ISSN 2044-3145

Continuous Mortality Investigation

Institute and Faculty of Actuaries

High Age Mortality Working Party

WORKING PAPER 122

Final report of the High Age Mortality Working Party

June 2019

Summary

This is the final report of the CMI High Age Mortality Working Party (HAMWP). The Working Party was set up to investigate potential issues with mortality data at high ages, focussing principally on age 90 and above, and propose methods that other CMI committees could use in their investigations. The Working Party has so far published three working papers, summarised in Section 1.1. Following the publication of this paper, the Working Party will cease although the CMI will continue to consider whether further work in this area is appropriate.

Use of the Working Party's high age extension method by CMI investigation committees

The Working Party published its proposed approach to extending mortality tables to high ages in Working Paper 106. This work has so far been used or considered by the CMI's Self-administered Pension Schemes (SAPS), Annuities and Assurances Committees:

- The SAPS Committee used the method to produce high age extensions for the "S3" Series mortality tables. These were published in draft in Working Paper 107, and confirmed in Working Paper 113, following consultation with CMI Subscribers.
- The Annuities Committee tested the method as part of its indicative graduations in Working Paper 112. It found that the method produced reasonable results, after choosing appropriate parameters that differed from those used by SAPS, and it expects to adopt the method when it issues draft tables later this year.
- The Assurances Committee considered whether the method could have been applied to the "T08" dataset for term assurance data, with extensions being applied from around age 70. It was, however, unable to apply the method successfully, due to a lack of convergence between mortality rates for the "T08" dataset and the general population at age 70. The method was of course designed for application at much higher ages than 70 where mortality convergence is widely observed, so it is perhaps not surprising that it performed less well in this context.

The HAMWP method was designed for extensions that apply from high ages, around age 90. Work by CMI Committees has showed that the method works well in this case, but may not be appropriate for extensions from younger ages.

Analysis of data for large pension schemes

The need for a high age extension method, rather than simply relying on mortality experience data, is driven by weaknesses in typical datasets at high ages. For example, the SAPS "S3" dataset shows apparent falls in mortality at ages over 100. This pattern of mortality is considered implausible, as discussed in Working Paper 85 and in Working Paper 100, which set out strong independent evidence for continuing increases in mortality into the highest ages based on analysis of national population data.

The Working Party was keen to explore whether particular (large) pension schemes showed more plausible mortality patterns at high ages, and use this to inform views on high age mortality. To do this, it considered data collected for the SAPS investigation from eighteen large pension schemes, together with questionnaires on the data collection practices employed by eight of those schemes. It hoped that it would be able to identify a high-quality subset of the SAPS data that showed more plausible mortality patterns at high ages, and use this to inform views on high age mortality.

While it has been of interest to investigate the large schemes dataset, the analysis has unfortunately provided limited additional insight:

- The large schemes dataset, taken as a whole, does not seem to have better quality data than other schemes, as it shows a fall in mortality rates similar to the full SAPS dataset.
- The fall in mortality rates at high ages does not appear to be due to late-reported deaths, as the pattern is similar if we remove the final year's submission for each pension scheme.
- It also appears uncorrelated with the administration practices of contributing pension schemes, insofar as these can be judged from the subset of schemes that responded to the data questionnaire.
- The analysis does not give the Working Party reason to change the approach to high age extensions proposed in Working Paper 106.

Contents

Reliances and limitations

The purpose of this paper is to provide an update to interested parties on the Working Party's recent work. Anyone using this paper should exercise judgement on its suitability for their particular purpose.

The CMI aims to produce high-quality outputs and takes considerable care to ensure that the results and the accompanying documentation are accurate. However:

- We cannot guarantee their accuracy (see the Disclaimer on the last page of this document).
- There is a reliance on the underlying data, in particular the pension scheme data that is supplied at scheme level by actuarial consultancies. We carry out reasonableness checks on the data but ultimately rely on the data provider for its accuracy.
- We have applied judgement and assumptions in deciding how to illustrate mortality rates at high ages in the paper; for example, the method used to calculate five-year averages of mortality, and the comparative table for general population mortality.

1. Introduction

The CMI High Age Mortality Working Party (HAMWP) was set up to investigate potential issues with mortality data at high ages (focussing principally on age 90 and above), consider existing research on this subject, and propose methods that other CMI committees could use in their investigations.

The Working Party has so far published three working papers, summarised in Section 1.1.

This paper reports on:

- how CMI investigation committees have used the HAMWP method for extending mortality graduations to high ages; and
- analysis of data collected from large self-administered pension schemes.

Following the publication of this paper, the Working Party will cease although the CMI will continue to consider whether further work in this area is appropriate.

1.1 Previous work

Working Paper 85 described the Working Party's initial research into the features of high age mortality:

- There are various concerns with data quality at high ages for all data sources considered by the Working Party. It modelled the potential impact of late reporting of deaths and age mis-statement, which may explain some of the features observed in pension scheme data.
- Analysis of extinct cohorts indicated that population figures produced by the Office for National Statistics (ONS) for the over-90 age group had been overstated historically, leading to an understatement of mortality rates at those ages.
- Existing research on the level and shape of mortality at high ages is inconclusive, and different tables, published by the CMI and others, show wide variation.

Working Paper 100 described further research:

- It proposed a range of improvements that could be made to the Kannisto-Thatcher method used by the ONS to estimate population data at high ages.
- After further reviewing existing research and carrying out its own analysis, the Working Party considered that a mortality curve that makes allowance for deceleration at advanced ages is supportable for period mortality at population level, in particular for England & Wales.
- It described a framework for extending mortality graduations to high ages. Under this approach, mortality rates for the specific population are used at ages where the data is considered credible, and then converge towards mortality rates for the general population at high ages.

