%	-2.6%	-2.3%

Appendix 3 Lexis adjustments

Chart A3.1 shows a sample of the Lexis adjustments $D_U(x,t)$ / D(x,t), for males, covering ages 60 to 120 and calendar years 2000 to 2015, that result from the analysis described in Section 9. You can see that the assumption of a 50/50 split between Lexis triangles becomes less accurate as we move to the older ages, and there is a clear cohort pattern. Although other bodies (for example, the HMD) adopt a similar principle of applying Lexis adjustments, the precise adjustments used are likely to differ from those we have determined (due to differences in the detail of our approaches).

Chart A3.1: An illustration of Lexis adjustments DU(x,t) / D(x,t) for ages 60-120 and calendar years 2000-2015 (ONS E&W male data)

Age \ Year	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015
60	52%	52%	49%	50%	49%	54%	48%	50%	53%	52%	52%	51%	51%	50%	50%	51%
61	51%	52%	52%	49%	50%	49%	54%	48%	50%	53%	52%	52%	51%	51%	50%	51%
62	51%	51%	52%	52%	49%	50%	49%	54%	48%	50%	53%	52%	52%	51%	51%	50%
63	51%	51%	51%	52%	52%	49%	50%	49%	54%	48%	50%	53%	52%	52%	51%	51%
64	51%	51%	51%	51%	52%	52%	49%	50%	49%	54%	48%	50%	53%	52%	52%	51%
65	51%	51%	51%	51%	51%	52%	52%	49%	50%	49%	54%	48%	50%	53%	52%	52%
66	50%	51%	51%	51%	51%	51%	52%	52%	49%	50%	49%	54%	48%	50%	53%	52%
67	52%	50%	51%	51%	51%	51%	51%	52%	52%	49%	50%	49%	54%	48%	50%	53%
68	51%	52%	50%	51%	51%	51%	51%	51%	52%	52%	49%	50%	49%	54%	48%	50%
69	52%	51%	52%	50%	51%	51%	51%	51%	51%	52%	52%	49%	50%	49%	54%	48%
70	51%	52%	51%	52%	50%	51%	51%	51%	51%	51%	52%	52%	49%	50%	49%	54%
71	51%	51%	51%	51%	52%	50%	51%	51%	51%	51%	51%	52%	52%	49%	50%	49%
72	50%	51%	51%	51%	51%	52%	50%	51%	51%	51%	51%	51%	51%	51%	49%	50%
73	52%	50%	51%	51%	51%	51%	52%	50%	51%	51%	51%	51%	51%	51%	51%	49%
74	51%	52%	50%	51%	51%	51%	51%	52%	50%	51%	51%	51%	51%	51%	51%	51%
75	51%	51%	52%	50%	51%	51%	51%	51%	52%	50%	51%	51%	51%	51%	51%	51%
76	51%	51%	51%	52%	50%	51%	50%	51%	51%	52%	50%	51%	51%	51%	51%	51%
77	51%	51%	51%	51%	52%	50%	51%	50%	51%	51%	52%	50%	51%	51%	51%	51%
78	52%	50%	51%	51%	51%	51%	50%	51%	50%	51%	51%	51%	50%	51%	51%	51%
79	52%	52%	50%	51%	51%	51%	51%	50%	51%	50%	51%	51%	51%	50%	51%	51%
80	44%	52%	52%	50%	51%	51%	51%	51%	50%	51%	50%	51%	51%	51%	50%	51%
81	52%	44%	52%	52%	50%	51%	51%	51%	51%	50%	51%	50%	51%	51%	51%	50%
82	50%	51%	44%	52%	52%	50%	51%	50%	50%	51%	50%	50%	50%	51%	51%	51%
83	53%	50%	51%	44%	52%	51%	50%	51%	50%	50%	51%	49%	50%	50%	51%	51%
84	50%	53%	50%	51%	43%	52%	51%	50%	50%	50%	50%	51%	49%	50%	50%	51%
