ISSN 2044-3145

Continuous Mortality Investigation

Institute and Faculty of Actuaries

Mortality Projections Committee

WORKING PAPER 90

CMI Mortality Projections Model consultation

August 2016

This Working Paper was originally issued in June 2016.

A version released in July made a number of changes relating to the timing of life expectancies calculated using the proposed method. The specific changes are to Charts 9E and 9F, Tables 9.1 and 9.2, and to the text below Table 9.2. We have also taken the opportunity to add comparisons at age 75 in Table 9.2.

This version has been updated to refer to Working Paper 91, now that it has been released, and to show the revised deadline of 30 September 2016 for responses to the consultation.

NOTE: This document is being made available publicly and its use is not subject to CMI's Terms and Conditions for Subscribers or Terms and Conditions for Academics and CMI Committee Members.

Executive Summary

The CMI Mortality Projections Committee has been critically reviewing the CMI Mortality Projections Model ("the Model") and proposes a number of changes. Subject to consultation, these changes would be made in the next version of the Model, CMI_2016, which is planned to be released in March 2017.

The elements of our proposal for CMI_2016 that are materially different from current practice are:

1. Calibrating the Model to data for the United Kingdom rather than England & Wales; subject to verifying the feasibility of this with the relevant national statistical bodies.

2. Simplifying the method used to adjust exposure data. The new method retains the broad principle of the previous method – that underlying mortality rates are smooth, and outliers are indicative of artefacts in the data.

3. Defining mortality improvements in terms of $\log m_{x,t}$ but with results from the Model still being expressed using the existing definition in terms of $q_{x,t}$.

4. Using a new Age-Period-Cohort Improvement (APCI) model to determine components of historical improvements. This means that:
    a. we can fit historical mortality rates and determine mortality improvements, both in aggregate and split into age-period and cohort components, in a single step;
    b. the fitting process is much quicker in terms of run-time; and
    c. we can implement the Model entirely within Microsoft Excel using Visual Basic for Applications (VBA), which makes it more accessible to users.

5. Removing the "step-back" from the edges of the data when determining historical improvements and, instead, requiring the Model itself to deal with the issue of stability.

6. Allowing and encouraging users to adjust the responsiveness of the Model to new data by using a single "period smoothing parameter".

7. Enabling users to express the pattern of convergence in terms of the slope of mortality improvements ("direction of travel") as an alternative to the current approach of proportion remaining at mid-point.

8. Tapering the long-term rate of age-period mortality improvements to zero between ages 85 and 110, rather than between ages 90 and 120.

9. Shortening cohort convergence periods for the youngest cohorts.

10. Removing the "Constant Addition to Mortality Improvements" parameter, which we think is little-used.

11. Adding an Intermediate layer of parameters to make it easier to set and communicate certain Advanced parameters.

Section 14 contains a list of consultation questions. We would like to receive responses to these from any interested parties by **30 September 2016**.

Contents

1. Introduction

The CMI published a new mortality projections model, CMI_2009, in November 2009. This is described in CMI Working Papers 38, 39 and 41 and was intended to:

- reflect the latest experience on trends in mortality;
- be relatively straightforward to understand, describe and use;
- allow users the flexibility to modify projections tailored to their own views and purpose; and
- be regularly updated over time to reflect emerging experience.

Since the publication of CMI_2009, the CMI Mortality Projections Model ("the Model") has been updated as each new year of data has become available. The structure of the Model has not changed, but default parameter values have been updated. Changes to parameters mainly reflect the incorporation of newly-published data, but some improvements were made to the method used to calibrate CMI_2014 and CMI_2015.

The Model has been successful – it is widely-used for analysis, disclosure, and benchmarking of mortality improvements; both in the UK and overseas. However expectations of mortality modelling are increasing, and there are some concerns with the current Model. The CMI Mortality Projections Committee has been critically reviewing the Model and considering whether any alternatives can offer sufficient advantages to justify a change.

The review has taken place during a period of unusual mortality experience in the UK. Aggregate mortality improvements since 2011 have been much lower than in the preceding decade, and it is unclear to what extent this is a short-term blip, or the start of a longer-term trend, and how the Model should respond.

This Working Paper summarises the Committee's investigations and its proposals for future versions of the Model and seeks the views of users. It is intended to be relatively concise and accessible. We will publish a further Working Paper (Working Paper 91), referred to in this paper as the "Technical Working Paper", during the consultation period. This will contain further technical detail and results for our proposals as well as other approaches that we considered but decided not to pursue. We also intend to release illustrative software that will allow users to test our proposed approach themselves.

1.1. Consultation timetable

The intended timetable from the release of this paper to the planned release of the next version of the Model is shown below. We refer to that version as CMI_2016 in this paper, but question 11.1 consults on its name.

29 June 2016	Public meeting at the Royal College of Physicians in Edinburgh
11 July 2016	Public meeting at Staple Inn in London
August 2016	Publication of the "Technical Working Paper" (Working Paper 91) and the illustrative software referred to above
30 September 2016	Deadline for responses to the consultation
November 2016	Publication of a further working paper, summarising responses to the consultation and any revisions to our proposals
March 2017	Next version of the Model, CMI_2016, published to CMI subscribers[1]

[1] Academics can gain free access, for non-commercial use, to the Model and other CMI outputs on application to the CMI.

1.2. Contents of this paper

This Working Paper is organised in the following sections:

1. This introduction.
2. The principles and objectives that we have considered throughout our review.
3. Description of the current Model and a number of concerns that have been raised about it.
4. Discussion of recent mortality, updating the analysis of Working Paper 83.
5. Considerations for the data used to calibrate the Model.
6. An alternative definition of mortality improvements that is used within our proposed model (although the current definition will be retained for final output).
7. An introduction to the Age-Period-Cohort Improvement (APCI) model, which we propose to use instead of the current approach to calibration.
8. Alternative approaches to projecting mortality improvements.
9. Results for our proposed model, including sensitivities to some parameters, and comparisons with the existing Model.
10. Proposals for a change to the Core parameters of the Model and a new layer of Intermediate parameters.
11. A naming convention for versions of the Model.
12. Brief discussion of alternative methods that we considered but decided not to pursue at this time.
13. A summary of our proposals.
14. A list of the consultation questions. (These are also listed at the relevant point during the paper.)

1.3. Acknowledgements

The members of the Mortality Projections Committee involved in the production of this Working Paper are Tim Gordon (Chair), Steve Bale, Piero Cocevar, Matthew Fletcher, Steven Rimmer, Neil Robjohns and Brian Sewell.

The Committee would like to thank:

- Jonathan Hughes, who was a Committee member during part of our review;
- Kishore Ananda and Simon Donnelly, who contributed to the Committee's investigations into alternative models;
- Kevin Armstrong, David Bartlett, Andrew Cairns, Deborah Cooper, Adrian Gallop, Bob Howard, Jonathan Hughes (again), Rob Kairis, Larry Pinzur and Richard Willets, who commented on our proposals at an earlier stage; and
- the Secretariat for its work in putting this paper together; in particular, Jon Palin, who contributed a large component of the intellectual content, implemented the modelling, wrote the software, and without whom we simply would not have been able to deliver this paper in the time available.

2. Principles

Before describing details of the current Model and considering alternatives, we discuss the way in which the Model is used, and the principles that have informed our investigations and decisions.

2.1 Pragmatism

Having looked at more statistically-based approaches, we have retained the pragmatic approach of the original version of the Model on the grounds that:

- by its very nature, past mortality data has limitations when used to predict future mortality improvements (because the past drivers of improvement will necessarily fall away);
- there is no single "correct" model so we expect and encourage users to adjust the Model by incorporating their own views on both short- and long-term future mortality improvements; and
- the Model covers improvements over a wide age range with complex patterns of past improvement, which are difficult to capture in a single statistical model.

2.2 Plausibility

When considering alternative models we have preferred those which produce mortality rates, and components, which show plausible features. For example, we have rejected some models which have cohort components that seem implausible to us, and some which appear to use cohort parameters to model age or period features.

2.3 Rates versus improvements

The current Model is a model of mortality improvements rather than mortality rates. This is helpful as the improvements can be applied to any base table of the user's choosing. However most other (stochastic) mortality models tend to model mortality rates, and it would be useful to have a model that could consistently consider mortality rates and improvements. Such a model would be consistent with the existing Model, but also more readily comparable with other mortality models.

2.4 Default parameters

The Model does not currently give a default answer, as users of the Model are required to specify at least the long-term rate of mortality improvements. We think that it is helpful that there is only one parameter that users have to specify. This means that it is obvious which of two models that use the Core parameters has the higher level of improvements and liability values.

2.5 Recent mortality and responsiveness

As discussed in Section 4, the national mortality experience of England & Wales since 2011 has seen significantly lower mortality improvements than the preceding decade. At this stage it is unclear to what extent this may be a change in medium- or long-term trends rather than just a short-term blip, and we encourage users of the Model to consider this point themselves and assess whether the default parameterisation of the Model is in line with their beliefs.

Current projections from the Model could be inaccurate in two distinct ways:

- a projection could "over-react" – responding promptly to a perceived trend which is not real, when it might have been preferable to be more stable; or
- a projection could "under-react" – being too slow to respond to a genuine new trend, when it might have been preferable to respond faster.

Users of the Model may have different views on recent mortality experience, and different preferences on how it should respond; perhaps depending on whether under-reacting or over-reacting is a bigger risk for them. For

example, trustees of a closed defined-benefit pension fund may value stability more highly than the management of an insurer that is currently selling annuities in a competitive market.

2.6 Complexity

Users may also have different views on how detailed and complex the Model should be. For some, the Model is what they use to set assumptions. For others it is merely a common language used to communicate assumptions that are set in some other way. We aim to strike a balance between a model that is detailed enough to capture important features of mortality, and one that is simple enough to be understood and communicated, including to lay trustees and senior colleagues from other professional backgrounds.

2.7 Justification for changes

We are aware that many users see value in stability between versions of the Model, and we have only proposed a change where we believe it offers a significant advantage over the current approach. As such, our recommendation draws on the existing Model, rather than starting from scratch.

2.8 Transparency

We intend to make the entire process of the Model, from initial data to final results, open, transparent and accessible to Subscribers. We also intend to make software (written in Excel) available so that Subscribers can replicate results or easily apply the Model to alternative datasets, with an initial version to be released alongside the Technical Working Paper during the consultation period.

3. The current Model

This section describes the current Model and a number of concerns that have been raised about it.

3.1 How the Model works

The current Model is primarily a model of mortality improvements rather than mortality rates. Users of the Model can apply mortality improvements from the Model to their own choice of base mortality table to obtain projected mortality rates.

The basic approach of the current Model is to project rates of mortality improvement by interpolating between current rates, which are estimated from historical data, and assumed long-term rates, which are set by users of the Model. This process is carried out separately for age-period and cohort components, and these are summed to give the overall mortality improvements. The Model has a default ("Core") set of parameters, but users can also choose to use their own ("Advanced") parameters.

There has been broad support for this method in the CMI's previous consultations (reported in Working Papers 41 and 69). It is seen as intuitive, striking a good balance between flexibility and simplicity; and the use of a single key long-term rate parameter provides a "common currency" for comparing projections.

There are a number of steps in the implementation of the current Model to derive projected mortality improvements from population mortality data. It is helpful to think of these in two broad phases: "Calibration" and "Projection". The Projection phase is contained in the published Excel workbook, while the Calibration phase provides the parameters that it requires.

The steps in these processes are summarised below, and further detail can be found in the user guides that accompany the release of each Model. The specific calendar years mentioned below relate to the most recent version of the model, CMI_2015.

Calibration

- Fit a penalised spline (p-spline) model of central mortality rates (m_x) to death and exposure data for ages 18-102 and calendar years 1975-2015. (For CMI_2015 the data for 2015 was based on actual data to 31 July and an estimate for the remainder of the year.) This approach requires choices of the order of the splines; their knot spacing and positions; the penalty direction (age-cohort rather than age-period); and the order of penalty functions. The fitting is based on a Poisson model. We determine the degree of smoothing by optimising the Quasi-Bayesian Information Criterion (QBIC), which includes an assumption regarding overdispersion.
- Convert m_x rates to q_x rates, and calculate mortality improvements.
- Use an age-period-cohort (APC) model to split the mortality improvements into smooth age, period and cohort components. These components are constructed from basis splines, the fitting criterion is unweighted least-squares, and we impose identifiability constraints to obtain a unique solution.
- Calculate the "residual" – the difference between fitted mortality improvements in 2012 for ages 20-100 (a "step-back" from the edges of the data) and the sum of the corresponding age, period and cohort components.
- Calculate combined age-period and cohort components in 2012, splitting the residual between them.

Having done this, we have age-period and cohort components that can be used in the Projection phase.

Projection

- The inputs from the Calibration phase are the historical improvements to 2012, and the "initial improvements" – a split of the improvements in 2012 into age-period and cohort components.
- The age-period and cohort components are projected separately using the same form of convergence function to match desired long-term rates at and beyond the end of their convergence periods.

- Under the standard "Core" assumptions:
  - The long-term rate is nil for cohort components.
  - The long-term rate for age-period components is flat up to age 90, and then tapers linearly to nil at age 120. The rate below age 90 has to be specified by a user of the Model; this is the one assumption for which no default value is provided.
  - The convergence periods for age-period and cohort vary according to age in 2012.
  - The mortality improvement halfway through the convergence period is the mean of the initial improvement and the long-term rate.
- The age-period and cohort components of mortality improvements are then summed to give the overall mortality improvement for each year from 2013 onwards.
- Users can use these projected mortality improvements directly, although the Model also enables the improvements to be used with base tables to calculate annuities and life expectancies.

As a matter of policy, the Model is not fully specified – users must, as a minimum, select a value for the long-term rate, the intention being to encourage user engagement in determining future longevity improvement assumptions. Users do not need to consider the full Calibration process if they are happy to use the Core assumptions for initial improvements.

3.2 Sensitivity to assumptions

Working Paper 84, which accompanied the publication of CMI_2015, shows the sensitivity of results to certain assumptions. We reproduce some of these sensitivities in Table 3.1 using annuity values for a 65-year old male calculated using S2PMA and CMI_2015 with Core assumptions, with a net discount rate of 3% p.a.