Working Paper 106 developed the work of Working Paper 100, and presented firm proposals for:

- alternative population estimates at high ages, adjusting the figures produced by the ONS; and
- extending mortality graduations to high ages, with reference to mortality rates for the general population.

1.2 Structure of this paper

Section 2 reports on how the high age extension method from Working Paper 106 has been applied to the CMI's SAPS, Annuities and Assurances datasets.

Section 3 describes the Working Party's analysis of data collected from large pension schemes.

1.3 TAS compliance

This paper is intended to provide an update to interested parties on the Working Party's recent work. This paper complies with the principles of the Financial Reporting Council's Technical Actuarial Standard "TAS 100: Principles for Technical Actuarial Work". Any person using this paper should exercise judgement over its suitability and relevance for their purpose

1.4 Acknowledgements

The members of the Working Party involved in the preparation of this paper were Steve Bale (Chair), Mark Cooper, Andrew Gaches, Adrian Gallop, Andy Harding and Richard Lamb.

2. Application of the high age extension method

As noted above, Working Paper 106 described and proposed a method for extending mortality rate graduations to high ages, with the intention that it would be used throughout the CMI. We refer to this as the "HAMWP method" in the remainder of this paper. In this section we describe how it has been used by the SAPS, Annuities and Assurances Committees to date.

2.1 Summary of the HAMWP method

Before describing specific applications of the HAMWP method, we briefly summarise it. The method is described in detail in Section 4 of Working Paper 106.

High age extensions are required when mortality rates for the specific dataset being graduated are considered unreliable above some age. In order to apply the HAMWP method, we require a set of "reference" mortality rates that are considered to be sufficiently reliable at high ages. The CMI uses UK general population mortality as its reference. Working Paper 106 describes how we adjust the ONS population figures for this dataset, and extend its mortality rates to high ages.

Once we have established the reference mortality rates and graduated the specific dataset for ages which are considered reliable:

- For ages up to a specified extension age (x_0), the mortality rates for the specific dataset are used without any adjustment.
- For ages above x_0, mortality rates converge asymptotically from the specific graduated rates towards the reference graduated rates. The shape of this convergence is controlled by two parameters, a convergence interval (in years) N and a rate of convergence (per convergence interval) c_N.

2.2 "S3" Series pension scheme tables

The first application of the HAMWP method was in Working Paper 107, which contained proposed graduations of the "S3" Series mortality tables, for SAPS data from 2009-2016. This was published alongside Working Paper 106.

The "S3" tables include tables which relate to specific amount bands, as well as "all-amounts" tables. Most of the all-amounts tables used an extension age of 95, although some tables used a lower value in order to avoid implausible results such as sharp kinks in mortality rates, or related tables crossing over.

The "S3" tables that relate to a specific amount band were extended with reference to the corresponding all-amounts table, rather than the UK general population, and typically used an extension age of 90. As for the all-amounts tables, some tables used a lower extension age to avoid implausible results.

The SAPS Committee considered the empirical rates of convergence between the various S3 tables and the general population. It decided to use the same convergence parameters for all extensions, with values of $N = 1$ and $c_N = 15\%$, reflecting the average rate of convergence across the various tables.

Working Paper 107 launched a consultation on the proposed "S3" Series rates, including two questions on the high age extension method. Some respondents asked for clarification on the use of a lower extension age for some tables in order to avoid implausible rates, but there was broad support for the HAMWP method, and there were no changes between the proposed and final rates.

2.3 Annuities

Working Paper 112 contained various analyses of the experience of pension annuities in payment (PAIP) from 2011-2014, including indicative graduations of that dataset. Section 3.10 of the paper used the HAMWP method to extend the indicative graduations to high ages.

The Annuities Committee initially considered using convergence parameters of $N = 1$ and $c_N = 15\%$ (as for the "S3" Series tables) and an extension age of 90. It found that these values led to "a distinct kink" in mortality rates

at the extension age. After considering actual convergence patterns in the PAIP dataset, it found that reducing c_N to 5% and/or reducing the extension age to 85 produced a smoother progression of mortality rates by age. The ability to vary the convergence parameters in the HAMWP method was intended to allow users of the method to reflect different rates of convergence in different datasets in this way.

The Committee notes that "it is unclear to us whether it is desirable to lower the age at which the extension starts if the data at these ages appears reliable" and it "would be helpful to understand whether potential users of the tables would be concerned with a kink in the rates."

Working Paper 112 states that the Annuities Committee expects to adopt the HAMWP method when it comes to produce formal graduations.

2.4 Assurances

The Assurances Committee has considered whether the HAMWP method could usefully be applied to the extension of term assurances data – specifically the "T08" graduations. Unlike the results in Sections 2.2 and 2.3, this analysis has not been published previously.

The T08 graduations used data up to age 75, reflecting the age range over which there were sufficient policyholders to have mortality experience data. The Committee considered extension ages of 68, 72 or 75, varying by gender and smoker status. It found that it was not possible to find convergence parameters which produced extensions that showed both:

- sufficient convergence between tables at high ages; and
- a smooth progression of mortality rates before and after the extension age.

A high value for the convergence parameter would lead to sufficient convergence but produce kinks in mortality rates, while a low value would avoid kinks but lead to an implausibly low level of convergence.

The Working Party considers "high age" as meaning age 90 and above, so it is perhaps not surprising that an extension method that was designed for those ages does not work so well when applied at younger ages at which rates can still be diverging from the reference population. The extension method may still be suitable for Assurances data relating to whole life mortality tables, but this has yet to be assessed.