85	51%	50%	53%	49%	51%	43%	51%	51%	50%	50%	50%	50%	51%	49%	50%	50%
86	50%	51%	49%	53%	49%	51%	43%	51%	51%	50%	50%	50%	50%	51%	49%	50%
87	49%	49%	50%	49%	52%	49%	51%	43%	51%	51%	50%	50%	50%	50%	51%	49%
88	49%	49%	49%	50%	49%	52%	49%	50%	43%	51%	51%	49%	50%	50%	50%	50%
89	49%	49%	49%	49%	50%	49%	52%	49%	50%	43%	51%	51%	49%	50%	50%	50%
90	49%	49%	49%	48%	49%	50%	49%	52%	48%	50%	42%	50%	50%	49%	49%	49%
91	49%	49%	49%	48%	48%	49%	50%	48%	52%	48%	50%	42%	50%	50%	49%	49%
92	47%	49%	49%	48%	48%	48%	48%	49%	48%	51%	48%	50%	42%	50%	50%	49%
93	48%	47%	48%	48%	48%	48%	48%	48%	49%	48%	51%	48%	50%	42%	50%	50%
94	47%	48%	47%	48%	48%	48%	48%	48%	48%	49%	48%	51%	48%	49%	42%	50%
95	48%	47%	48%	47%	48%	48%	48%	48%	47%	48%	49%	48%	51%	47%	49%	41%
96	47%	47%	47%	47%	46%	47%	47%	47%	47%	47%	47%	48%	47%	50%	47%	49%
97	47%	47%	47%	46%	47%	46%	47%	47%	47%	47%	47%	47%	48%	47%	50%	47%
98	46%	46%	47%	47%	46%	47%	46%	47%	47%	47%	47%	46%	47%	48%	47%	50%
99	46%	46%	46%	46%	46%	46%	46%	45%	47%	47%	46%	46%	46%	46%	47%	46%
100	46%	46%	46%	46%	46%	46%	45%	46%	45%	46%	46%	46%	46%	46%	46%	47%
101	46%	46%	45%	45%	45%	46%	46%	45%	46%	45%	46%	46%	46%	46%	45%	46%
102	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%	44%	46%	46%	46%	45%	45%
103	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	44%	45%	45%	45%	45%
104	46%	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	44%	45%	45%	45%
105	44%	46%	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	44%	45%	45%
106	46%	44%	46%	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	44%	45%
107	45%	46%	44%	46%	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	44%
108	46%	45%	46%	44%	46%	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%
109	44%	46%	45%	46%	44%	46%	46%	45%	45%	45%	45%	45%	45%	45%	45%	45%
110	46%	44%	46%	45%	46%	44%	46%	46%	45%	45%	45%	45%	45%	45%	45%	45%
111	45%	46%	44%	46%	45%	46%	44%	46%	46%	45%	45%	45%	45%	45%	45%	45%
112	46%	45%	46%	44%	46%	45%	46%	44%	46%	46%	45%	45%	45%	45%	45%	45%
113	46%	46%	45%	46%	44%	46%	45%	46%	44%	46%	46%	45%	45%	45%	45%	45%
114	45%	46%	46%	45%	46%	44%	46%	45%	46%	44%	46%	46%	45%	45%	45%	45%
115	45%	45%	46%	46%	45%	46%	44%	46%	45%	46%	44%	46%	46%	45%	45%	45%
116	44%	45%	45%	46%	46%	45%	46%	44%	46%	45%	46%	44%	46%	46%	45%	45%
117	45%	44%	45%	45%	46%	46%	45%	46%	44%	46%	45%	46%	44%	46%	46%	45%
118	45%	45%	44%	45%	45%	46%	46%	45%	46%	44%	46%	45%	46%	44%	46%	46%
119	44%	45%	45%	44%	45%	45%	46%	46%	45%	46%	44%	46%	45%	46%	44%	46%
120	44%	44%	45%	45%	44%	45%	45%	46%	46%	45%	46%	44%	46%	45%	46%	44%