Table 3.1: Sensitivities to assumptions

	Long-term rate of 1%		Long-term rate of 2%	
Base case	16.00	-	16.48	-
Initial age-period improvements 1% lower	15.68	−2.0%	16.15	−2.0%
Initial age-period improvements 1% higher	16.32	+2.0%	16.81	+2.0%
Initial cohort improvements 1% lower	15.50	−3.1%	15.96	−3.2%
Initial cohort improvements 1% higher	16.50	+3.1%	16.99	+3.1%
Proportion remaining at mid-point of 0%	15.88	−0.8%	16.55	+0.4%
Proportion remaining at mid-point of 100%	16.12	+0.8%	16.41	−0.4%
Long-term rate 1% lower	15.57	−2.7%	16.00	−2.9%
Long-term rate 1% higher	16.48	+3.0%	17.01	+3.2%

Table 3.1 shows that the impacts of changing the initial and long-term rates of improvements are broadly similar. Choices concerning the shape of convergence (i.e. the proportion remaining at mid-point) are less material in cases where the initial and long-term rates are similar. Indeed if the initial and long-term rates were identical then the convergence assumption would have no impact.

3.3 Concerns

In this section we describe a number of concerns that have been raised about the current Model. Some are specific to the Model while some are more general. They are taken from a number of sources, including previous

consultations, public meetings, and the Committee's own views. In Section 12 we discuss how our proposals address some of these concerns.

Data

1. The Core Model is calibrated to data for the general population. In practice, the outputs of the Model are typically applied to subsets of the population (e.g. annuitants and pension scheme members) that may experience different levels of mortality improvements. This creates socio-economic basis risk.
2. There is also geographical basis risk as the Model is calibrated to data for England & Wales, but may be applied to people in other countries, or to a local region.
3. Concerns have been raised about the quality of the Office for National Statistics ("ONS") data used to calibrate the Model. The CMI has taken steps to adjust the exposure data but these do not completely solve the problem. There is particular concern about the accuracy of ONS exposure data at high ages, particularly ages 90 and above.

Calibration

4. The use of age-cohort penalties in the p-spline model, rather than age-period or combined age-period-cohort penalties, may exaggerate the strength of the cohort terms.
5. It is easier to test the goodness-of-fit to historical data than to test the predictive power of the Model. There is a risk, as with any model, that the Model may over-fit to the data and this may affect the quality of its projections.
6. There is potential for inconsistency between the shapes of typical base tables (including those produced by other CMI committees) and the more complex shapes of the mortality projections.
7. The decomposition of mortality improvements into age, period and cohort terms has a large residual item, relative to the two-dimensional p-spline fit, that has become more material in recent years due to high mortality improvements at young adult ages.
8. The split of improvements between age, period, and cohort terms (including reallocation of the residual) requires some subjective judgements.
9. The calibration process makes no explicit allowance for annual volatility. The objective function assumes Poisson variation only, with no tendency for a particular year to have light or heavy mortality across a range of ages. In practice, we see that year-to-year variation is larger than Poisson noise.
10. Our sense is that the majority view among users is that the Model may be too responsive – that is, the size of changes in Model outputs between successive versions because of new data is too great. In particular, some users would like to be able to adjust the Model's responsiveness.

Projection

11. Under the Core assumptions for convergence, age-period and cohort components of mortality improvements are initially level and then converge monotonically to their long-term rates. This can be inconsistent with historical data. Some users would like the Model to allow for the "direction of travel" of mortality improvements, so that historical and projected improvements have a continuous slope.
12. The convergence periods for age-period and cohort components are subjective and they differ. While this difference was justified in Working Paper 39 based on analysis of historical mortality improvements, the difference between them means that results are sensitive to the split of initial improvements into age-period and cohort components, which is itself subjective. If the convergence periods were the same then this sensitivity would be removed.
13. The Model is inconsistent between versions in the sense that even if mortality experience from year to year were as expected, the Model's projections would change. For example, the cohort component of mortality improvements for the 1960 cohort in CMI_2010 is assumed to converge to the long-term rate over 40 years, by 2047. However the convergence period for the same cohort in CMI_2015 would also be 40 years, so convergence would be assumed to be complete in 2052. The age-period convergence period is similarly reset with each version of the Model.

Inconsistency between Calibration and Projection

14. The objective function for the initial p-spline model is penalised log-likelihood, and so places more weight on ages and years with more deaths. However the age-period-cohort model uses a sum of squares approach, which gives equal weight to all ages and years.
15. The "step-back" in the Calibration phase means that there is a difference between the projected rates and the p-spline model for the latest years of the calibration period.
16. The Calibration phase assumes that age, period and cohort effects are constant during the period of the calibration data; but these effects are varied in the Projection phase.

Other

17. Although the parameters can be varied to reflect different views, doing so typically requires setting a vector of parameters (e.g. age-period convergence periods for all 131 ages from 20-150). It would be helpful to be able to have an intermediate level of control (e.g. to scale up or down all convergence periods with one parameter which is widely understood by users).
18. The entire process of calibrating and projecting the model, from data to projected improvements, requires several pieces of software, written in R and Excel, and is not easy for users to replicate.

Q 3.1. Do you have any concerns about the current Model that are neither mentioned in Section 3.3 nor addressed by our proposals?

4. Recent mortality

Working Paper 83 analysed and discussed mortality in England & Wales between 1975 and 2015, with a focus on the unusually low mortality improvements since 2011. That paper used data to 31 July 2015, with an estimate for the remainder of 2015. In this section we update the analysis to use data to 31 December 2015 for calendar-year analysis, and to 3 June 2016 for analysis of weekly deaths, using the same methods as in Working Paper 83. We find that both one-year and four-year mortality improvements to 2015 are lower than at any other time in the 41-year data period used to calibrate the Model.

4.1 Annual analysis

Chart 4A plots the standardised mortality ratio (SMR) since 2000 for ages 18-102, together with an extrapolation of a trend fitted to 2000 to 2011. It shows that mortality fell fairly steadily between 2000 and 2011, always being close to the trend over that period. However improvements have slowed considerably since then and mortality in 2015 was at a similar level to that in 2011, 10% above the projected trend.

Chart 4B plots four-year average mortality improvements derived from the SMR for three age ranges and their aggregate. The period 2011-2015 has had the lowest improvement of all the periods for all of the age bands shown. While the fall in mortality improvements has been seen across a wide age range, it has affected the oldest ages the most. This may be a challenge for some mortality models, as the recent experience is not just a simple period effect.

Chart 4A: Standardised mortality ratio (SMR) compared to 2000-2011 trend

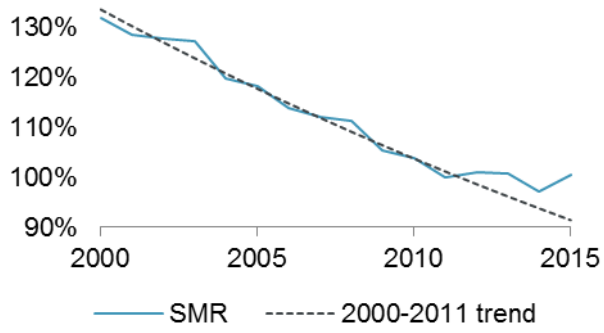

Chart 4B: Annualised four-year average mortality improvements for different age ranges

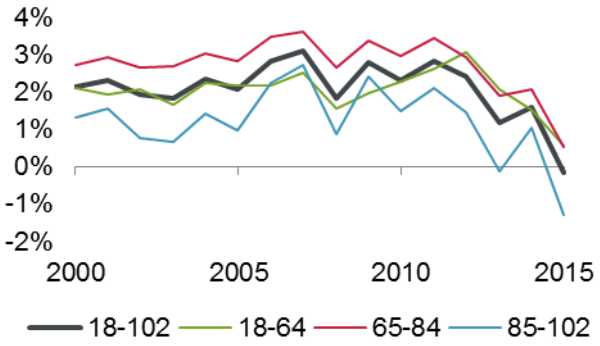

Both of these charts show a similar picture to the corresponding charts in Working Paper 83 (Charts 2 and 6), although the mortality improvements are even lower now that we allow for actual data to the end of 2015. As noted above, both the one-year and four-year mortality improvements to 2015 are lower than at any other time in the 41-year data period used to calibrate the Model.

4.2 Blip or trend?

We can think of changes in observed crude mortality rates as a combination of a number of factors including:

- long-term, typically smooth, underlying influences (e.g. lifestyle, medical and economic trends),
- short-term variations (e.g. cold winters and infectious diseases),
- artefacts of the data (e.g. inaccurate population estimates) that do not represent true mortality events, and
- random (Poisson) variation, particularly in small populations.

For the Model we are particularly interested in smooth long-term underlying changes to mortality rates and improvements, rather than short-term effects. We would like our model to deal with short-term blips without

masking or removing changes in longer-term trends; however it is difficult to tell from the annual data whether recent experience is due to short term blips, or a change in the longer-term trend. Analysis of patterns of mortality within years, as in Section 4.3, may help to address this.

As there is considerable uncertainty over the current level of underlying mortality improvements, in Working Paper 84 the Committee advised users of the CMI_2015 version of the Model to "consider carefully the appropriateness of the Core parameters and whether they should be amended to reflect users' own views and needs".

4.3 Seasonal analysis

Chart 4C is an extended version of Charts 11 and 12 of Working Paper 83. The solid blue line shows the seasonal variation in the smoothed standardised mortality ratio, compared to the trend for 2000-2011. The dashed grey line shows the seasonal average – the average of the corresponding weeks of the year from 2000 to 2011 – and is identical within each year. For example, the dashed grey line shows that smoothed mortality in early-January is typically 17% above the 2000-2011 trend, and smoothed mortality in the summer is typically 10% below the trend. January 2000 was exceptional, with actual smoothed mortality (the solid blue line) being 35% above the trend. Individual weeks have been more extreme, either because of genuine mortality events or artefacts of recording data.

Chart 4C: Seasonal variations in smoothed standardised mortality ratio; compared with 2000-2011 trend

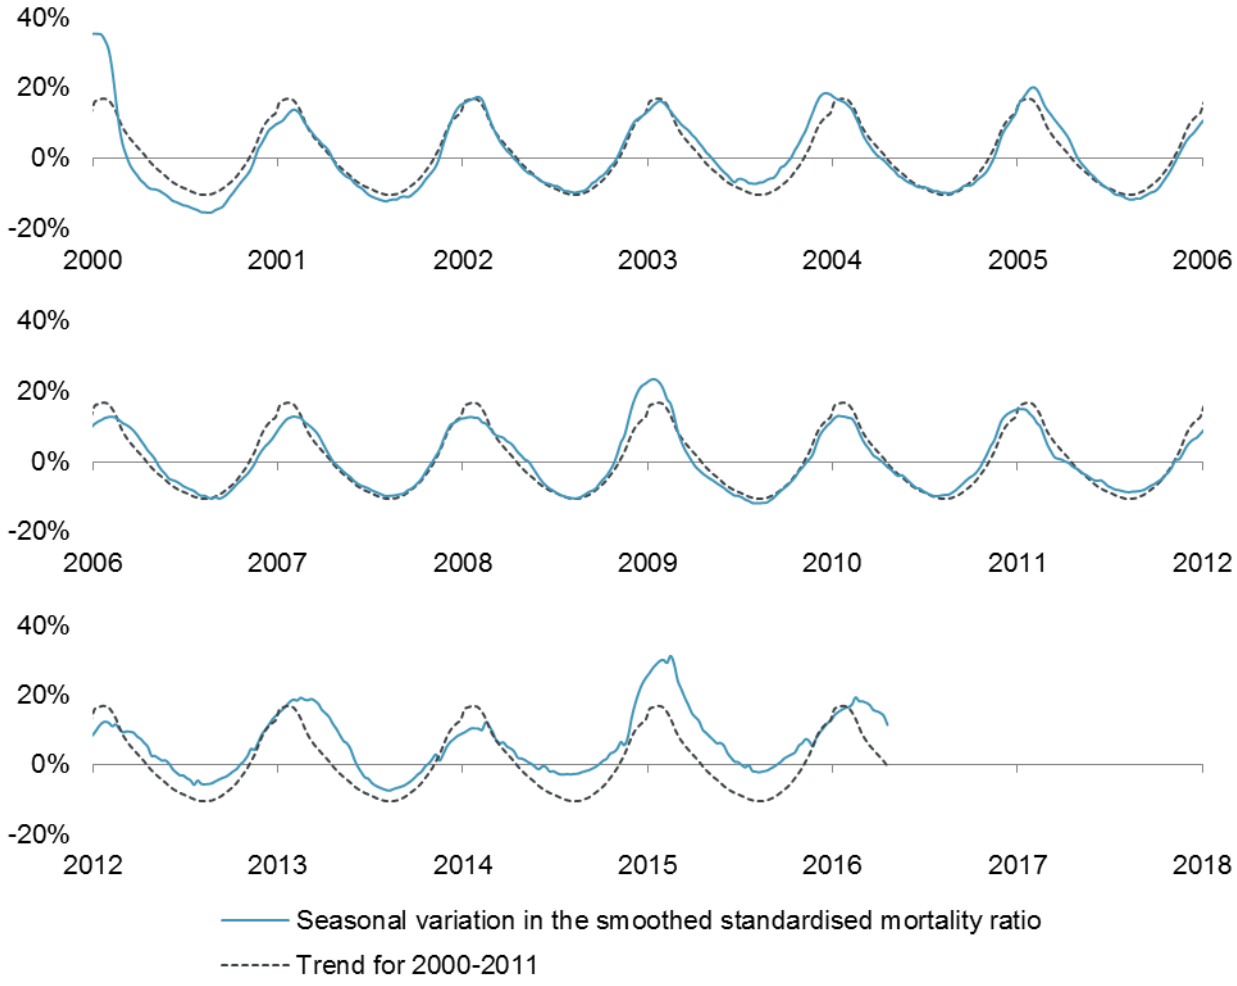

Chart 4C emphasises the messages of Charts 4A and 4B – that mortality has been heavier than expected since 2011. The corresponding analysis in Working Paper 83 used data to 31 July 2015. Mortality has remained above its long-term trend since then, except for a brief period at the turn of the year, when the weather was exceptionally mild. (Only the winter of 1868/69 is known to have been warmer than the winter of 2015/16, according to the Met Office Hadley Centre Central England Temperature Data[2]).

The seasonal analysis is helpful when considering possible reasons for heavy mortality. We tend to see greater uncertainty in mortality rates in the winter, due to temperature and influenza. But the fact that mortality has been above its long-term seasonal average in the last four summers suggests that the higher mortality is not purely due to short-term winter effects and could indicate the start of a medium- or long-term trend.

[2] http://www.metoffice.gov.uk/hadobs/hadcet/ssn_HadCET_mean_sort.txt

5. Calibration data

All the models we have considered need to be calibrated to historical data. We view the choices of model and calibration data to be separate and independent; i.e. it is unnecessary to make different decisions about the data depending on the chosen model.