2.5 Summary

The application of the high age extension method to SAPS, Annuities and Assurances data has shown that:

- The method works well for SAPS data, and has been met with approval from CMI Subscribers.
- The method can also work well for Annuities data, as long as the method's parameters are tailored to the specific dataset.
- The method relies on having convergence between the specific and reference mortality rates at the extension age. While this will typically be the case for genuine "high age" extension ages, it may not work well for lower extension ages where rates are still diverging, as for the Assurances T08 rates.

3. Analysis of data for large pension schemes

The Working Party has used mortality data collected from large UK pension schemes, together with information on their data collection practices, and has analysed these to see what insights it can draw. These large UK pension schemes are included within the SAPS dataset. In collecting this data, the Working Party is seeking to understand whether there is a higher-quality subset of the SAPS dataset, which could lead to a better understanding of high age mortality.

3.1 High age data in the "S3" Series dataset

The "S3" Series dataset consists of mortality data for a large number of UK pension schemes, of various sizes, for the period 2009-2016.

Chart 3A shows crude mortality rates for male and female pensioners on a lives basis (i.e. for the datasets underlying the S3PML and S3PFL tables) and an amounts basis (i.e. for the datasets underlying the S3PMA and S3PFA tables), each with 95% confidence intervals (CI), compared with English Life Table 17 (ELT).

Note that English Life Table 17 is lives-weighted and is based on data for 2010-12, so has an effective date of 2011. This differs from the effective date of 2013 for the HAMWP dataset. We have not adjusted ELT17 as we are only showing it as a broad indicator of the shape of mortality in the general population, and mortality improvements from 2011 to 2013 were small.

The chart shows that mortality rates for these datasets fall at very high ages on both a lives and an amounts basis. This is implausible and suggestive of data quality issues, given the strong independent evidence from other data sources that mortality rates into higher ages continue to increase – including the analysis in Section 3 of Working Paper 100, which shows that while mortality rates tend to decelerate at high ages (i.e. they increase more slowly than exponentially) we do not expect them to fall. An apparent fall in mortality rates is likely to be due to some deficiency in data quality, such as late-reported or unreported deaths, or mis-statement of some ages. These features were explored and modelled in Working Paper 85.

We would typically expect pensioners with higher pensions to experience lighter mortality than those with lower pensions, so that amounts-weighted mortality is lower than lives-weighted mortality. While that is the case for the lower ages shown, it does not hold at older ages.

Chart 3A: Five-year-average crude mortality rates for the S3PxL and S3PxA datasets

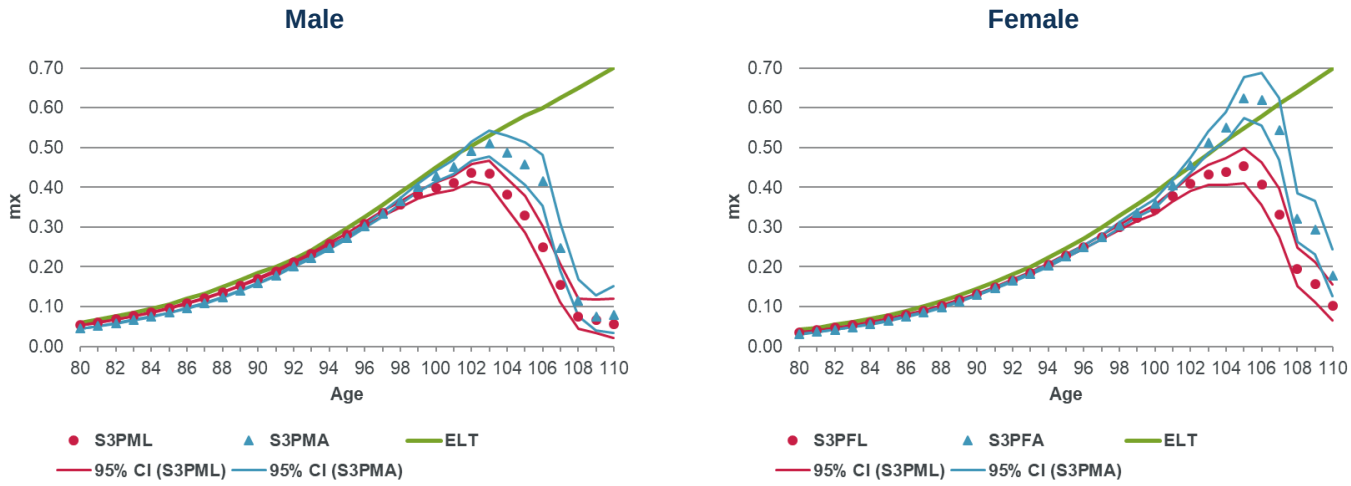

As a result of the unreliability of the SAPS dataset at high ages, it was not credible to derive high-age mortality rates directly from the data in the production of the "S3" Series tables, and it was instead necessary to apply high age extensions. These were based on the UK general population.

The Working Party sought to obtain a subset of the SAPS data that was of higher quality, in order to understand high age mortality in the SAPS dataset better. This could be used directly, to inform graduation of the SAPS dataset at higher ages, or indirectly, to inform wider views on high age mortality.

3.2 Collection of large scheme data

The CMI Secretariat identified large schemes in the SAPS dataset and engaged with key data contacts within data contributors to:

- seek permission to use these schemes' data in the HAMWP analysis; and
- identify whether there were any other large schemes for which data could be supplied.