Appendix 4 Diagnostic tests – additional information

Impact diagnostics

The first question is whether a variant actually affects the population exposures to a material extent.

To address this, we show

- period life expectancy estimates by age (based on the raw average mortality rates implied for calendar years 2011-2015 inclusive) and
- population estimates by age (for calendar year 2015) relative to the current ONS approach.

The life expectancy figures in our main analysis are based on ungraduated mortality rates. This approach gives a better picture of the underlying effects introduced in the mortality grid (before these are smoothed out by graduation). We have truncated the calculation of the life expectancy figures to end at age 106, to avoid excessive volatility from sparse data at the very highest ages.

We have also shown the impact of our final proposal on life expectancy based on graduated mortality rates for 2011-2015, as a better guide to the impact users may see in practice. In addition, we present the results for the period 2006-2010 to assess the robustness of our proposal.

Performance diagnostics

We have prepared the following diagnostics as a guide to help determine whether variants perform better or worse than one another.

Final year balancing adjustment

This is the adjustment implied by the correction factor c in the KT method, i.e. equal to $c - 1$.

In general we prefer final year balancing adjustments close to zero, as these imply a lower discrepancy between the underlying KT method and the official population total to which it is constrained in the final year. 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.

Annual scaling adjustments

For each variant, and for each calendar year over the period 1972 to 2015, we have calculated the percentage adjustment required to scale the raw KT method exposures to match the official population total above the join age (in line with the constraint applied by the ONS). Our analysis shows the average magnitude of this adjustment across all years. We have used the average as a simple unit. We recognise that variants on this metric may also be informative but have not explored this further.

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.

There is also a converse element to this. For the earlier part of the history included in the average, the majority of cohorts in the high age population will now be extinct (or very nearly extinct). This means that the population estimates generated by the KT method are effectively being constructed purely from the deaths dataset (under the method of extinct generations) with very little dependence on the projection of survivor ratios. As such, they are almost entirely data driven and likely to be very accurate (given the evidence that death registration data in E&W is accurate[5]). Over this region, the size of the annual scaling adjustment is therefore more a diagnostic on the accuracy of the official total population estimate above the join age, to which we are constraining (insofar as it measures the agreement of this total to an accurate independent calculation derived from the deaths data).

A small average annual scaling adjustment therefore implies that the historic official total population figures above the join age must be relatively accurate based on the method of extinct generations. This is an important diagnostic when comparing alternative totals to use for the constraint (in particular, in analysing different choices of join age).

Cohort inconsistency metric

An internally consistent population should exhibit decrements from year to year which match the number of deaths reported. We have calculated a simple metric to quantify this.

For a given age x (last birthday) and calendar year t, we define the cohort inconsistency as:

$$\left(P_{x-1,t-1} - P_{x,t} - D\right)/P_{x,t}$$

where

- $P_{x,t}$ denotes the mid-year population estimate at age x (last birthday) for calendar year t, and
- D is an estimate of the number of deaths between mid-year $t-1$ and mid-year t (who were aged x last birthday at mid-year t), given by
  $$D = \left(3D_{x-1,t-1} + D_{x,t-1} + D_{x-1,t} + 3D_{x,t}\right)/8$$
  where $D_{x,t}$ denotes the number of deaths aged x (last birthday) during calendar year t.

We then calculate the average magnitude of this metric across all ages 71 to 104 inclusive, and all calendar years 1972 to 2015 inclusive. We prefer an average cohort inconsistency close to 0%.

The fundamental KT method should deliver cohort consistency (precisely because it constructs its population estimates from the deaths data) but this can be distorted by aspects of the implementation such as:

- conversion of death input data to 1 January timing, and of 1 January population estimates to mid-year estimates,
- constraining the KT estimates to official population totals for each year independently, and
- the transition between census-based estimates and KT estimates at the join age.

The cohort inconsistency metric should capture these kinds of distortion, although we note that the metric itself is not perfect, in particular because the death count D is an approximate estimate that does not allow for differing sizes of the various cohorts which contribute.

Smoothness of mortality across join age

We expect mortality rates for each calendar year to vary smoothly by age – in fact, we expect $\log m_{x,t}$ to be locally linear in x (where $m_{x,t}$ is the crude central mortality rate for age x and year t).

A particular concern is the smoothness in rates at the *join age*, where we transition from the ONS's census-based population estimates at younger ages to the (constrained) KT method at higher ages.

To measure this, we define the deviance in $\log m_{x,t}$ for age x and year t as

$$\frac{1}{3}\sum_{y=x-2}^{x+2}\left(\log m_{y,t} - \log m'_{y,t}\right)^2$$

where $\log m'_{y,t}$ is the estimate of $\log m_{y,t}$ derived by fitting a straight line in $\log m_{y,t}$ to the 5 age points $y = x-2$ to $y = x+2$.

We then calculate the average deviance in $\log m_{x,t}$ at the join age, across all calendar years 1984 to 2015 inclusive. An average deviance closer to 0% is preferred, as this implies a smoother transition of mortality rates across the join age.