For the current Model we take the following approach:

1. Calibrate to data for England & Wales.
2. Use ONS data for deaths and population estimates.
3. Calibrate to ages 18-102 and a 41-year period.
4. Make our own estimates of old-age exposures, to have these before the ONS figures are published.
5. Adjust the ONS exposures to address concerns about the data.
6. Estimate exposures by age for the final year of the calibration period.
7. Estimate deaths by age for the final year, based on weekly deaths for the year to date.

We consider each of these in turn, and indicate our proposals for future versions of the Model, assuming the continued use of data from the ONS or equivalent national bodies. We then discuss concerns with exposure data, that could lead to a change in data source at some future date, in Section 5.8 and describe our proposed new method for exposure adjustment in Section 5.9.

5.1 Countries

To date, the Model has been calibrated to data for England & Wales and not included data for Scotland or Northern Ireland. The original decision was primarily due to the timings of releases of data; the data for England & Wales was released first, and waiting for other countries would have delayed the Model beyond the intended September release date.

Following an earlier consultation, the outcome of which is described in Working Paper 80, we intend to release CMI_2016 in March 2017 and later versions in March of each subsequent year. Each Model version will use calibration data up to the end of the calendar year (e.g. up to 31 December 2016 for CMI_2016) based on provisional data. Because of this change we are consulting on whether users would like the calibration data to be extended to include other countries.

The Committee proposes that, if the timing and availability of data allows, the Model should be calibrated to the UK population (including England, Wales, Scotland and Northern Ireland). This would better reflect the populations that UK actuaries typically apply the Model to. There is value in having a single standard Model so we recommend that the CMI produces only a UK model, rather than a UK model and an England & Wales model.

However the Committee notes the following points:

- The results of the Model may not change greatly, given the relative size of the datasets. (According to the 2011 Census, the England & Wales population is 89% of the UK population.) Table 7 of Working Paper 74 showed life expectancies for a range of ages based on models calibrated to England & Wales and UK data; the differences ranged from −0.4% to +0.1%.
- Publication of the Model relies on the calibration data, and using data provided by three separate organisations rather than just one increases the risk of delay due to circumstances outside the Committee's control.
- It is unlikely that the population estimates by single year of age at ages 90 and above derived by aggregating the estimates for England & Wales, Scotland and Northern Ireland would be the same as those derived by applying the methodology used to UK data as a whole.

We intend to publish software that would enable users to calibrate the Model to a population of their choice.

Q 5.1. Would you prefer the Model to be calibrated to data for England & Wales (as is done currently) or the United Kingdom (England, Wales, Scotland and Northern Ireland)?

In the rest of this section we use the phrase "ONS data" for consistency with the current Model, but this should be interpreted to include data from equivalent bodies in Scotland and Northern Ireland, if we produce a model calibrated to the United Kingdom.

5.2 Use of ONS data

The Model is calibrated to national population data, but it is usually applied to specific groups, such as annuitants or pension scheme members, which differ from the general population. For example, they may have different socio-economic characteristics or be concentrated in a particular geographic area. This leads to basis risk.

In Section 5.8 we note concerns over the exposure data and describe ongoing work by the CMI and the ONS to consider and address these concerns. As we do not yet know what conclusion will be reached, we intend to continue to use national population data for the time being. We will continue to adjust the exposures, as described in Section 5.9.

5.3 Age and year range

Pending the outcome of further research described in Section 5.8, we intend to calibrate the Model to an age range of 20-100, and to extrapolate to older ages.

We note that a model that was fitted to older ages only (e.g. to ages 60-100 for a portfolio of annuitants or pensioners) could give different results. However, many users do need to produce results at younger ages, and we think there is considerable value in having a single standard model that covers a wider age range, rather than having multiple models. Also, a model calibrated only to older ages would not have data to allow such reliable estimation of cohort effects for the youngest of those ages.

We intend to continue our current practice of using 41 calendar years of data, which gives rise to 40 years of annual mortality improvements. The period used will move forward each year, so we propose to use data for 1976-2016 for CMI_2016, 1977-2017 for CMI_2017, etc.

5.4 Own estimates of old-age exposures

We have previously needed to make our own estimates of old-age exposures because the official ONS data was typically released after the target publication date of the Model. Our own estimates were intended to closely match the ONS method. As the release date of the Model has changed to March, the ONS data should be available in time, and we would no longer need to make our own estimates.

5.5 Exposure adjustments

Working Paper 74 noted concerns over the quality of the ONS exposure data, particularly for cohorts with unusual patterns of birth, and described adjustments that we made to it in CMI_2014. The same method was used in CMI_2015.

The concerns remain but we propose to take a simpler approach to making the adjustments. This is described in detail in Section 5.9.

5.6 Exposure estimate for the final year

When we publish CMI_2016 in March 2017, the ONS exposure data for calendar year 2016 will not be available. As a result we will need to estimate this. We propose to use the same method as was used for CMI_2014 and CMI_2015, described in Working Paper 74.

5.7 Deaths data for the final year

For CMI_2014 and CMI_2015 we have estimated the number of deaths for the final year of the Model (2014 and 2015 respectively) based on weekly deaths data for the first part of the year. The change of publication date of the Model means that it should no longer be necessary to do this, as we should have weekly deaths data for the full year.

However the weekly deaths data is provided by broad age group and, unless we can obtain more detailed data from the ONS, it will still be necessary for us to allocate this data to individual years of age for the final year of the calibration data. We propose to do this in the same way as for CMI_2014 and CMI_2015, by using a standard mortality table to allocate the deaths.

5.8 Concerns with ONS exposure data

Concerns with ONS exposure data have been reported by Cairns et al (2014) and by the CMI in Working Papers 74, 77 and 85. There are a number of related issues, many of which particularly affect older ages.

Unusual patterns 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. The 1919 and 1920 cohorts are particularly affected by the unusual pattern of births following the return of troops at the end of World War I.

Cairns et al noted that 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, as mentioned above, 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.

"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; until correction at the next census. This problem is particularly acute at older ages, where mortality rates are high. A similar problem of "negative phantoms" holds if a census underestimates the true population.

Exposure estimates at age 90 and above

The ONS publishes estimates for the England & Wales 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 Kannisto-Thatcher survival ratio methods, constrained to the official mid-year population estimate of the total population aged 90 and over. This approach can lead to a discontinuity in the estimates for ages 89 and 90.

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 and inter-censal periods.

Differences between ONS and HMD data

Analysis by the CMI High Age Mortality Working Party in Working Paper 85 indicates 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 Human Mortality Database ("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.

The Working Party is seeking to better understand the approaches followed for very high ages by ONS and HMD, as well as potential variants; and the ONS is also reviewing its own methods. At this stage, the Committee considers that there is no compelling reason to move away from the current ONS dataset (after making the exposure adjustments described below) until this work is complete. It also notes that HMD data for the United Kingdom is currently only available up to 2013 and so is not suitable for modelling very-recent mortality improvements; and that if the difference between ONS and HMD data is stable over time, this may not greatly affect mortality improvements.

5.9 Proposed new exposure adjustment method

As noted above, the Committee has concerns about the quality of the exposure data in the ONS dataset. As a result, we made adjustments to the exposure data used to calibrate CMI_2014 and CMI_2015 to reduce the impact of these concerns. That method is described in Working Paper 74. We propose to continue to adjust the exposure data, but to use a modified approach. The new method has been applied to the ONS dataset prior to the modelling in the rest of this paper.

The principle of our proposed method is that we expect mortality rates to vary smoothly with age. While there may be some inaccuracy in the deaths data, for example the age recorded at death, we expect deaths data to be much more reliable then exposure data; so any outliers from the assumption of smooth mortality rates suggest a problem with the exposure data.

We start with ONS data consisting of registered deaths $D_{x,t}$ and exposure estimates $E_{x,t}$ for a range of ages x and years t. For each specific combination (X,T) of age and year we want to decide whether to use the existing exposure $E_{X,T}$ or to adjust it.

We assume that the smoothed mortality rate $m_{x,T}$ in the age range $[X-n, X+n]$ in year T is exponential and so can be expressed as:

$$\log m_{x,T} = a_{X,T} + b_{X,T}x$$

for some parameters $a_{X,T}$ and $b_{X,T}$. We fit these parameters using least squares regression over that age range, to minimise the expression:

$$\sum_{x\in[X-n,X+n]}\left(a_{X,T} + b_{X,T}x - \log\left(\frac{D_{x,T}}{E_{x,T}}\right)\right)^2$$

The approach taken means that our estimate of the smoothed mortality rate $m_{X,T}$ for the specific point that we are considering is given by:

$$\log m_{X,T} = \frac{1}{2n+1}\sum_{x\in[X-n,X+n]}\log\left(\frac{D_{x,T}}{E_{x,T}}\right)$$

We then calculate the deviance residual $r_{X,T}$ as:

$$r_{X,T} = \text{sign}(D_{X,T} - E_{X,T}m_{X,T})\sqrt{2\left(D_{X,T}\log\left(\frac{D_{X,T}}{E_{X,T}m_{X,T}}\right) - \left(D_{X,T} - E_{X,T}m_{X,T}\right)\right)}$$

If our assumption that smoothed mortality is exponential in the age range $[X-n, X+n]$ in year T holds, then we would expect the deviance residual to be Normally-distributed with a mean of zero and variance of one. If this is not the case then this suggests a potential problem with exposure data.

We write Φ for the cumulative distribution function of the standard normal distribution. Then:

- if $|r_{X,T}| \leq \Phi^{-1}\left(1-\frac{p}{2}\right)$ we use the unadjusted exposure $E_{X,T}$.
- However if $|r_{X,T}| > \Phi^{-1}\left(1-\frac{p}{2}\right)$ we use the adjusted exposure $E'_{X,T} = \frac{D_{X,T}}{m_{X,T}}$.

In our work we have used parameters $n = 2$ and $p = 1\%$. For ages at and near the edges of the data we need to use a lower value of n, e.g. for ages 21 and 99 we use $n = 1$, and for ages 20 and 100 we make no adjustment.

Chart 5A shows crude mortality improvements derived from the raw ONS England and Wales data before and after adjusting exposures using our proposed methods.

Chart 5A: Crude mortality improvements, before and after exposure adjustment

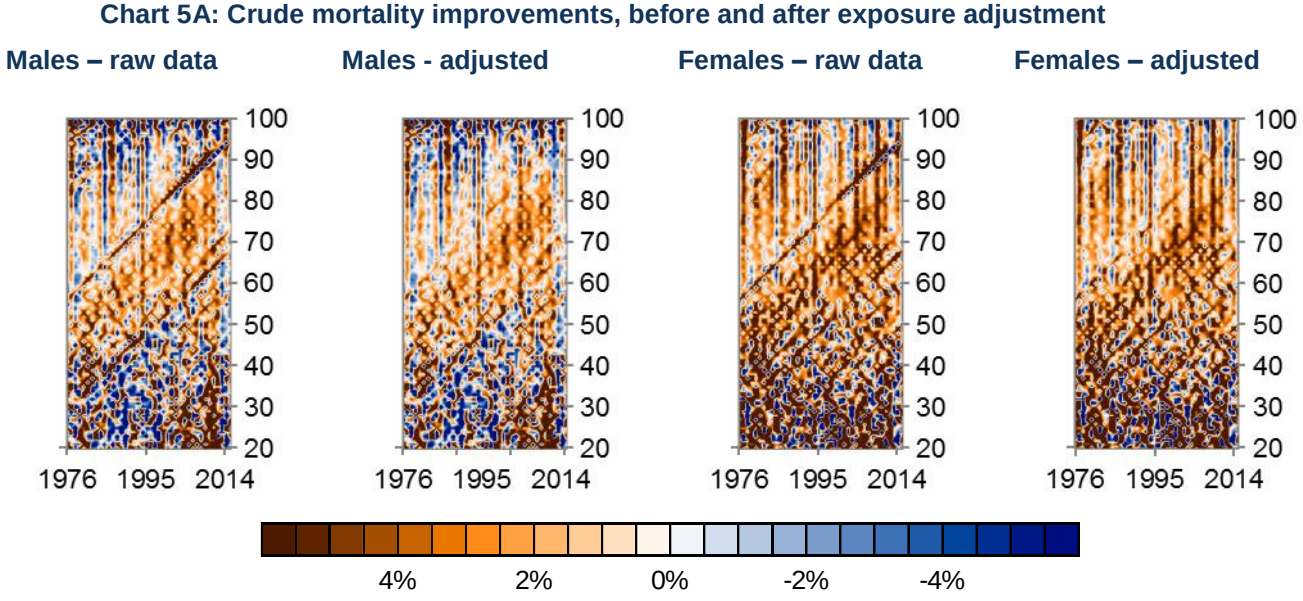

Comparing raw and adjusted charts for each gender, the main result of the approach has been to reduce the strong diagonal lines around the 1919 and 1920 cohorts, and to a lesser extent those around the 1947 cohort. These post-war cohorts were identified in earlier work as being of greatest concern.

While the new method is similar in principle to our previous approach it has a number of advantages:

- It is based solely on the assumption of a smooth progression of mortality rates by age. There is no assumption of smoothness over time, so the method is unaffected by annual noise.
- It is not dependent on the form of the model used subsequently. In particular, it is helpful to move away from the previous use of a p-spline model for exposure adjustment as we propose not to use a p-spline model for fitting mortality rates.
- It is not affected by the number of years included within the dataset, and only the highest and lowest n ages are affected by the age range used.
- It is quick and easy to apply, and can be replicated by users with spreadsheet formulae.

We will consider the method in more detail in the Technical Working Paper.

Q 5.2. Do you agree with the proposed new exposure adjustment method?

6. Definition of mortality improvement

The definition of mortality improvement in the current Model dates back at least as far as Working Paper 1 and is expressed in terms of mortality rates, $q_{x,t}$:

$$MI_{x,t} = 1 - \frac{q_{x,t}}{q_{x,t-1}}$$

In this section we describe an alternative definition in terms of central mortality rates, $m_{x,t}$. This has appealing properties when used within our proposed model. However, we recognise that many actuaries, and actuarial computer systems, are more familiar with mortality rates $q_{x,t}$ and the final output of the proposed model will use the current definition of mortality improvement.

6.1 Alternative definition of mortality improvement

When fitting a model to central exposures and using a Poisson model for deaths it is more natural to work with central mortality rates $m_{x,t}$ and many models, including the p-spline model that underlies the current Model, express mortality in terms of $\log m_{x,t}$ (where $\log$ denotes the natural logarithm).