The Working Party requested pensioner and dependant data from the schemes covering the period 1 January 2006 to 31 December 2015. Each scheme taking part in the data analysis exercise was also asked to complete a data questionnaire. This sought to collect information on administration practices including mortality tracing and treatment of suspended cases to allow consideration of the impact on data quality and to compare the practices of different schemes in the analysis.

We received data from 18 pension schemes[1] and received responses to the questionnaire for eight of them. We refer to these 18 schemes as the "HAMWP dataset".

Table 3.1 shows the exposure in person-years at ages 80 and above for the SAPS S3 dataset, the HAMWP dataset, and the subset of the HAMWP dataset relating to schemes that completed the questionnaire (denoted HAMWP_Q). We also show comparisons between them – e.g. "HAMWP ÷ S3" shows the size of the HAMWP dataset as a percentage of the S3 dataset.

Table 3.1: Relative sizes of the SAPS S3 and HAMWP datasets (see text for details)

Member type	Male pensioners	Female pensioners	Female dependants
SAPS S3 exposure	2,047,329	1,472,420	1,939,792
HAMWP exposure	1,005,007	798,157	962,882
HAMWP_Q exposure	412,267	109,661	444,103
HAMWP ÷ S3	49%	54%	50%
HAMWP_Q ÷ S3	20%	7%	23%
HAMWP_Q ÷ HAMWP	41%	14%	46%

The HAMWP dataset is around half the size of the SAPS S3 dataset for all three member types. Schemes that responded to the questionnaire have 41% of the exposure of the full HAMWP dataset for male pensioners, but only 14% of the exposure for female pensioners.

3.3 Questionnaire responses

There were a number of encouraging responses that suggested that the data included in the analysis should be of good quality:

- Data tracing
  - Data tracing is the process of actively checking the continued existence of every pension scheme member in receipt of a pension.
  - All schemes responded that data tracing is conducted on a monthly basis.

[1] Note that although we refer to "schemes" in this paper, one of the "schemes" is actually a combined submission for two pension schemes that the CMI is unable to distinguish between.

  - Some schemes responded that untraced pensioners are issued with Certificate of Existence letters whilst others responded that pensions are suspended to prompt communication from the member.

- Death trigger events
  - All schemes list notification from family as their most common trigger for the death verification process. The second most common trigger is as the result of data tracing.
  - Some schemes responded that deaths are recorded on the system once the verification process is complete (i.e. once the death certificate is received and processed), whilst others mark the record as a death at the point of notification.

- Processing deaths
  - All schemes responded that the actual date of death is recorded for deaths, rather than the date of notification.
  - The treatment of suspended cases recorded pending resolution differed between schemes. Five schemes record the pension as suspended until confirmation is received, however, three keep the record showing as in-force and one records them as deaths.

3.4 Methods

We focussed our analysis on crude mortality rates averaged over five consecutive years of age. We calculate this measure by first calculating the total exposure and numbers of deaths within the five-year age range, and then taking the quotient of these; i.e.:

$$\sum_{i=-2}^{i=+2} D_{x+i} \div \sum_{i=-2}^{i=+2} E_{x+i}$$

We considered, but rejected, other measures:

- Because of the low data volumes at high ages for individual pension schemes, crude mortality rates for single ages had wide confidence intervals which made it difficult to draw meaningful conclusions. The same was true for averages over a smaller number of years.
- Our chosen method allows confidence intervals to be calculated easily[2], but taking an average of (say) the crude mortality rate over five years as $\frac{1}{5}\sum_{i=-2}^{i=+2}(D_{x+i} \div E_{x+i})$ does not.

3.5 Datasets

We considered the following datasets:

- Individual pension schemes;
- “HAMWP” is the HAMWP dataset for all eighteen schemes;
- “HAMWP_Q” is a subset of the HAMWP dataset for the eight schemes for which we have questionnaire responses;
- “Good schemes” is a subset of the HAMWP dataset which we believe has good quality data. This is based on our assessment of five-year-average crude mortality rates for individual schemes by each pensioner type and dataset period (described in more detail below). We also considered judging "good schemes" by reference to the questionnaire responses, but we couldn't find any meaningful correlation between the responses and observed features of the data so this approach was not pursued.

For each of these datasets, we considered two variants:

- the “Full” variant including all data for the period 2009-2016; and
- a “Restricted” variant that excludes the final year of data, on a scheme-by-scheme basis, in order to mitigate the potential impact of late-reported deaths.

[2] We calculate the confidence intervals assuming a Poisson distribution for deaths. If deaths at individual ages have independent Poisson distributions, then the sum of deaths over five years of age also has a Poisson distribution.

We show results for male pensioners, female pensioners and female dependants. We do not show results for male dependants, as the low volume of data for this group makes it challenging to draw a meaningful conclusion.

All analysis has been done on a lives-weighted, rather than an amounts-weighted basis. As mortality rates for the full SAPS dataset showed similar implausible patterns at high ages on both lives-weighted and amounts-weighted bases (shown in Chart 3A), we did not see any merit in considering an amounts-weighted basis for the HAMWP dataset.

All datasets, and both variants, have been adjusted to have effective dates in 2013, consistent with the SAPS "S3" Series dataset. The method is described in Appendix 2 of Working Paper 107.

3.6 HAMWP and HAMWP_Q datasets

We first analysed crude mortality rates in both the HAMWP and HAMWP_Q datasets. Charts 3B-3D show:

- five-year average mortality rates and 95% confidence intervals for the Full and Restricted datasets;
- single-year crude mortality rates for the corresponding "S3" Series SAPS dataset; and
- English Life Table 17 (ELT17).