Cairns Blake Dowd Kessler (CBDK) diagnostics

In their paper 'Phantoms never die', Cairns et al proposed a set of visual diagnostics to help identify potential data anomalies in the E&W exposure data. We have replicated some of the diagnostics from that paper and use them in our analysis to assess whether our variant methods are able to resolve the anomalies concerned.

A brief summary of each diagnostic (along with a replica of the chart for the official ONS exposures under the current estimation method) is provided below.

CBDK 1 – mortality rates by cohort

This is a plot of crude central mortality rates by age, for each of 5 adjacent birth cohorts. The test is whether the pattern of mortality rates by age looks similar for these successive cohorts. For example, Chart A4.1 shows this diagnostic applied to the official ONS male data for the 5 birth cohorts 1917 to 1921. We would expect a consistent (log-linear) upward trend in mortality by age across all 5 cohorts. However, for birth cohorts 1917 to 1921 there is a clear discrepancy, with the mortality curve for the 1919 cohort bending materially away from the other curves. This highlights a potential mismatch between the deaths and exposure estimates for these cohorts, making them particularly sensitive to the methods used to estimate the mid-year population figures (both in census years and in rolling forward between census years).

Chart A4.1: CBDK 1 – mortality rates by cohort (for official ONS male data)

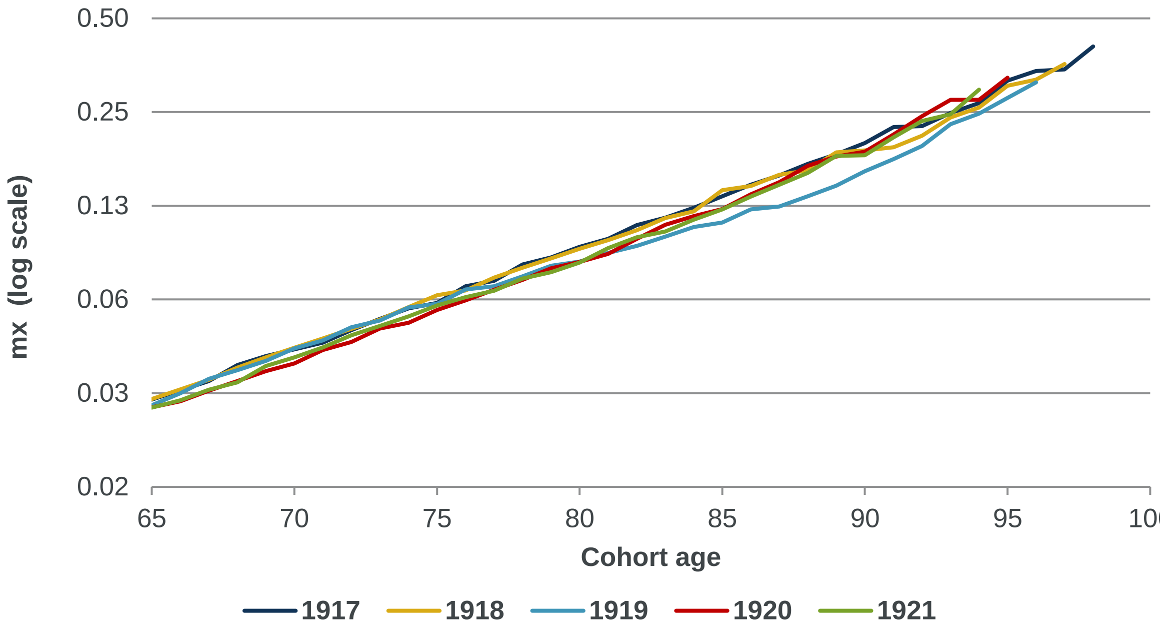

CBDK 2 – concavity by cohort

This is a plot of the empirical concavity function $C(x,t)$ by time for a given birth cohort, where

$$C(x,t) = \log m_{x,t} - \frac{1}{2}\left(\log m_{x-1,t} + \log m_{x+1,t}\right)$$

This function measures the log-linearity of mortality rates by age for a given calendar year t at age x. We expect $\log m_{x,t}$ to be locally linear, so the test is whether the concavity function stays close to zero, without any systemic bias. Chart A4.2 shows the result of this diagnostic for birth cohorts 1919, 1920, 1924 and 1947, based on the official ONS male data. The plot for the 1924 cohort is as expected for a dataset with no anomalies – there is apparently random variation in the concavity function around zero, with no systemic bias.