This suggests an alternative definition of mortality improvement that we will use in this paper and propose to use in future versions of the Model. It is defined in terms of central mortality rates $m_{x,t}$:

$$MI^*_{x,t} = -\nabla_t \log m_{x,t} = \log m_{x,t-1} - \log m_{x,t}$$

where ∇_t is a backward difference operator, with $\nabla_t f(t) := f(t) - f(t-1)$

There are several similarities between the two definitions:

- The timing is the same: the improvement at time t is based on mortality rates at times $t-1$ and t.
- A positive mortality improvement means that mortality rates are falling.
- A constant rate of mortality improvements means an exponential fall in mortality rates (either m or q).
- Numerical values of $MI_{x,t}$ and $MI^*_{x,t}$ are similar but not identical, with larger differences at older ages.

To see the similarity, we can relate $MI_{x,t}$ and $MI^*_{x,t}$ by:

$$MI_{x,t} = 1 - \frac{q_{x,t}}{q_{x,t-1}} \approx 1 - \frac{m_{x,t}}{m_{x,t-1}} = 1 - \frac{\exp(\log m_{x,t})}{\exp(\log m_{x,t-1})} = 1 - \exp\left(\nabla_t \log m_{x,t}\right) \approx -\nabla_t \log m_{x,t} = MI^*_{x,t}$$

where the approximation $1 - \exp\left(\nabla_t \log m_{x,t}\right) \approx -\nabla_t \log m_{x,t}$ uses a Maclaurin expansion.

6.2 Numerical impact

Table 6.1 shows examples of the size of the difference between the two definitions of improvements $MI_{x,t}$ and $MI^*_{x,t}$. To calculate this, we:

(1) Specify a base table $q_{x,0}$.

(2) Convert from q to m using the approximation $m_{x,0} = -\log(1 - q_{x,0})$

(3) Specify an improvement $MI^*_{x,1}$ for each age using the alternative method.

(4) Project mortality rates using (2) and (3), so that $\log m_{x,1} = \log m_{x,0} - MI^*_{x,1}$.

(5) Convert from m to q using the approximation $q_{x,1} = 1 - \exp(-m_{x,1})$, consistent with (2).

(6) Calculate $MI_{x,1} = 1 - \frac{q_{x,1}}{q_{x,0}}$

We consider two base tables (S2PMA and S2PFA) and two options for $MI_{x,1}$ (either 1% or 2% up to age 85, tapering to nil at age 110, as proposed for the shape of the long-term rate (LTR) in the proposed model in

Section 8.5). Table 6.1 shows the difference between $MI^*_{x,1}$ and $MI_{x,1}$ in basis points for the four cases; e.g. for the first row of results, $MI^*_{90,1}$ is 0.800% and $MI_{90,1}$ is 0.729%; a difference of 0.071% or 7.1 basis points.

Table 6.1: The difference between $MI_{x,1}$ and $MI^*_{x,1}$ in basis points for different choices for $MI^*_{x,t}$

Age	20	30	40	50	60	70	80	90	100	110	120
S2PMA, 1% LTR	0.5	0.5	0.6	0.6	0.8	1.3	3.2	7.1	8.6	0.0	0.0
S2PFA, 1% LTR	0.5	0.5	0.5	0.6	0.7	1.1	2.5	5.7	7.4	0.0	0.0
S2PMA, 2% LTR	2.0	2.0	2.1	2.2	2.6	3.6	7.2	14.7	17.3	0.0	0.0
S2PFA, 2% LTR	2.0	2.0	2.1	2.2	2.5	3.1	5.9	12.0	15.0	0.0	0.0

Table 6.1 shows that the difference between $MI_{x,1}$ and $MI^*_{x,1}$ depends on the mortality rate (with a larger difference for higher mortality rates) and the size of $MI^*_{x,1}$ (with a larger difference for higher improvements). The proposed taper of the long-term rate means the difference peaks at ages in the late-90s and then steadily reduces to zero at age 110.

6.3 Conversion

The new definition of mortality improvement will prove to be particularly useful in conjunction with the Age-Period-Cohort Improvement (APCI) model, described in Section 7. The definition enables us to consistently fit in one step: mortality rates, aggregate mortality improvements, and mortality improvements split by age-period and cohort. That does not seem possible using the current definition.

Although we propose to define mortality improvements within the APCI model in terms of m rather than q, and the long-term rate will be defined in the same way, we will still output the final improvements in terms of q. These improvements will be calculated as described below:

(1) Fit the APCI model to the historical calibration data in terms of $\log m_{x,t}$.

(2) Project mortality improvements $MI^*_{x,t}$.

(3) Project future mortality rates $\log m_{x,t}$, using (1) and (2).

(4) Convert future mortality rates from $\log m_{x,t}$ to $q_{x,t}$ using the approximation: $q_{x,t} = 1 - \exp(-m_{x,t})$.

(5) Calculate mortality improvements $MI_{x,t}$ based on $q_{x,t}$ from (4).

Q 6.1. Do you agree with the use of mortality improvements in terms of $\log m_{x,t}$ within the APCI model, and in the definition of the long-term rate?

Q 6.2. Do you agree with the method, described in Section 6.3, of converting to mortality improvements in terms of $q_{x,t}$?

7. The Age-Period-Cohort Improvement (APCI) model

The Committee has considered a large number of possible models, both our own inventions and models published by others. We will consider these in detail in the Technical Working Paper but in this paper we focus on our preferred option, the Age-Period-Cohort Improvement (APCI) model.

The APCI model described below has the following advantages compared with the Calibration phase of the current Model:

- It fits mortality rates and improvements and determines age, period and cohort components in one consistent step; rather than the multi-step approach of the current Model.
- Smoothing is achieved by regularisation (described below) rather than splines; avoiding concerns about possible artefacts due to the spline basis.
- The choice of weights applied to the regularisation penalties allows users to control responsiveness.
- It is consistent with the approach of commonly-used stochastic models, and could itself be used as a stochastic model.
- The entire modelling process runs quickly and can be done in a single Excel workbook, rather than using the current combination of R and several Excel workbooks. This brings the advantage that it can be made available to users by the CMI, aiding transparency.

7.1 Notation

We have used the following convention, which is intended to be consistent with the existing academic literature, where possible.

Age, calendar year and birth year:

x	is age at last birthday
t	is time; i.e. calendar year
$\bar{t}$	is the mean of the years within the calendar year range that is used to fit the model; e.g. if we calibrate to years 1975 to 2015, then $\bar{t}$ is 1995
c	is cohort, with $c = t - x$. Note that this does not correspond exactly to birth year.

We generally prefer to write cohort explicitly as $t - x$, but we use c where this is clearer, e.g. as the index to a sum by cohort.

Mortality rates and improvements:

$m_{x,t}$	is the central mortality rate
$MI^*_{x,t}$	is the mortality improvement, defined as $MI^*_{x,t} = -\nabla_t \log m_{x,t} = \log m_{x,t-1} - \log m_{x,t}$

Parameters:

α_x	for terms by age relating to mortality rates
β_x	for fitted terms by age relating to mortality improvements
κ_t	for terms by period (i.e. calendar year)
γ_{t-x}	for terms by cohort (i.e. birth year)

Other:

λ	is used, with a subscript, for the weight placed on a smoothing penalty; e.g. λ_γ applies to γ_{t-x}
S	is used to express a smoothing weight λ in a more user-friendly form e.g. $S_\kappa := \log_{10} \lambda_\kappa$
θ_i	is used when describing identifiability transforms

7.2 Model structure

The current Model derives mortality improvements in terms of age, period and cohort components. The APCI model is designed to be broadly consistent with this – its mortality improvements are expressed as the sum of age, period and cohort terms – and this motivates its name.

The APCI model is defined by:

$$\log m_{x,t} = \alpha_x + \beta_x(t - \bar{t}) + \kappa_t \ + \gamma_{t-x}$$

This leads to improvements (under our definition that $MI^*_{x,t} = -\nabla_t \log m_{x,t}$) of:

$$MI^*_{x,t} = -\beta_x - \nabla_t \kappa_t \ - \nabla_t \gamma_{t-x}$$

and since:

$$\nabla_t \gamma_{t-x} = \gamma_{t-x} - \gamma_{t-1-x} = \gamma_c \ - \gamma_{c-1} = \nabla_c \gamma_c$$

we can write this equivalently as:

$$MI^*_{x,t} = -\beta_x - \nabla_t \kappa_t \ - \nabla_c \gamma_c$$

This justifies the motivation and description of the model as having age ($-\beta_x$), period ($-\nabla_t \kappa_t$) and cohort ($-\nabla_c \gamma_c$) components to its improvements.

7.3 Identifiability

In common with many mortality models, the parameters of the APCI model are not identifiable; i.e. there are multiple sets of parameters that could give exactly the same fit for $\log m_{x,t}$. Specifically, the following transformations leave the values of $\log m_{x,t}$ unchanged for any values of $\theta_1, \dots, \theta_5$:

$$\alpha_x \longmapsto \alpha_x + \theta_1 - \theta_2(x - \bar{x}) + \theta_3(x - \bar{x})^2 + \theta_4$$

$$\beta_x \longmapsto \beta_x - 2\theta_3(x - \bar{x}) + \theta_5$$

$$\kappa_t \longmapsto \kappa_t + \theta_2(t - \bar{t}) + \theta_3(t - \bar{t})^2 - \theta_4 - \theta_5(t - \bar{t})$$

$$\gamma_{t-x} \longmapsto \gamma_{t-x} - \theta_1 - \theta_2((t - \bar{t}) - (x - \bar{x})) - \theta_3((t - \bar{t}) - (x - \bar{x}))^2$$

So that the parameter values are uniquely determined, we use the following five identifiability constraints:

$\sum_t \kappa_t = \sum_t t\kappa_t = 0$ i.e. a linear fit to κ_t would be zero for all years t

$\sum_c \gamma_c = \sum_c c\gamma_c = \sum_c c^2\gamma_c = 0$ i.e. a quadratic fit to γ_c would be zero for all cohorts c.

There are various choices that we could have made here. Because the Model projects age-period and cohort improvements differently, the set of constraints on γ_c is more important as it affects the split between age-period and cohort components. We choose to remove as much of a systematic pattern from the cohort terms as is possible, so that improvements are expressed primarily through age and period terms. The transformations noted above allow us to require that a quadratic fit to γ_c would be zero. This is consistent with the approach of Cairns et al (2009) for the Cairns-Blake-Dowd M7 model.

The set of constraints affecting κ_t is less important (and somewhat cosmetic) as, although it affects the split between age and period improvements, these will be projected in the same way.

7.4 Regularisation

The APCI model defined so far is similar to a number of stochastic mortality models (e.g. those in Cairns et al (2009)) and it could be projected in a similar way, by fitting parameters to minimise the deviance of the model, and choosing suitable time series for κ_t and γ_{t-x}, in order to project them.

For our purpose we are focussed on a smooth central projection rather than a stochastic projection. To achieve this, we want each of the sets of fitted parameters α_x, β_x, κ_t and γ_{t-x} to be smooth.

Currie (2013) and Delwarde et al (2007) smooth the parameters of mortality models by using p-splines. We instead achieve smoothness by the use of regularisation penalties; i.e. the objective function to be optimised is the sum of the deviance and the penalties.

For example the proposed penalty for α_x is:

$$\text{Penalty}(\alpha_x) = \lambda_\alpha \sum_x (\nabla_x^3 \alpha_x)^2 = \lambda_\alpha \sum_x (\alpha_x - 3\alpha_{x-1} + 3\alpha_{x-2} - \alpha_{x-3})^2$$

where:

∇_t^3 is a third-order difference operator, with $\nabla_t^3 f(t) = \nabla_t \left(\nabla_t \left(\nabla_t f(t)\right)\right)$

λ_α is a hyperparameter (i.e. one that is specified rather than fitted) that controls the amount of smoothing. A higher value of λ_α will lead to a smoother series for α_x

the sum is over those ages for which the difference operator can be calculated; e.g. for an age range of 20-100 the sum is for ages 23-100.

The use of regularisation penalties dates back to Whittaker (1923). It is similar in principle to the use of penalties in p-spline models, but without the need to use a spline basis. Applying regularisation directly to the parameters rather than using splines means that there is no need to make a decision about knot-spacing and knot-placement and there is no risk of artefacts caused by these. It can also mean that identifiability criteria can be easier to apply.

We use a third-order penalty for age and cohort parameters and a second-order penalty for κ_t; i.e.:

$$\text{Penalty}(\alpha_x) = \lambda_\alpha \sum_x (\nabla_x^3 \alpha_x)^2 = \lambda_\alpha \sum_x (\alpha_x - 3\alpha_{x-1} + 3\alpha_{x-2} - \alpha_{x-3})^2$$

$$\text{Penalty}(\beta_x) = \lambda_\beta \sum_x (\nabla_x^3 \beta_x)^2 = \lambda_\beta \sum_x (\beta_x - 3\beta_{x-1} + 3\beta_{x-2} - \beta_{x-3})^2$$

$$\text{Penalty}(\kappa_t) = \lambda_\kappa \sum_t \left(\nabla_t^{(2)} \kappa_t\right)^2 = \lambda_\kappa \sum_t \left(\kappa_t - 2\kappa_{t-1} + \kappa_{t-2}\right)^2$$

$$\text{Penalty}(\gamma_c) = \lambda_\gamma \sum_c (\nabla_c^3 \gamma_c)^2 = \lambda_\gamma \sum_c (\gamma_c - 3\gamma_{c-1} + 3\gamma_{c-2} - \gamma_{c-3})^2$$

The use of a second-order penalty for κ_t means that increasing λ_κ makes the period mortality improvements flatter; if λ_κ was infinite then the period component of mortality improvements would be flat. This is because $\text{Penalty}(\kappa_t)$ will be zero if $\kappa_t - 2\kappa_{t-1} + \kappa_{t-2} = 0$ i.e. if $\nabla_t \kappa_t = \nabla_t \kappa_{t-1}$. This is a helpful property as it means that mortality improvements will be flatter at the edges of the data where we have less evidence on their direction of travel. If we instead used a third- or higher-order penalty for λ_κ then the direction of travel of fitted mortality improvements would persist even at the edges of the data, leading to an over-responsive model.

The choices of the hyperparameters, λ, will allow users of the proposed model to control the smoothness of different features of the model. Of these, λ_κ is of particular importance as this controls the smoothness by calendar year; i.e. the responsiveness of the model to new data.