Across all three pensioner types, we see falls in mortality rates at high ages, similar to those seen in the corresponding "S3" dataset. The age at which the drop-off occurs is around age 102 for male pensioners and age 104 for female pensioner and dependants.

Excluding the final year of data for each scheme, in the "Restricted" variant, does not avoid a fall in mortality rates at high ages. This suggests that modest delays in reporting of deaths has not materially affected data quality.

Chart 3B: Five-year-average crude mortality rates for male pensioners

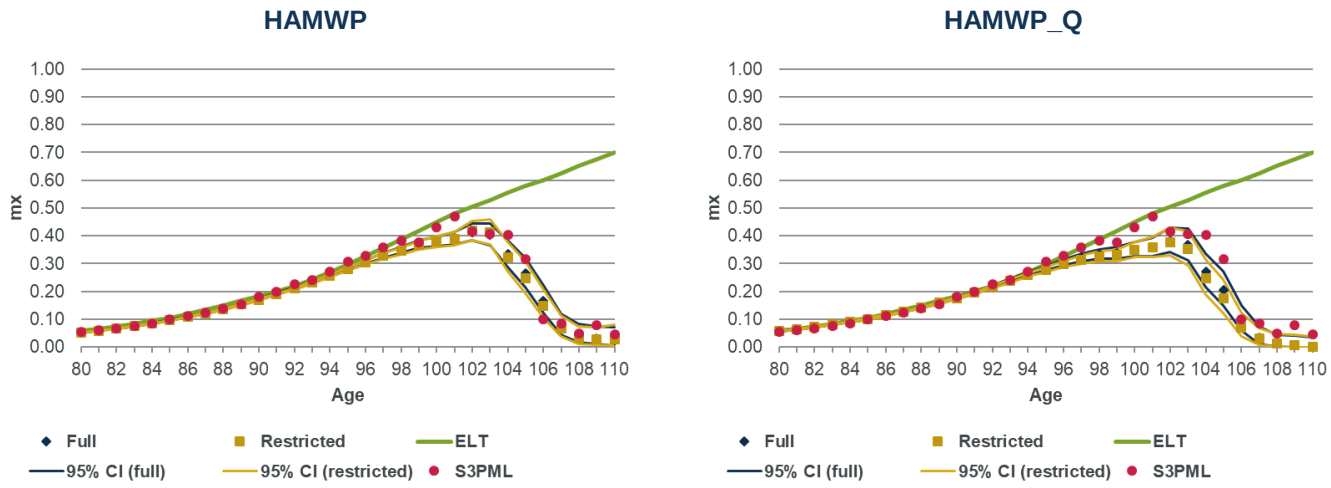

Chart 3C: Five-year-average crude mortality rates for female pensioners

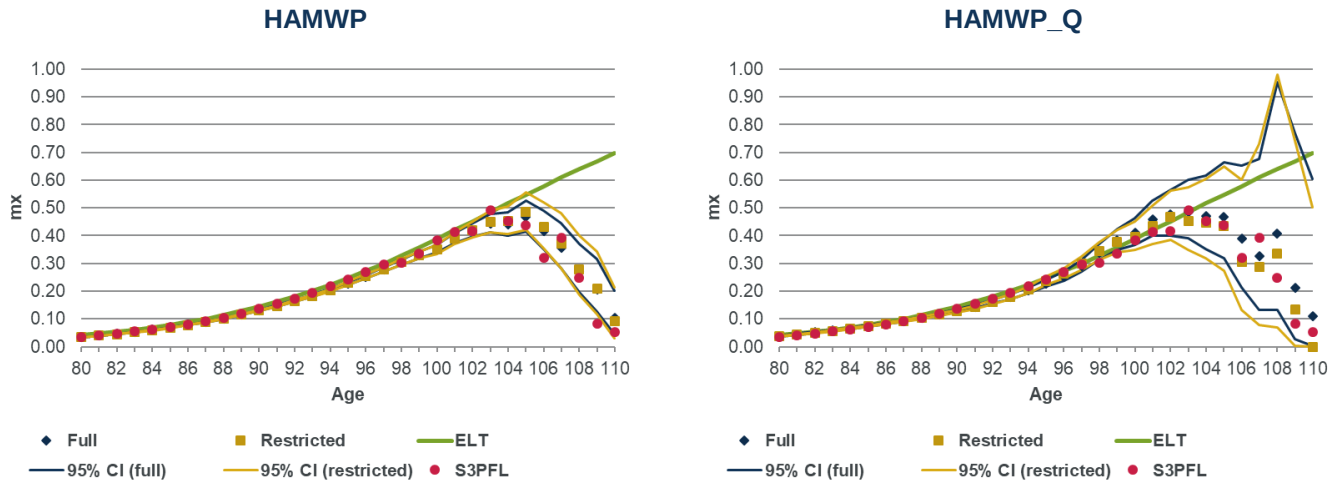

Chart 3D: Five-year-average crude mortality rates for female dependants

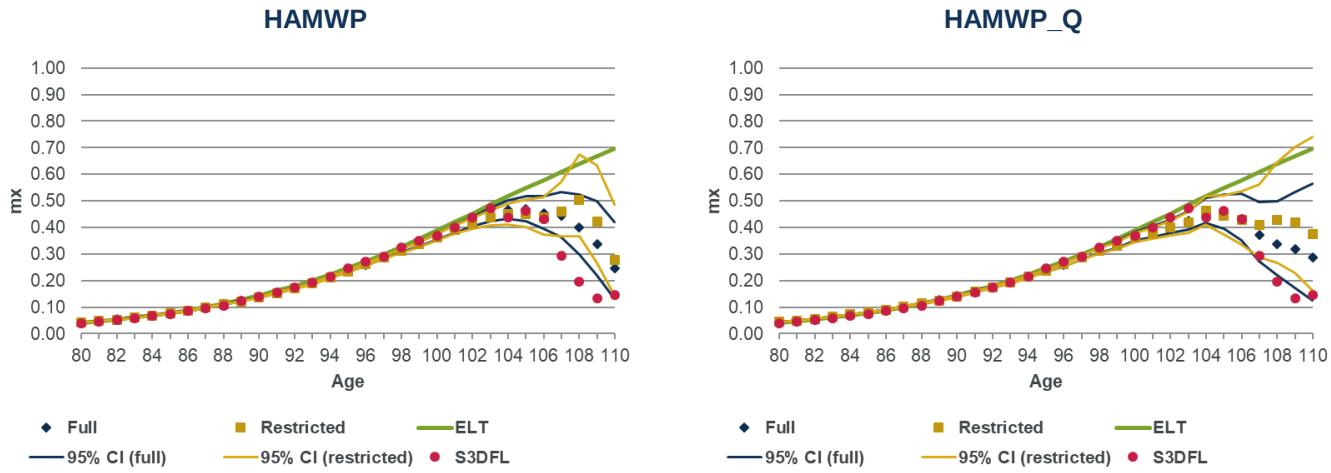

3.7 "Good schemes" combined

As noted above, we constructed a subset of "Good schemes" which show a more plausible pattern of mortality rates. This was done by considering five-year-average crude mortality rates for each combination of scheme and pensioner type (male pensioner, female pensioner and female dependant), and including each in the "Good schemes" dataset only if it showed a plausible pattern of mortality rates at high ages.

We consider mortality rates to be plausible if rates either continue to increase broadly exponentially with age, or show a modest deceleration; and implausible if they fall with increasing age, or decelerate sharply. This judgement is consistent with our findings in Working Papers 85 and 100. When making this assessment, we take account of the volume of data and the confidence intervals at each age, so that ages with very little data (e.g. ages 105 and above for most schemes) have little impact on our judgement.

We are conscious that a decision on the plausibility of mortality rates is somewhat subjective. To reduce the degree of subjectivity, two people first made their own independent assessments, before comparing results. In 46 of the 54 cases (18 schemes, each with three pensioner types) their initial assessments agreed, giving some comfort that many of the decisions were clear-cut.

Table 3.2 shows the number of schemes with each combination of Good or Bad data for the three pensioner types; e.g. there were four schemes where data was Good for all three types, and one scheme where data was Good for male and female pensioners but Bad for female dependants.

Table 3.2: Number of schemes with each combination of good or bad data by pensioner type

Male pensioners	Female pensioners	Female dependants	Number of schemes
Good	Good	Good	4
Good	Good	Bad	1
Good	Bad	Good	2
Good	Bad	Bad	2
Bad	Good	Good	2
Bad	Good	Bad	2
Bad	Bad	Good	3
Bad	Bad	Bad	2