In contrast, the plots for 1919 ,1920 and 1947 show systematic bias above or below zero (for 1919 and 1920 in particular these effects are seen prior to around 1992). The gold dashed line in Chart A4.2 shows the “convexity adjustment ratio – 1”. 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. For 1919, 1920 and 1947, which are all years with an unusual pattern of births, it mostly explains the systematic bias seen in the empirical concavity function. However after 1992 the 1919 and 1920 birth years exhibit some drift in the empirical concavity function. Cairns et al noted that the method used to roll-forward the 2001 census day population estimates to the middle of 2001 did not allow for the unusual distribution of births for these cohorts. This could have resulted in inaccuracies in the 2001 mid-year population estimates which could also affect years back to 1992 and forward to 2010 due to the methods used to determine population estimates for non-census years.

Chart A4.2: CBDK 2 – concavity by cohort (for official ONS male data)

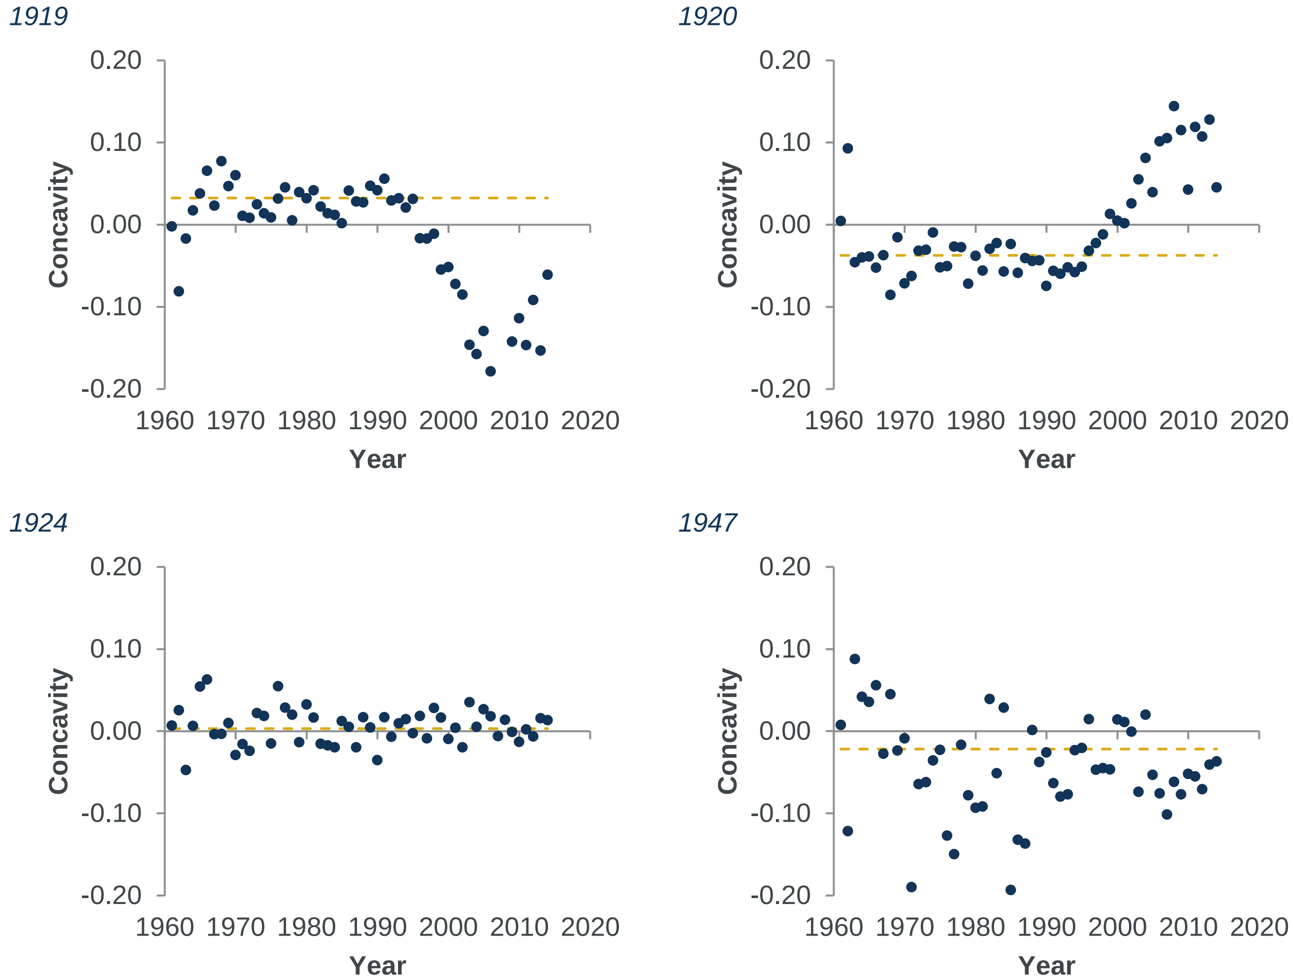

CBDK 2 – concavity heatmap