The values of λ will typically be very large. It is convenient to refer to the base-10 logarithm of these instead, so we define $S_i = \log_{10} \lambda_i$; e.g. $S_\kappa = \log_{10} \lambda_\kappa$

Note that the amount of smoothing implied by particular choices of the values of λ depends on the numbers of ages and years in the calibration data. For example if the number of ages used for calibration is reduced then the number of terms in $\text{Penalty}(\kappa_t)$ will not change, but the number of terms in the other three penalties and in the deviance will. So reducing the number of ages would make $\text{Penalty}(\kappa_t)$ relatively more important if the values of λ were unchanged, and lead to greater smoothing by time.

7.5 Fitting the model

The model is fitted by minimising an objective function that is the sum of the deviance and the penalty functions; i.e. the objective function is:

$$2\sum_{x,t}\left(D_{x,t}\log\left(\frac{D_{x,t}}{E_{x,t}m_{x,t}}\right)-D_{x,t}+E_{x,t}m_{x,t}\right)+\lambda_\alpha\sum_x(\nabla_x^3\alpha_x)^2+\lambda_\beta\sum_x(\nabla_x^3\beta_x)^2+\lambda_\kappa\sum_t\left(\nabla_t^2\kappa_t\right)^2+\lambda_\gamma\sum_c(\nabla_c^3\gamma_c)^2$$

where the first term expresses the goodness of fit, and the remaining terms express the smoothness of sets of parameters.

Note that this is applied to the full rectangle of data, including cohorts which have few data points. We are aware that it is common practice when fitting certain stochastic longevity models to exclude cohorts with few data points, but we have found such a precaution to be unnecessary when we apply regularisation.

The objective function is minimised iteratively using Newton's method. The method is to repeatedly:

1. Improve α_x (holding other parameters constant)
2. Improve β_x (holding other parameters constant)
3. Improve κ_t (holding other parameters constant)
4. Improve γ_{t-x} (holding other parameters constant)
5. Apply identifiability constraints
6. Recalculate the objective function and exit the loop when it has changed by less than 0.00001.

While the identifiability constraints could be added into the objective function using Lagrange multipliers, we have found that applying them explicitly instead improves the speed of the fitting process.

When applying the constraints we:

(1) Find the values of $\theta_1, \theta_2, \theta_3$ (in the notation of Section 7.3) so that the constraints applying to the cohort parameters are satisfied i.e. $\sum_c \gamma_c = \sum_c c\gamma_c = \sum_c c^2\gamma_c = 0$.

(2) Apply the adjustments related to $\theta_1, \theta_2, \theta_3$ to all of the parameters.

(3) Find the values of θ_4, θ_5 so that the constraints applying to the period parameters (after the adjustment in (2)) are satisfied i.e. $\sum_t \kappa_t = \sum_t t\kappa_t = 0$.

(4) Apply the adjustments related to θ_4, θ_5 to $\alpha_x, \beta_x, \kappa_t$.

The identifiability constraints have no impact on the deviance or on the penalties that use third-order differences. However the identifiability constraints do affect the penalty for κ_t , so the objective function can increase after applying the adjustment related to θ_3. As the impact appears to be minor (less than one part in 10,000 for the datasets we have considered to date) we take no specific action to address this.

Q 7.1. Do you agree with the use of the APCI model to determine components of historical mortality improvements?

Q 7.2. Do you agree with the identifiability constraints used in the APCI model?

Q 7.3. Do you agree with the approach taken in smoothing the APCI model; i.e. the use of regularisation penalties with a different order used for the period terms?

8. Projection and convergence

As noted earlier, projected mortality improvements in the current Model are the sum of projected age-period and cohort components. Each of these components is projected by non-linear interpolation between initial and long-term improvements. To do this, it is necessary to specify the following, for each age for age-period components, and for each birth year for cohort components:

1. The initial rate of improvements (based on historical data for the Core assumption)
2. The long-term rate (nil for the cohort component for the Core assumption)
3. The convergence period (i.e. the time at and after which the long-term rate applies)
4. The "proportion remaining at mid-point" (to control the pace of convergence)

The key concerns with this method are that:

- It does not easily allow for "direction of travel"; i.e. the initial slope of mortality improvements.
- It is inconsistent; i.e. if experience is as expected, the results of next year's Model will not agree with this year's.

8.1 Allowing for direction of travel

We can slightly amend the current approach to convergence to allow for direction of travel.

For a particular age or cohort, we write $f(t)$ for the mortality improvement at time t, where $f(t)$ is defined by:

$$f(t) = L + (I - L)\left(1 - 3\left(\frac{t}{T}\right)^2 + 2\left(\frac{t}{T}\right)^3\right) + Dt\left(1 - \frac{t}{T}\right)^2 \qquad \text{for } 0 \le t \le T$$

$$f(t) = L \qquad \text{for } t > T$$

with parameters:

L	the long-term rate of mortality improvements
I	the initial rate of mortality improvements
D	the "direction of travel"; i.e. the initial slope of mortality improvements
T	the convergence period

The function $f(t)$ has been chosen so that:

$f(0) = I$	matching the specified initial rate of improvements
$f(t) = L$ for $t \ge T$	matching the specified long-term rate of improvements
$f'(0) = D$	matching the specified direction of travel
$f'(t) = 0$ for $t \ge T$	improvements are flat beyond time T

Under the current approach, users of the model specify the proportion remaining at mid-point, p. Under the proposed approach users could choose to specify D directly to specify the initial slope of mortality improvements. This would allow for "direction of travel".

If we set $D = \frac{1}{T}(8p - 4)(I - L)$ we obtain:

$$f(t) = L + (I - L)\left(1 + (8p - 4)\left(\frac{t}{T}\right) + (5 - 16p)\left(\frac{t}{T}\right)^2 + (8p - 2)\left(\frac{t}{T}\right)^3\right) \qquad \text{for } 0 \le t \le T$$

and so:

$$f\left(\frac{T}{2}\right) = L + p(I - L).$$

This shows that we can retain the current approach to convergence of specifying the proportion remaining at mid-point, p. Setting $p = \frac{1}{2}$ is equivalent to setting $D = 0$ and gives the Core parameterisation.

Chart 8A shows examples of convergence to a long-term rate of 1.5% p.a. with and without allowance for direction of travel.

Chart 8A: Illustrations of convergence with and without direction of travel

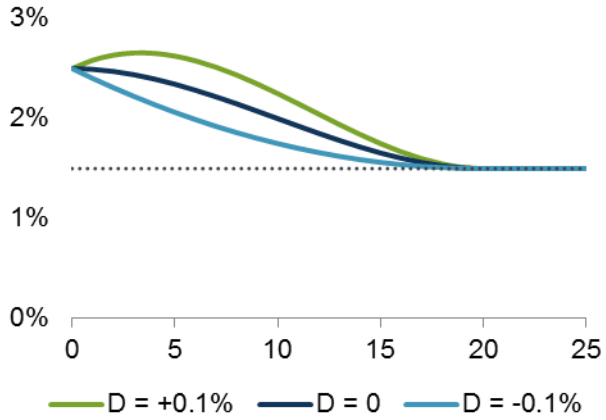

The function $f(t)$ is exactly equivalent to the specification of convergence in the current Model, introduced in Section 5.3.2 of Working Paper 39. The only difference is that here $f(t)$ represents mortality improvements directly, rather than as a weight to be placed on the initial rate, as in Working Paper 39 and the Model.

The chosen representation of $f(t)$ enables us to see the relationship between specifying a proportion remaining at mid-point, as in the current Model, or an initial slope to mortality improvements to preserve "direction of travel".

8.2 Difficulty of determining direction of travel

We have seen that re-expressing the convergence function would allow us to reflect direction of travel. But it is difficult to reliably estimate the direction of travel. Chart 8B shows the period component of mortality improvements, using the current method, for calibration periods of 1965-2005, 1966-2006, …, 1975-2015. The direction of travel, i.e. the slope of each line at its right-hand end, varies greatly. The direction of travel, allowing for the two-year step-back and measured in basis points per year, was +7 in 2005, fell to +3 in 2008 and 2009, rose to +9 in 2011 and is −15 in 2015. Given this uncertainty, the current Model's approach of effectively assuming nil direction of travel does not seem unreasonable.

Chart 8B: Period component of improvements for different calibration periods (current method)

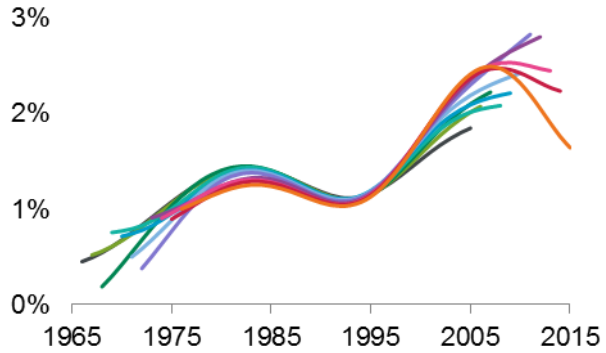

Note that this chart has no legend as it is only intended to illustrate the range of outcomes.

However using second-order penalties for κ_t in the APCI model (as described in Section 7.4) acts to dampen the direction of travel compared to third-order penalties or the current Model. As a result, we propose to amend the Model so that users can specify the shape of the convergence function either by using the proportion remaining at mid-point (as now) or by using the direction of travel under the approach outlined above. Also, we will publish the value of the direction of travel in the APCI model, a single value for all ages. However the Core parameterisation will not allow for direction of travel; i.e. it will have $D = 0$ or, equivalently, $p = 50\%$.

Q 8.1. Do you agree that we should give users the option of specifying convergence as they choose: either in terms of direction of travel or proportion remaining at mid-point?

Q 8.2. If you agree with allowing users to specify convergence in terms of direction of travel, should the Core model allow for this or not?

8.3 Inconsistency

We noted earlier that the projection method is inconsistent in the sense that even if mortality experience from year to year were as expected, the Model's projections would change. We have considered a number of ways to address this concern; these are discussed briefly in Section 13 and will be considered in more detail in the Technical Working Paper. Making the model consistent in this sense would seem to require us to make two amendments to the current Model:

1. Change from a finite convergence period (i.e. mortality improvements reach the long-term rates after a finite number of years) to asymptotic convergence (i.e. mortality improvements get closer to the long-term rates but never reach them); and
2. Allow for direction of travel to be potentially different for each age and cohort.

We concluded that while consistency would be theoretically desirable, and may make results of the Model easier to present, it would introduce unnecessary complexity. We therefore propose to maintain the same projection method, albeit with the ability to express convergence using direction of travel.

If experience is largely as expected, and consistency is considered to be important, then a user could choose to continue to use the previous version of the Model.

8.4 Initial rates of mortality improvement

In the current Model, initial rates of mortality improvement are those that apply two years before the end of the calibration data. If this "step-back" was not applied then the volatility of the Model from year to year would be excessive, due to the sensitivity of the p-spline model to an additional year of data.

In the current Model, the step-back effectively provides additional smoothing of mortality rates. We find this to be unnecessary in the proposed model due to the use of a second-order penalty for period terms (to flatten improvements in the most recent years of the calibration data) and the ability to control the degree of smoothing through S_κ.

When using the APCI model we propose that the initial rates will be determined from the last year of the calibration without any step-back:

Age-period component	$-\beta_x - \nabla_t \kappa_T$	where T is the last year of the data
Cohort component	$-\nabla_c \gamma_c$.	

8.5 Long-term rate of mortality improvement

Under the Core version of the current Model, users are required to specify a long-term rate of age-period mortality improvements that applies up to age 90, and this reduces linearly to a value of zero at ages 120 and above.

Based on analysis of historical mortality improvements, which we intend to present in the Technical Working Paper, we propose that the specified long-term rate should be used up to age 85, and then be reduced linearly to zero at ages 110 and above.

Q 8.3. Do you agree with the proposed Core assumption, to taper the long-term rate to zero between ages 85 and 110?

8.6 Older ages

The APCI model is fitted to ages 20-100. We also need to determine initial rates of mortality improvement for ages above 100. For consistency with the proposal for the long-term rate above, we propose that both age-period and cohort components of initial mortality improvements are interpolated linearly between their calibrated values at age 100 and nil at ages 110 and above. This is consistent with the current Model.

8.7 Cohort convergence periods

In the current Model the cohort component of mortality improvements is constrained to be nil for the youngest cohorts, so the cohort convergence period is irrelevant. However the APCI model assigns a large part of mortality improvements at young ages in recent years to cohort effects. Consequently the length of the cohort convergence period for young cohorts becomes a material choice.

This has led us to reconsider that assumption. We think that the current assumption, a 40-year convergence period for the youngest cohorts, may not be appropriate:

(i) The main causes of death at young ages (accidents and suicide) are different to the main causes at older ages (circulatory diseases, cancer etc), so the cohort pattern of improvements may be expected not to persist as members of the cohort age.

(ii) It is not clear how reliable the improvements for the youngest cohorts are as the number of deaths is relatively small, and high levels of net migration mean that the population estimates are uncertain.

We propose the Core assumptions in Table 8.1 for the cohort convergence period, where the ages refer to ages in the final year of the calibration data.

Table 8.1: Current and proposed cohort convergence periods (years)

Age (x)	Current	Proposed
20 to 49	40	$x - 10$
50 to 60	40	40
61 to 94	$100 - x$	$100 - x$
95 to 105	5	5
106 to 109	5	$110 - x$
110 and above	5	0

As well as a reduction for cohorts ages 20-49, for the reasons discussed above, we have also reduced the cohort convergence period for cohorts aged 106 and above in the final year of the calibration data. This is a minor “tidying-up” change, motivated by a desire to make mortality improvements nil at age 110 and above. For example, under the previous assumption, a cohort aged 108 at the start of the projection would have a five-year convergence period and so mortality improvements for that cohort would only reach nil at age 113.

Chart 8C shows the current and proposed convergence periods for cohort improvements, and the (unchanged) assumption for age-period improvements. Chart 8D is similar, but shows the age at which convergence is achieved.