If the difference in data quality was purely due to the data collection and processing practices of each pension scheme, then we might expect most schemes to have either Good data for all pensioner types, or Bad data for all pensioner types. However the numbers of schemes with each of the eight combinations in Table 3.2 are quite similar. This suggests that schemes do not have systematically Good or Bad data and there is a high degree of randomness in either whether they actually have Good or Bad data, or our ability to assess this. Furthermore, the attribution of Good or Bad status appeared to have little discernible correlation with the main variants in administration practice reported in the data questionnaire.

The Working Party discussed whether to assess Good data by considering each pensioner type independently, or whether to require a scheme to have plausible rates for all three pensioner types in order to be considered good. As shown in Table 3.2, taking the latter approach would have meant that we only included four schemes in the Good dataset. We instead considered each pension type independently. e.g. the Good dataset for male pensioners consists of the data for the nine schemes whose male pensioner data was considered Good

Charts 3E-3H show results for the Good datasets. The left-hand charts are in the same format as charts 3A-3D, while those on the right have a logarithmic scale to allow mortality deceleration to be seen more clearly.

Chart 3H shows results for all females, combining the datasets for female pensioners and female dependants.

We observe that:

- By construction, the “Good schemes” have a more plausible shape of mortality rates than the corresponding S3 datasets.
- Male and female pensioners have wide confidence intervals at the highest ages, whereas female dependants have narrower confidence intervals, reflecting higher data volumes at those ages.
- While there are some differences between “Full” and “Restricted” mortality rates at the highest ages, there is no clear pattern with “Restricted” rates being above “Full” for some pensioner types, and below for others, and the differences between them are within the confidence intervals.
- Mortality rates appear to flatten and fall at an earlier age for female dependants than for male or female pensioners.

As mentioned above, crude rates in the HAMWP and HAMWP_Q datasets start to fall from ages 102 to 104. Through our selection of “Good schemes”, we appear to gain confidence in crude rates for a few additional years for all pensioner types, although at this point confidence intervals widen materially, and mortality rates become erratic (as observed for male pensioners) or fall (as observed for female dependants).