A related diagnostic is the presentation of the concavity function $C(x,t)$ as a two-dimensional heatmap by age x and calendar year t. This should provide a more comprehensive picture of any structural concavity in the dataset. The test is whether this heatmap simply shows random variation around a concavity value of zero, or whether it exhibits structural features indicative of anomalies in the data.

Chart A4.3 shows the plot for the official ONS male data. Anomalies are clearly suggested by the diagonal cohort patterns (again, 1919/20 stands out) and by the horizontal patterns around the join age of 90.

Chart A4.3: CBDK 2 – concavity heatmap (for official ONS male data)

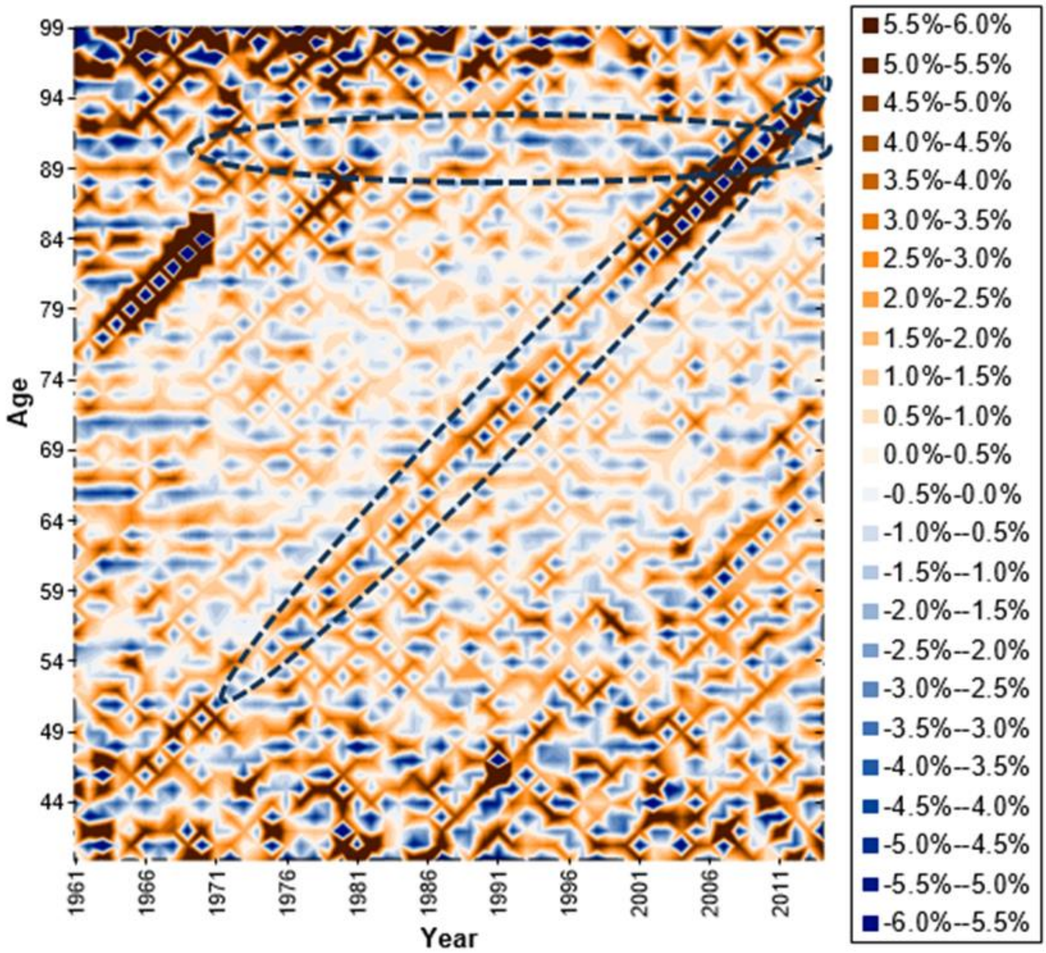

Synthetic population modelling

We can assess the accuracy of the KT method by using synthetic populations constructed for the age ranges under consideration. For these artificial datasets, assumed mortality and migration rates can be input and death counts determined in the format required for the KT method. The resulting single year of age population estimates derived using the KT method can then be compared against the known population figures to give an indication of the method's accuracy.