Chart 8C: Convergence periods

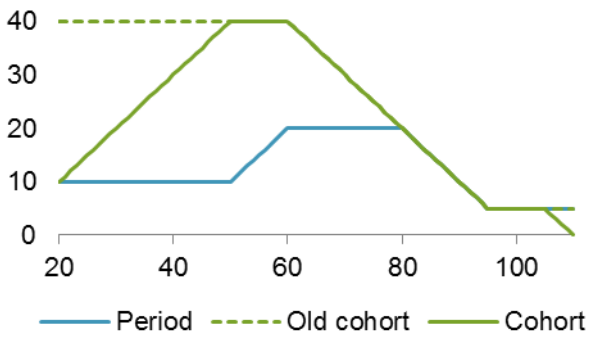

Chart 8D: Age at which convergence is reached

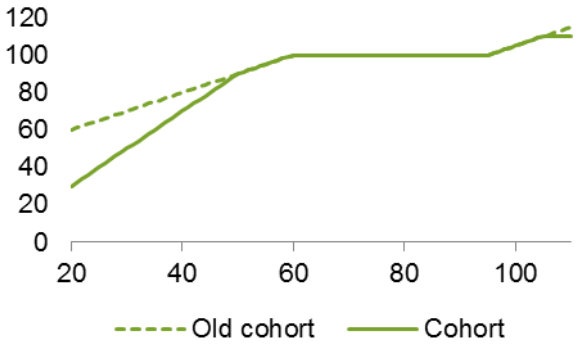

Q 8.4. Do you agree with the proposal to amend the cohort convergence periods in the way described in Section 8.7?

8.8 Future cohorts

The current Model produces mortality projections for "future cohorts" not contained in the calibration data – i.e. those aged below 20 in the last year of the calibration data. This is straightforward for the current Model: it constrains the cohort component of improvements to be nil for everyone below age 30, including future cohorts.

However the proposed model would have cohort improvements for all calibrated ages, so if we wanted to project improvements for future cohorts it would be necessary to make a cohort assumption for them. We have not proposed a method at this stage as we are unsure if users would find projections for future cohorts useful. If not, then we will not project them.

Q 8.5. Would you be happy for the Model to only make projections for cohorts contained within the calibration data; i.e. those aged 20 and above at the end of the calibration period?

9. Results

The results in this section are all based on ONS data for England & Wales. Data for 2015 has been estimated as described in Sections 5.6 and 5.7 and we have applied exposure adjustments as described in Section 5.9.

We show results for the APCI model and also for the current CMI method using the same data. This is intended to provide a comparison to the APCI model, using a consistent recent dataset. However the results shown for the CMI method are not an official release of the Model.

For the APCI model we use smoothing hyperparameters of $S_\alpha = 7$, $S_\beta = 9$, and $S_\gamma = 7$ and consider the effect of different values of S_κ.

The results in Sections 9.1 to 9.5 use the current assumptions for the shape of the long-term rate by age and cohort convergence periods, as the focus is on the impact of different options for calibration and either projection assumptions are irrelevant or it is helpful to compare calibration options using the same projection method.

The results in Sections 9.6 and 9.7 use the proposed projection assumptions for the proposed model, and the current assumptions for the current Model, in order to illustrate the difference between them.

9.1 Responsiveness

In Section 4 we used standardised mortality ratios (SMRs) to look at crude mortality rates and improvements over time across all ages. Here we have calculated SMRs based on mortality rates for fitted models to make comparisons between them and to compare with the crude data.

Charts 9A and 9B show mortality improvements based on those SMRs for models fitted to data for 1975-2015. They show results for the current CMI method and for the APCI model with four illustrative values of S_κ. We write "APCI(S_κ)" for the APCI model with a particular value of S_κ.

Chart 9A: Mortality improvements for 1975-2015 data – males

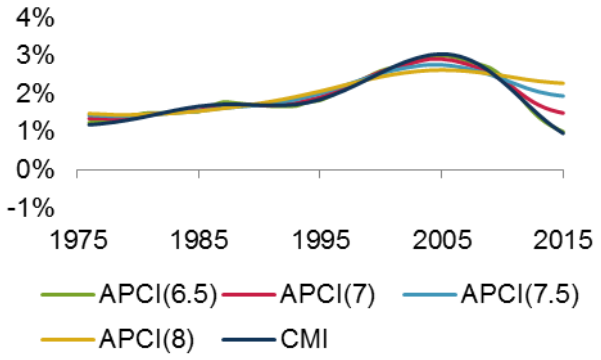

Chart 9B: Mortality improvements for 1975-2015 data – females

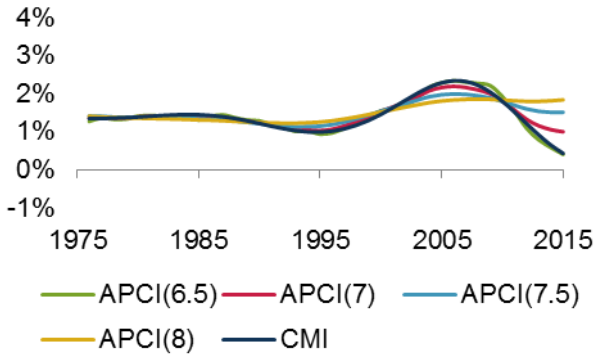

Charts 9A and 9B demonstrate that a higher value for S_κ leads to flatter mortality improvements. The difference is particularly noticeable in recent years, where a higher value for S_κ gives a flatter peak around 2005 and a lesser fall in improvements in recent years. The improvements for the current CMI method give results that are similar to those for APCI(6.5) for both males and females.

For the current CMI method we step back two years, so that the initial rates would be based on 2013. For the APCI model we do not propose to step back, so the initial rates would be based on 2015. The SMR-based improvement in 2013 for the CMI method is most similar to the improvement in 2015 for APCI(7). So although APCI(6.5) gives the closest improvements to the Model for the 1975-2015 dataset, APCI(7) would give the closest initial rates, and hence projection, because of the treatment of the step-back.

It is also instructive to compare the SMR-based mortality improvements from the fitted models with the corresponding crude mortality improvements. Charts 9C and 9D show this for four-year-average mortality improvements, as the volatility of annual improvements makes the picture unclear. The crude improvements are consistent with those shown in Chart 4B for males and females combined.

Chart 9C: Four-year average improvements for 1975-2015 data – males

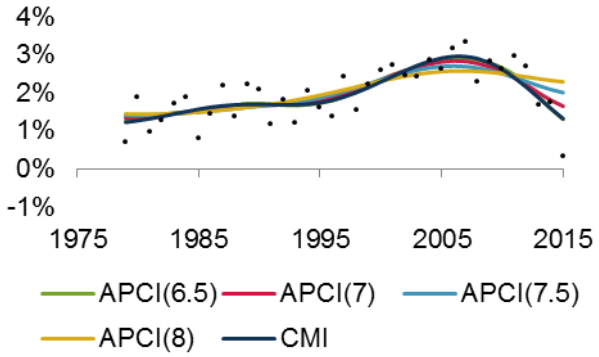

Chart 9D: Four-year average improvements for 1975-2015 data – females

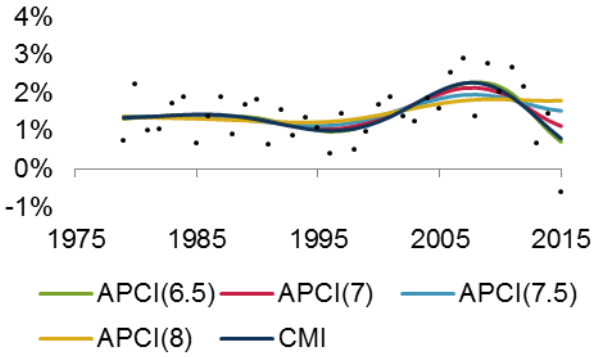

9.2 Progression of life expectancy

In Section 9.1 we showed how fitted mortality improvements vary according to the choice of S_κ for a single period of calibration date (1975-2015). We now consider the stability of results from successive versions of the APCI model, and assess this using expectations of life.

For each sex and choice of S_κ we have fitted eleven models to different datasets: 1965-2005, 1966-2006, …, 1975-2015. We plot cohort expectation of life for each dataset at age 65 at 31 December 2015, using S2PMA and S2PFA mortality tables projected using an illustrative long-term rate of 1.5% p.a.. The results shown for "CMI" are calculated by applying the CMI method to the same datasets, and are not equal to historical versions of the Model. As the focus of this section is on the impact of S_κ we have consistently used the existing projection method for both models (i.e. the current convergence periods and long-term rate shape, rather than our proposal for CMI_2016).

Chart 9E: Progression of cohort life expectancy at age 65 for calibration data ending in various years – males

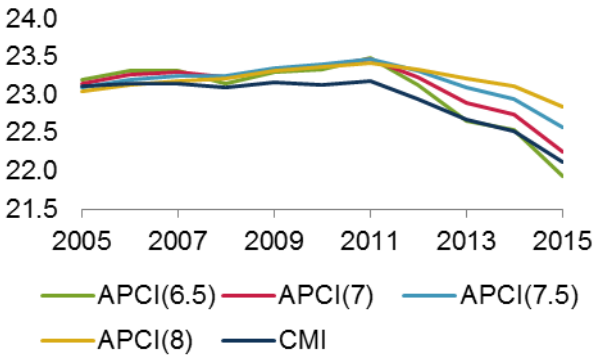

Chart 9F: Progression of cohort life expectancy at age 65 for calibration data ending in various years – females

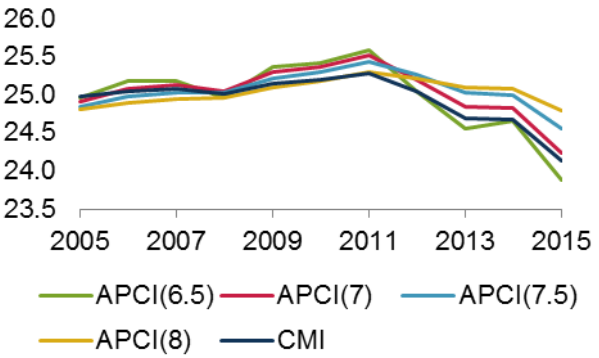

In each case, as expected, a higher value of S_κ gives a smoother progression of life expectancy. For males, the four parameter values for the APCI model produce broadly similar results up to 2011, but markedly different results afterwards. For females there is also a material difference in results before 2011.

Table 9.1 shows the fall in life expectancy at age 65 between 2011 and 2015 for the four parameter values for the APCI model, and for the CMI method.

Table 9.1: Fall in life expectancy (in years) at age 65 between 31 December 2011 and 31 December 2015

Model	Males	Females
APCI(6.5)	1.55	1.70
APCI(7)	1.22	1.29
APCI(7.5)	0.89	0.87
APCI(8)	0.58	0.51
CMI	1.06	1.15

9.3 Choice of S_κ

We now consider what would be an appropriate default value for S_κ, based on the analysis above. A value of 7 or less would give rise to a fall in life expectancy that is greater than the current CMI method. Our perception is that the majority of users think that CMI_2015 may have responded too strongly to recent data, so a value of 7 or less would seem to be unpopular as a default value. Conversely, Chart 9B shows that a value of 8 for S_κ would produce improvements for females that are marginally higher in 2015 than in 2011, despite the unprecedented low improvements of 2011-2015. That value seems to over-smooth the data. Consequently we propose a value of 7.5 for S_κ. In the absence of any changes to the projection method, that value would lead to slightly higher liabilities than CMI_2015 at older ages.

While 7.5 is intended to be a broadly reasonable figure for S_κ, other values are plausible and we encourage users to form their own views and adjust the model if they wish.

As noted in Section 7.4, the impact of S_κ depends on the number of ages and years in the calibration data. A particular value of S_κ that is deemed appropriate for one dataset may not remain appropriate if used for another that covers a different number of ages or years.

Q 9.1. Do you agree with the proposed Core assumption for S_κ of 7.5?

9.4 Goodness-of-fit

Cairns et al (2009) assess the goodness-of-fit of various models by plotting the patterns of positive and negative residuals. We do the same, in Charts 9G and 9H, for the APCI model with $S_\kappa = 7.5$. Each combination of age and calendar year is shown as light-grey if the actual number of deaths is higher than indicated by the fitted model, and dark-grey otherwise.

Chart 9G: Pattern of residuals – males

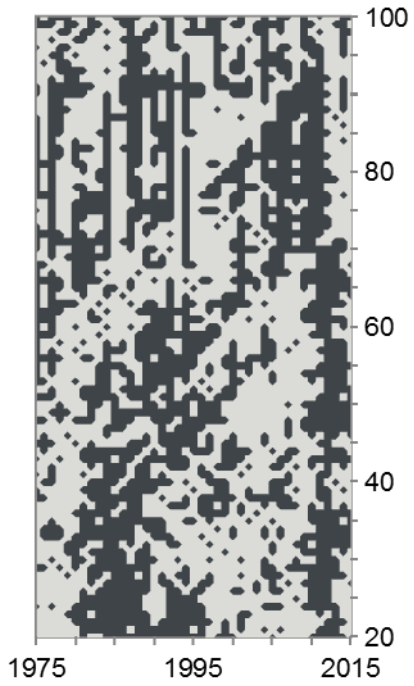

Chart 9H: Pattern of residuals – females

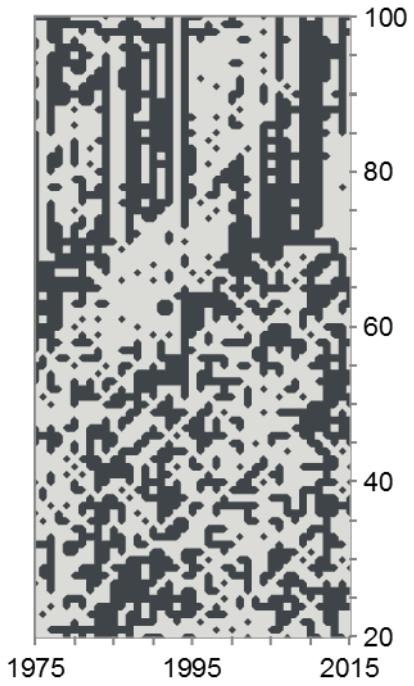

Cairns et al note that the pattern should be random if the residuals are independent and identically-distributed. In our case some patterns are apparent, particularly some vertical lines which are most prominent at older ages. We could improve the patterns of residuals by reducing the amount of smoothing in our proposed model, or including an explicit term to address annual volatility. However while these actions would make the model fit the data better, we do not think this would produce a better estimate of initial mortality improvements.

9.5 Comparison of components of improvements

In this section we compare results from the APCI model and the current CMI method. Both models are calibrated to the same data for 1975-2015, and we choose S_κ=7.5 for the APCI model, as noted in Section 9.4.

Charts 9I and 9J compare components of mortality improvements. Note that the improvements shown for the APCI model are in 2013, to be consistent with the two-year step-back of the current CMI method. As such the values shown for the APCI model are not those that we propose would be used as the initial rates in its projection.