Chart 3E: Five-year-average crude mortality rates for "Good" schemes for male pensioners

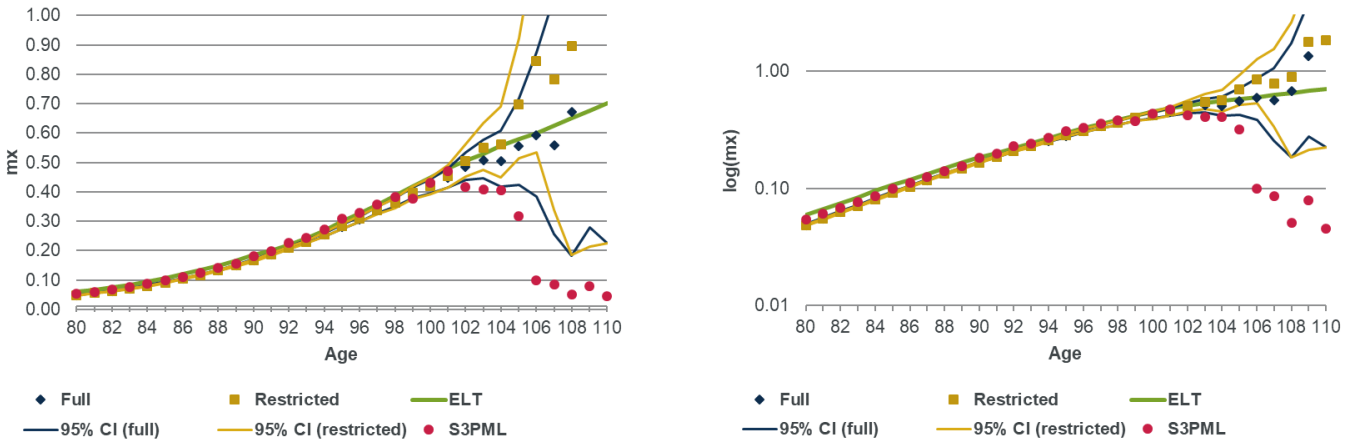

Chart 3F: Five-year-average crude mortality rates for "Good" schemes for female pensioners

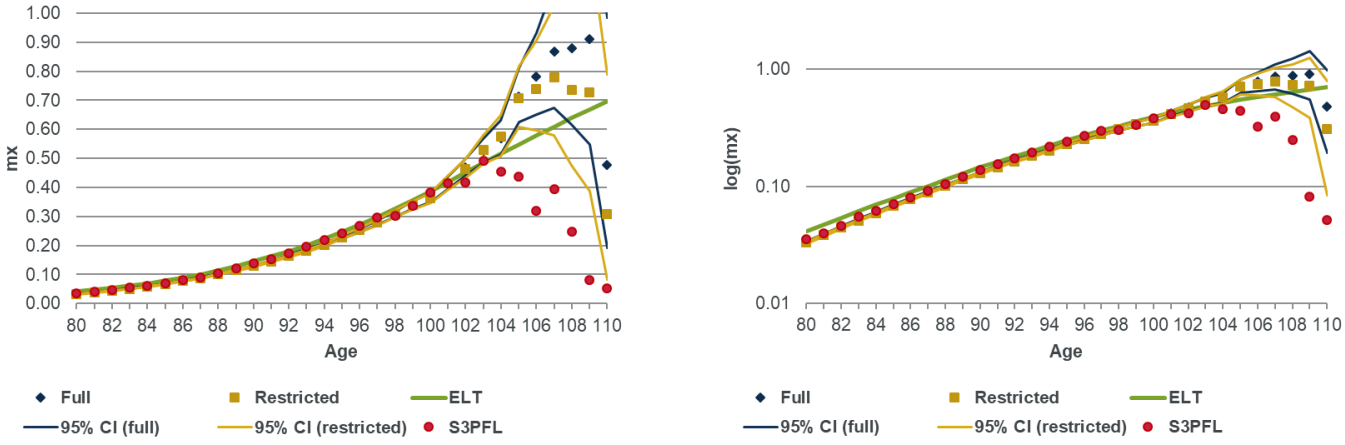

Chart 3G: Five-year-average crude mortality rates for "Good" schemes for female dependants

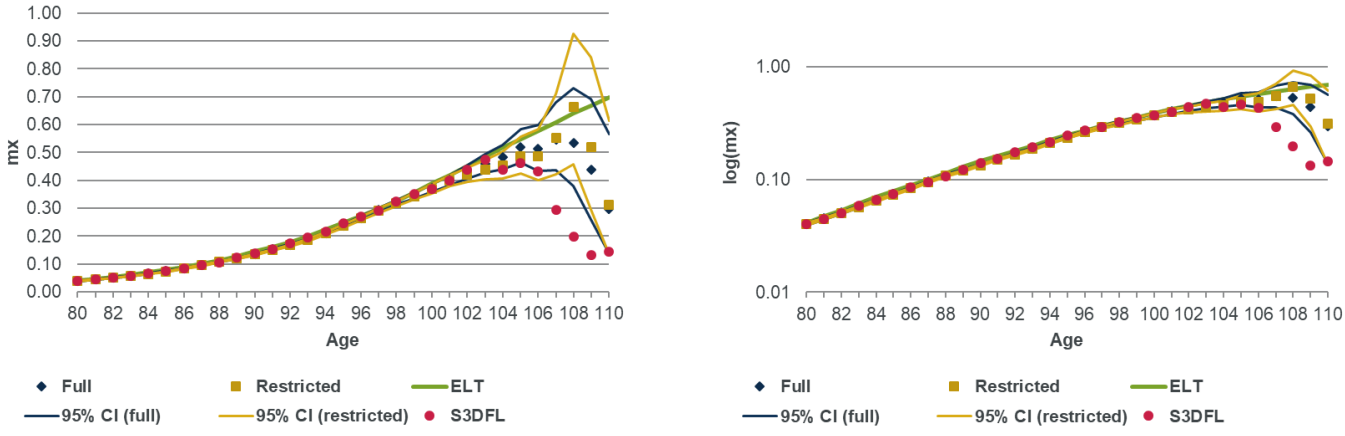

Chart 3H: Five-year-average crude mortality rates for "Good" schemes for all females