The key benefit of this approach is that it provides a means of testing the method's ability to estimate population figures correctly, for a range of hypothetical scenarios with potentially different features. A direct test of accuracy is, in contrast, not possible for the ONS dataset (because there is no reliable source of 'correct' figures against which to compare the KT estimates).

General approach

To carry out these comparisons, several synthetic populations were constructed for the period 1 January 1971 to 1 January 2015 using the following methodology:

- A starting population by age last birthday for ages 70 to 125 as at 1 January 1971 was assumed. This starting population was the same for all the synthetic populations derived.
- This population was then rolled forward to 1 January in subsequent years by applying assumed single year of age initial mortality rates ($q_{x,t}$) for each year, where x is age last birthday at 1 January in year t.
- The inflows in future years are the assumed numbers entering the population on 1 January each year at age 70 and net migration at older ages.

- In the synthetic populations modelled it was assumed that there was no net migration at any age during the period.
- Hence, the only movements assumed are the people entering the population on 1 January each year at age 70, and deaths.

Mortality assumed

As well as the same starting population, all the synthetic populations constructed had the same assumed age-specific mortality rates for 1971 (that is the probability of someone aged x at 1 January 1971 dying before reaching age $x+1$ at 1 January 1972 is the same in each model).

The mortality rates for 1972 and following years were then derived using various assumptions (specific to each synthetic population) for the improvements in mortality rates by age in years after 1971.

The resulting mortality rates were then applied to the relevant rolled-forward starting populations, together with the assumed number of persons entering at age 70 on 1 January in each future year, to project the resulting age-specific population at 1 January in each future year up to 1 January 2015.

Formulae

Let

- $P_{x,t}$ be the population estimate aged x last birthday at 1 January in year t;
- $q_{x,t}$ be the assumed probability that someone aged x last birthday at 1 January in year t dies before reaching age $x+1$ last birthday at 1 January in year $t+1$;
- $D_{x,t}$ be the number of deaths in year t aged x last birthday at 1 January in year t.

Then

$$P_{x+1,t+1} = P_{x,t} - D_{x,t}$$

and

$$D_{x,t} = P_{x,t} \times q_{x,t}$$

(with $D_{x,t}$ rounded to the nearest integer so that $P_{x,t}$ is an integer number for all x and t).

$q_{125,t}$ is assumed to be 1 in all years, so everyone aged 125 at the start of the year is assumed to die before reaching age 126 at the start of the following year.

$P_{70,t}$ are input assumptions.

Testing

The accuracy of the current KT method can then be tested by comparing the age-specific population estimates produced by the method (using deaths data from the synthetic populations) against the synthetic populations themselves, where of course the population figures are known to be 'correct'. This requires no caveats relating to the format of the underlying data (as all the population and deaths data are in the format required for the KT method) or assumptions on migration at the oldest ages. The synthetic populations can be used to assess the effectiveness of modifications applied to the KT method, to examine whether our alternative approaches may produce better results.

It should also be noted that the impact of rounding $D_{x,t}$ to the nearest integer has been tested and the impact of this is minor.

As well as comparing estimated vs actual populations (and estimated vs actual life expectancies) our analysis of the synthetic populations also includes a number of the consistency diagnostics outlined previously. We have excluded the CBDK diagnostics because the structure of the synthetic populations did not contain many of the idiosyncratic anomalies of the ONS data which these diagnostics were designed to pick up, as a result of which the plots turned out not to be very informative.

Although not covered in this paper, the synthetic populations could also be amended to allow for net migration and hence assess the sensitivity of the KT results to the levels of migration at older ages.

The analysis in this paper focuses on six synthetic populations, constructed from a range of alternative mortality improvement structures as set out in Section 4.3 of this paper.

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.