Chart 9I: Components of mortality improvements in 2013 – males

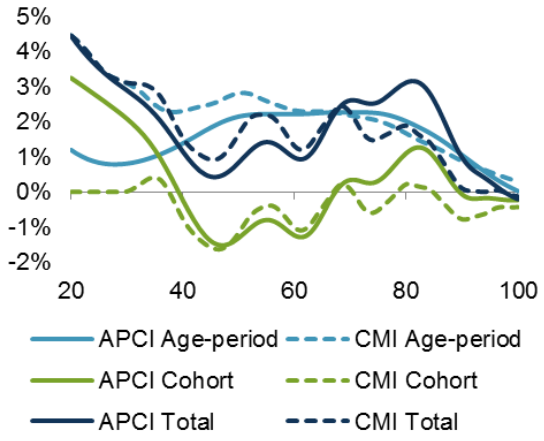

Chart 9J: Components of mortality improvements in 2013 – females

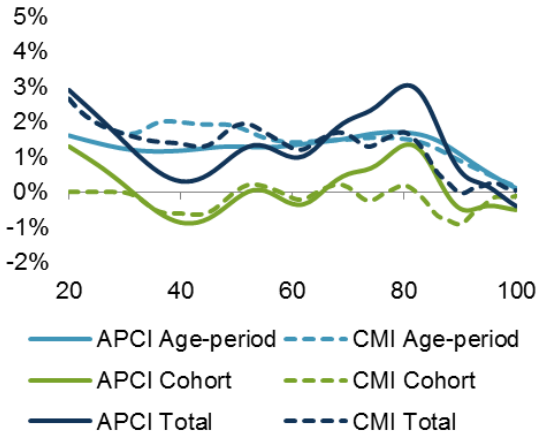

For both males and females, the APCI model gives higher mortality improvements than the current CMI method at older ages, with a stronger cohort effect around age 80, and lower improvements for broadly ages 30 to 60. The cohort improvements at young ages are prominent for the APCI model, while constrained to be nil for the CMI method (as discussed in Section 8.7).

It is noticeable that the ages at which peaks and troughs of cohort improvements are seen are broadly consistent between the two methods despite their different approaches to smoothing.

9.6 Comparison of life expectancies

Tables 9.2A and 9.2B show life expectancies at selected ages for different models; using data for 1975-2015. As in Section 9.2, these are as at 31 December 2015, using S2PMA and S2PFA mortality tables projected using an illustrative long-term rate of 1.5% p.a.

We consider:

a. The current CMI method. (Note that as we use data for the whole of 2015, these values do not correspond with values from CMI_2015.)

b. An intermediate stage using the proposed method for convergence, but the current projection assumptions. This is provided only to enable understanding of the relative importance of our proposed changes to calibration and projection.

c. The proposed method for both convergence and projection.

Table 9.2A: Life expectancies for the APCI and CMI Models – males

Age	25	45	65	75	85
a. Current.	64.51	42.30	22.12	13.30	6.60
b. Intermediate	64.66	42.35	22.58	13.84	6.82
c. Proposed	63.83	41.99	22.40	13.74	6.78
Change (current to proposed,%)	−1.1%	−0.7%	1.3%	3.3%	2.6%
Impact of calibration changes	0.2%	0.1%	2.1%	4.1%	3.3%
Impact of projection changes	−1.3%	−0.9%	−0.8%	−0.8%	−0.7%

Table 9.2B: Life expectancies for the APCI and CMI Models – females

Age	25	45	65	75	85
a. Current.	66.65	44.58	24.13	14.90	7.58
b. Intermediate	66.49	44.61	24.56	15.38	7.70
c. Proposed	65.68	44.14	24.33	15.24	7.64
Change (current to proposed,%)	−1.4%	−1.0%	0.8%	2.3%	0.7%
Impact of calibration changes	−0.2%	0.1%	1.8%	3.3%	1.5%
Impact of projection changes	−1.2%	−1.1%	−1.0%	-0.9%	−0.8%

At younger ages, the proposed life expectancies are lower than for the current method for both males and females, primarily because of changes to the projection method rather than the calibration method. For older ages, the proposed life expectancies are higher. This is primarily because of changes to the calibration method, consistent with the choice of value for S_κ (intended to give a smaller fall in life expectancy in recent years than the CMI method) but partly mitigated by the changes to the projection method. The impact at age 75 is notably higher than at ages 65 and 85.

9.7 Comparison of mortality improvements

Charts 9K to 9N show historical and projected mortality improvements for the current and proposed methods (i.e. options (a) and (c) from the list above). Charts 9O and 9P plot the difference in mortality improvements between the two approaches.

Chart 9K: Mortality improvements for the proposed model – males

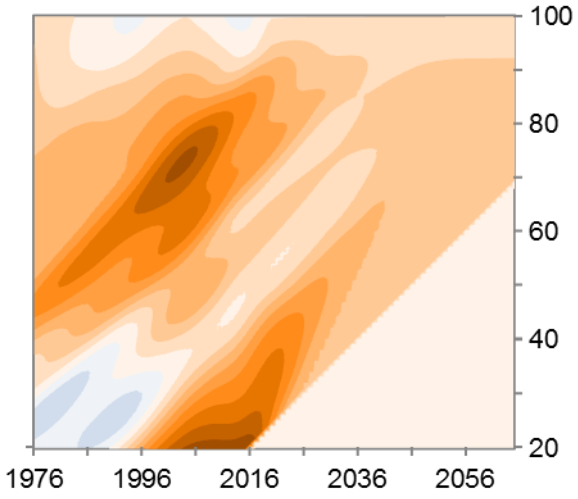

Chart 9L: Mortality improvements for the proposed model – females

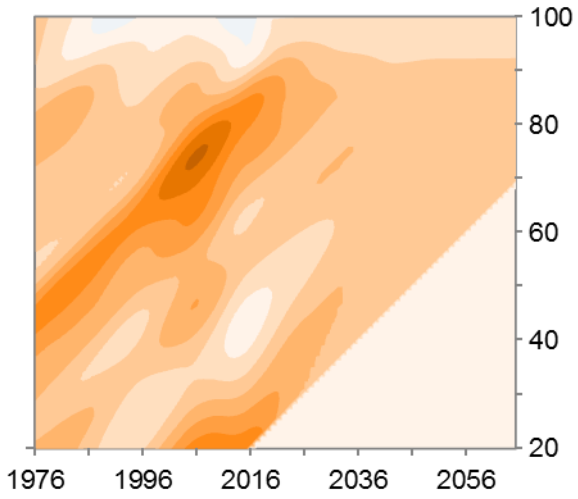

Chart 9M: Mortality improvements for the CMI method – males

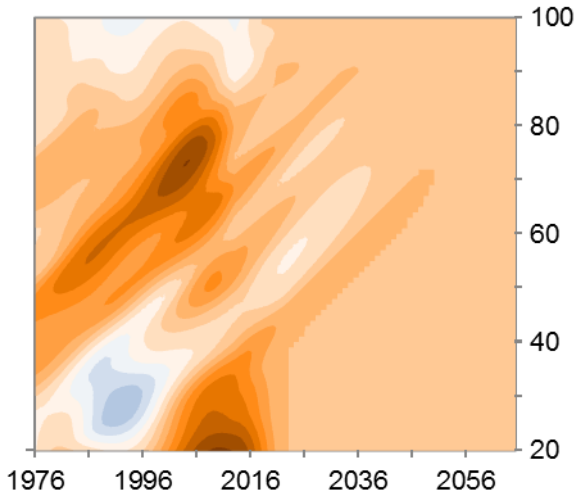

Chart 9N: Mortality improvements for the CMI method – females

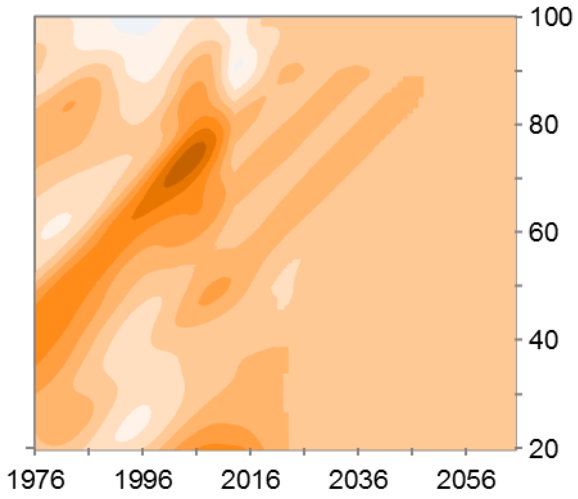

Chart 9O: Difference in mortality improvements proposed *minus* current – males

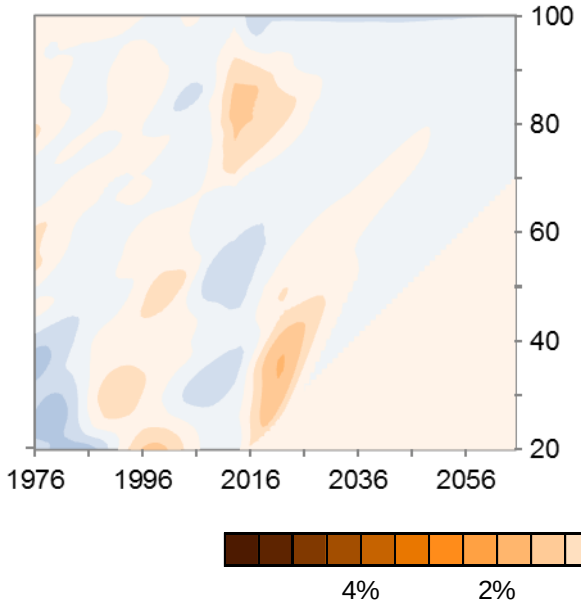

Chart 9P: Difference in mortality improvements proposed *minus* current – females

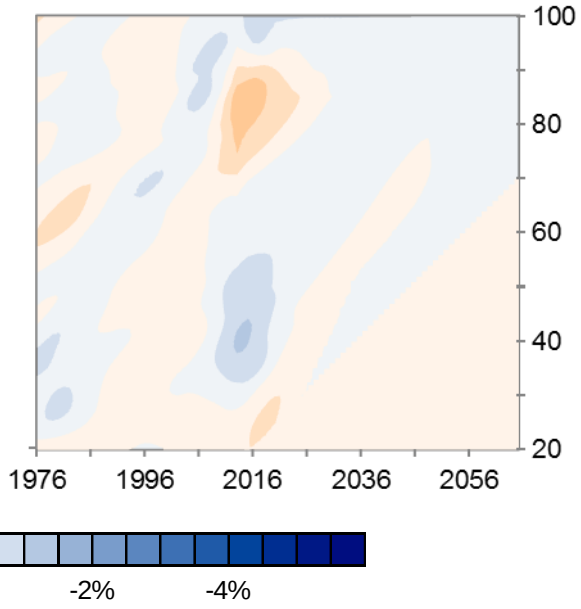

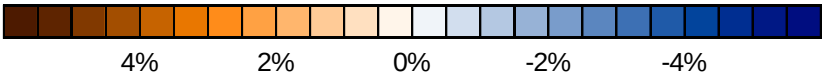

For the historical period the broad pattern of improvements is broadly similar for the two approaches. The most notable difference is the higher mortality improvements for the proposed approach at around age 80 in recent years.

For the projection, there are differences at younger cohorts, as the proposed approach assigns more of the improvements to cohort terms, and at older ages, as we taper the long-term rate between ages 85 and 110 rather than ages 90 and 120.

10. Model parameters

The current Model distinguishes between "Core" and "Advanced" parameters. We propose to remove one of the Core parameters, add new Advanced parameters, and introduce an additional layer of "Intermediate" parameters to make it easier to set and communicate the Advanced parameters.

10.1 Constant Additional Rate of Mortality Improvement

The "Constant Additional Rate of Mortality Improvement" is the second parameter in the Core version of the Model; intended to allow an easy means of adding a margin to the Model's unadjusted outcome, but with a default value of zero. We think this is little-used and propose removing it altogether.

Q 10.1. Do you agree with the removal of the Constant Additional Rate of Mortality Improvement parameter?

Q 10.2. If you disagree with the removal of the Constant Additional Rate of Mortality Improvement parameter, do you agree with making it an Advanced rather than a Core parameter?

10.2 Smoothing parameters

We have noted that the responsiveness of the APCI Model to recent data can be controlled by users being able to change the value of the period smoothing parameter. We propose that this should be a new Advanced parameter.

There are a number of choices for how this could be expressed:

1. Using λ_κ (e.g. with values like 10^7, $10^{7.5}$, 10^8)
    - This is very direct, but the numbers are large and do not have an obvious interpretation.
2. Using $S_\kappa = \log_{10} \lambda_\kappa$ (e.g. with values of 7, 7.5, 8)
    - The numbers are smaller and more manageable, but there is still no direct meaning to them.
3. Expressed relative to the default assumption for S_κ. (e.g. if the CMI assumption for S_κ is 7.5 then values of 7, 7.5 and 8 would be expressed as −0.5, 0 and +0.5).
    - The numbers are more intuitive, as positive means more smoothing and negative means less.
    - However the meaning of an assumption expressed in this way would vary if the CMI was to change its default assumption, so users may prefer an absolute rather than relative value.

We propose the second option, to use S_κ in absolute rather than relative terms, and to refer to this as the "period smoothing parameter".

Q 10.3. Do you agree that it is helpful for users of the Model to be able to control the responsiveness of the Model by varying the parameter S_κ?

Q 10.4. Do you agree that S_κ should be an Advanced rather than Core parameter?

Q 10.5. Which of the three options in Section 10.2 for expressing the smoothing parameter do you prefer?

We also propose that the smoothing parameters by age and cohort should be Advanced parameters, and that these should be expressed in the same form as for the period smoothing parameter. However the period smoothing parameter will be given more prominence, and typically included in the name of the Model, as described in Section 11.2.

Q 10.6. Do you agree that S_α, S_β and S_γ should be Advanced parameters?

10.3 Intermediate parameters

Currently the Model distinguishes between Core parameters (the long-term rate and the constant addition) and Advanced parameters (everything else).

Although the Advanced parameters can be varied to reflect different views, doing so typically requires setting a vector of parameters (e.g. age-period convergence periods for all 131 ages from 20 to 150). It may be helpful to have an intermediate level of control (e.g. to scale up or down all convergence periods with one parameter).