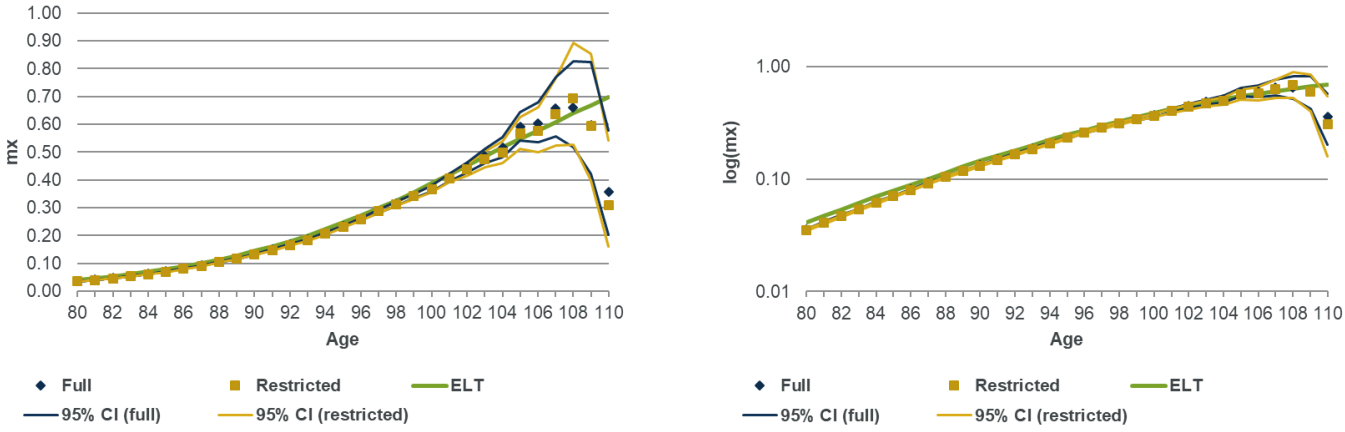

3.8 Discussion

The results appear to show a more plausible pattern of mortality rates for the Good datasets than for the full SAPS datasets or the combined HAMWP large schemes dataset, but we question whether the results are useful or merely "self-fulfilling".

It is perhaps unsurprising that a plausible shape of mortality rates can be achieved by filtering out datasets, and raises the question of whether we have successfully identified those schemes which have better quality data, or simply included the data that meets our preconceptions. If there was a clear differentiation in the quality of data between pension schemes then we would expect to see more schemes in the top and bottom rows of Table 3.2, and fewer in the middle.

Given the persistence of the falls in mortality observed at the very highest ages, even after restricting to Good datasets and attempting to manage late reporting of deaths, we did consider whether it could be a genuine effect. For example, a decline in aggregate mortality rates could *in theory* arise under pronounced survival bias within a heterogeneous population, even while mortality rates for each individual continue to increase (as discussed in Working Paper 85). However, we view this as unlikely given the credible independent evidence that aggregate population mortality rates do *not* decline in this way into the highest ages (examined in Working Paper 100). We expect it is more likely that the large schemes dataset is subject to residual data problems of the kind identified and analysed in Working Paper 85.

It is instructive to consider whether this analysis could have improved the CMI's most recently published mortality tables, the SAPS "S3" Series. Those tables were produced by graduating data up to age 95, and then producing high age extensions using the CMI's standard method, from Working Paper 106. The analysis in this paper does not suggest any change to rates up to age 95. For older ages, the evidence from the large schemes dataset does not feel sufficiently strong to rely on, although it is compatible with mortality rates in the reference table for the general population continuing to increase for ages 95+. It may be possible for others to gain greater insight into high-age mortality through the use of a different, higher-quality, dataset.

3.9 Summary

While it has been of interest to investigate the large schemes dataset, the analysis has provided little additional insight.

- The large schemes dataset, taken as a whole, does not seem to have better quality data than other schemes, as it shows a fall in mortality rates similar to the full SAPS dataset.
- The fall in mortality rates at high ages does not appear to be due to late-reported deaths, as the "Restricted" dataset showed similar results to the "Full" dataset.
- This analysis does not give us reason to change the approach to high age extensions proposed in Working Paper 106, and applied by CMI committees as discussed in Section 2.

References

Working Paper 85: “Initial report on the features of high age mortality” (2015)

Working Paper 100: “A second report on high age mortality” (2017)

Working Paper 106: “A proposed approach to closing off CMI mortality tables” (2018)

Working Paper 107: “Proposed graduations of the CMI SAPS 2009-2016 mortality experience based on data collected by 30 June 2017” (2018)

Working Paper 112: “Additional analyses of the experience of pension annuities in payment, 2011-2014” (2018)

Working Paper 113: “Final “S3” Series mortality tables” (2018)

These papers may be accessed and downloaded from the CMI section of the Institute and Faculty of Actuaries’ website: https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/other.

The three HAMWP papers, Working Papers 85, 100 and 106, are publicly available. Access to Working Papers 107, 112 and 113 is restricted to Authorised Users (i.e. to employees of subscribers; and to researchers, for non-commercial use).

This page is intentionally blank.

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: Two London Wall Place, 123 London Wall, London, EC2Y 5AU
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.