We are therefore seeking views of users as to whether they would find it helpful to have a standard "Intermediate" approach to adjusting certain parameters and, if so, what should be included. Our proposed list of Intermediate parameters is:

1. A flat addition to initial rates of age-period mortality improvements
2. A flat percentage scaling of age-period convergence periods
3. A flat percentage scaling of cohort convergence periods
4. A concise way to express a piecewise-linear shape to the long-term rate. This would specify the discontinuities in the shape, with rates being flat before the first discontinuity and after the last; e.g. with a long-term rate of 1.5%, the proposed shape would be denoted as "(1.5%@85,0%@110)".

These Intermediate parameters would not affect the range of choices that are available to users of the Model. Rather, they would make it easier to amend and express the existing Advanced parameters.

Q 10.7. Do you agree with the introduction of Intermediate parameters?

Q 10.8. Which of the four Intermediate parameters listed in Section 10.3 should be included?

Q 10.9. Are there any other Intermediate parameters that you would like to be included?

11. Version names

This section considers conventions for naming different versions of the Model.

11.1 Names of annual releases

Table 11.1 shows the names, release dates and calibration data for previous and proposed versions of the Model. Note that for CMI_2014 and CMI_2015 "Calibration data to" includes estimated data for the latter part of the final year.

Table 11.1: Model names

Name	Release date	Calibration data to
CMI_2009	25 November 2009	31 December 2008
CMI_2010	23 November 2010	31 December 2009
CMI_2011	16 September 2011	31 December 2010
CMI_2012	8 February 2013	31 December 2011
CMI_2013	13 September 2013	31 December 2012
CMI_2014	24 November 2014	31 December 2014
CMI_2015	28 September 2015	31 December 2015
CMI_2016	by 31 March 2017	31 December 2016
CMI_2017	by 31 March 2018	31 December 2017

The Committee has considered two options:

1. For most existing models, the year in the name corresponds to the year in which the Model was released. The exception is CMI_2012 which was released in early 2013 to allow for revisions to data following the 2011 Census. If this pattern was followed then the version to be released in March 2017 would be CMI_2017.
2. Option 1 would mean that there would be a jump in numbering of the Models, with no model called CMI_2016. Under Option 2 the Model numbers would be sequential, with the next version of the Model being CMI_2016. This would link the Model number to the final year of calibration data, at least from CMI_2014 onwards.

The Committee proposes to adopt Option 2.

Q 11.1. Do you agree with Model numbers being sequential, so the March 2017 release of the model would be CMI_2016?

11.2 Names of specific parameterisations of models

We recommend that Core projections of the current Model are referred to as "CMI_yyyy_x [a%] + c%" where:

- yyyy defines the version of the Model (as in Table 11.1)
- x is M for males or F for females
- a% is the long-term rate of mortality improvements
- c% is the constant addition to mortality improvements, and omitted if zero.

We propose to amend the naming convention to be “CMI_yyyy_x [a%;s]” where “s” is the value of the period smoothing parameter, S_κ.

Q 11.2. Do you agree with the naming convention “CMI_yyyy_x [a%;s]” where “s” is the period smoothing parameter, S_κ?

12. Summary of our proposals

This section summarises our proposals and assesses how it addresses the concerns in Section 3.3.

12.1 Proposed changes

The elements of our proposal for CMI_2016 that are materially different from current practice are:

1. Calibrating the Model to data for the United Kingdom rather than England & Wales; subject to verifying the feasibility of this with the relevant national statistical bodies.
2. Simplifying the method used to adjust exposure data. The new method retains the broad principle of the previous method – that underlying mortality rates are smooth, and outliers are indicative of artefacts in the data.
3. Defining mortality improvements in terms of $\log m_{x,t}$ but with results from the Model still being expressed using the existing definition in terms of $q_{x,t}$.
4. Using a new Age-Period-Cohort Improvement (APCI) model to determine components of historical improvements. This means that:
    a. we can fit historical mortality rates and determine mortality improvements, both in aggregate and split into age-period and cohort components, in a single step;
    b. the fitting process is much quicker in terms of run-time; and
    c. we can implement the Model entirely within Microsoft Excel using Visual Basic for Applications (VBA), which makes it more accessible to users.
5. Removing the "step-back" from the edges of the data when determining historical improvements and, instead, requiring the model itself to deal with the issue of stability.
6. Allowing and encouraging users to adjust the responsiveness of the Model to new data by using a single "period smoothing parameter".
7. Enabling users to express the pattern of convergence in terms of the slope of mortality improvements ("direction of travel") as an alternative to the current approach of proportion remaining at mid-point.
8. Tapering the long-term rate of age-period mortality improvements to zero between ages 85 and 110, rather than between ages 90 and 120.
9. Shortening convergence periods for the youngest cohorts.
10. Removing the "Constant Addition to Mortality Improvements" parameter, which we think is little-used.
11. Adding an Intermediate layer of parameters to make it easier to set and communicate certain Advanced parameters.

12.2 How our proposals address concerns

In Section 3.3 we described a number of concerns that had been expressed about the current Model. In this section we consider how our proposals address many of these concerns, and why we have not addressed all of them. Numbers in parentheses correspond to the numbering of the concerns in Section 3.3.

The improvements derived from the APCI model have terms in age, period and cohort which have equal prominence, and the smoothness of each can be controlled independently. This addresses (4) – that the use of age-cohort rather than age-period penalties in the p-spline model could exaggerate cohort effects.

The way that the improvements are derived under the APCI model avoids the explicit calculation and allocation of a residual (7). The split between age-period and cohort components is less subjective than under the current Model, depending on identifiability constraints that are consistent those used in the Cairns-Blake-Dowd M7 model (8) and – because of the integrated approach to fitting mortality rates and deriving improvements – the inconsistency of (14) is avoided.

The introduction of S_κ as a parameter allows users to control the impact of annual volatility (9) on the responsiveness of the Model (10), and avoids the need for a step-back (15).

The projection enables users to allow for the direction of travel (11) of mortality improvements, although no such assumption is made in the Core parameters.

The introduction of an Intermediate layer of parameters allows an easier way to set and communicate some of the Advanced parameters (17).

Finally, the entire calculation process, from raw data to projected results, can be achieved entirely within Excel, rather than requiring a combination of R and Excel (18).

12.3 Concerns not addressed

We have not directly addressed the question of basis risk (1) and (2) as we propose to continue to use national population data. However simplifications to the process and the software will make it easier for users to calibrate the Model to their own choice of dataset. We still have concerns about the quality of ONS exposure data (3) and the CMI is investigating this, partly through the High Age Mortality Working Party.

It is difficult to test the predictive power of the Model (5). We spent some considerable time on this issue at the outset of our review, but we moved towards a more pragmatic approach on the grounds outlined in Section 2.1.

We have not explicitly addressed potential inconsistencies between shapes of base tables and projected improvements (6), although this could be an area of future research.

We do not propose to harmonise the convergence periods for age-period and cohort improvements (12). As the length of convergence periods are somewhat subjective, we do not feel that we have strong enough evidence to move away from current practice, other than the changes described in Section 8.7, and we do not propose to address the inconsistency in (13). While it would be possible to do so, this would require fundamental change that would be disproportionate (e.g. requiring that convergence becomes asymptotic rather than over a finite period). We propose to still have constant age, period and cohort effects in the historical calibration period, while running them off in projection (16). While this could be perceived as a concern due to inconsistency it could equally be considered to be a valid choice in constructing the Model.

13. Other models considered

In the course of our research we have considered a wide range of options. The majority of this paper focusses on our proposals for CMI_2016, but this section briefly notes areas of research that we have considered but decided not to pursue, for now at least. The Technical Working Paper will contain more detail on some of these areas.

13.1 Alternative models for calibration

We considered a large number of alternative mortality models in order to determine initial rates of mortality improvements, including those in Cairns et al (2009). The main reasons for rejecting alternative models are:

- The lack of a cohort component.
- A cohort component that appeared to model non-cohort effects.
- Instability of fitted parameters when fitting to a different time period or age range.

We also considered variations of the APCI model, including the use of the terms ϕ_t and/or ψ_t in the formula:

$$\log m_{x,t} = \alpha_x + \beta_x(t - \bar{t}) + \kappa_t \ + \phi_t + \psi_x \gamma_{t-x}$$

The term ϕ_t is intended to explicitly model annual volatility, and to be distinguished from the smooth κ_t term by the use of identifiability constraints. While this approach had some promise, the distinction between κ_t and ϕ_t was somewhat subjective. Also, while use of the ϕ_t term allowed us to isolate annual noise, it was not clear that this made projections any more reasonable.

We refer to the effect of the term ψ_x as "age-moderation of cohort". Typically ψ_x would decline with age, so the importance of the γ_{t-x} term would decline as a cohort ages. This approach was promising for the APCI model, but required the inclusion of an approximate identifiability constraint. As for the annual volatility term, the fit of the model was improved by an additional term, but it was not clear that this led to better projections.

13.2 Integrated approach

Our proposed method retains the current Model's division between calibrating initial improvements and then using a separate procedure to project them forward. We considered a number of options to make this a single integrated process.

If we fit our proposed model to a time period including future years, with no data for deaths or exposures, then the time series κ_t would naturally be extrapolated to comply with the penalty function. For our proposed second-order penalty function, this would mean that period improvements would be constant in the projection period. We considered an alternative third-order penalty function that incorporates the long-term rate in order to force convergence. While this approach has some promise, it means that the choice of the long-term rate would affect views of historical improvements. There is some theoretical justification for this, but we felt that this would represent a material step away from current practice that could be seen as detrimental by many users.

As part of our investigations of an integrated approach we also considered a projection method using critically damped convergence that would remove the inconsistency noted in concern (13). In addition to being a departure from the current approach, this approach would mean that convergence would be asymptotic, with improvements tending ever-closer to the long-term rate, rather than being achieved over a finite period.

13.3 State space model

State space models characterise all the information known at time t in a vector of numbers together with equations that define how the state vector relates to the observed data and how it is projected forward. State space models have some desirable features, including:

- a very well-developed associated modelling literature,

- by their nature requiring clarity over what is and is not known at time t and hence what is and is not a prediction, which could help address concern (5), and
- automatically solving the issue of model inconsistency (13).

While some members of the Committee regard the state space framework as highly promising for predictive stochastic mortality models, given our time and resource constraints, the potential additional educational burden, and, in particular, the issues highlighted in Section 2.1, we decided to not pursue this strand of research.

14. Summary of consultation questions

This section contains a list of the consultation questions, most of which are included in earlier sections.

We encourage you to respond to the consultation whether or not you agree with our proposals, so that the responses that we receive are representative of all users. It would also be helpful to provide reasons for your responses, particularly where you disagree with the Committee's proposals.

Please send your responses to projections@cmilimited.co.uk. The same address can be used for any other correspondence regarding the Model. The Committee would like to receive responses by **30 September 2016**.

Q 3.1. Do you have any concerns about the current Model that are neither mentioned in Section 3.3 nor addressed by our proposals?

Q 5.1. Would you prefer the Model to be calibrated to data for England & Wales (as is done currently) or the United Kingdom (England, Wales, Scotland and Northern Ireland)?

Q 5.2. Do you agree with the proposed new exposure adjustment method?

Q 6.1. Do you agree with the use of mortality improvements in terms of $\log m_{x,t}$ within the APCI model, and in the definition of the long-term rate?

Q 6.2. Do you agree with the method, described in Section 6.3, of converting to mortality improvements in terms of $q_{x,t}$?

Q 7.1. Do you agree with the use of the APCI model to determine components of historical mortality improvements?

Q 7.2. Do you agree with the identifiability constraints used in the APCI model?

Q 7.3. Do you agree with the approach taken in smoothing the APCI model; i.e. the use of regularisation penalties with a different order used for the period terms?

Q 8.1. Do you agree that we should give users the option of specifying convergence as they choose: either in terms of direction of travel or proportion remaining at mid-point?

Q 8.2. If you agree with allowing users to specify convergence in terms of direction of travel, should the Core model allow for this or not?

Q 8.3. Do you agree with the proposed Core assumption, to taper the long-term rate to zero between ages 85 and 110?

Q 8.4. Do you agree with the proposal to amend the cohort convergence periods in the way described in Section 8.7?

Q 8.5. Would you be happy for the Model to only make projections for cohorts contained within the calibration data; i.e. those aged 20 and above at the end of the calibration period?

Q 9.1. Do you agree with the proposed Core assumption for S_κ of 7.5?

Q 10.1. Do you agree with the removal of the Constant Additional Rate of Mortality Improvement parameter?

Q 10.2. If you disagree with the removal of the Constant Additional Rate of Mortality Improvement parameter, do you agree with making it an Advanced rather than a Core parameter?

Q 10.3. Do you agree that it is helpful for users of the Model to be able to control the responsiveness of the Model by varying the parameter S_κ?

Q 10.4. Do you agree that S_κ should be an Advanced rather than Core parameter?

Q 10.5. Which of the three options in Section 10.2 for expressing the smoothing parameter do you prefer?

Q 10.6. Do you agree that S_α, S_β and S_γ should be Advanced parameters?

Q 10.7. Do you agree with the introduction of Intermediate parameters?

Q 10.8. Which of the four Intermediate parameters listed in Section 10.3 should be included?

Q 10.9. Are there any other Intermediate parameters that you would like to be included?

Q 11.1. Do you agree with Model numbers being sequential, so the March 2017 release of the model would be CMI_2016?

Q 11.2. Do you agree with the naming convention "CMI_yyyy_x [a%;s]" where "s" is the period smoothing parameter, S_κ?

Q 14.1. Do you have any other comments or questions for the Committee?

With your response, please also include:

- Your name
- Your contact details, in case we wish to clarify a response
- The name of your organisation (if any)
- The type of organisation (e.g. individual, insurer, reinsurer, pension fund, consultancy)

All responses will be shared with the Committee. A summary of responses and a list of respondents may then be made more widely available but will not attribute comments to particular companies or individuals.

References

CMI documents

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

CMI Working Paper 38: "A Prototype Mortality Projections Model: Part One – An Outline of the Proposed Approach" (2009)

CMI Working Paper 39: "A Prototype Mortality Projections Model: Part Two – Detailed Analysis" (2009)

CMI Working Paper 41: "CMI Mortality Projections Model: Feedback on Consultation and Issue of 'CMI_2009'" (2009)

CMI Working Paper 69: "The CMI Mortality Projections Model, CMI_2013, and feedback on the consultation on the future of the CMI Library of Mortality Projections and the CMI Mortality Projections Model" (2013)

CMI Working Paper 74: "The CMI Mortality Projections Model, CMI_2014" (2014)

CMI Working Paper 77: "Report of the Graduation and Modelling Working Party" (2015)

CMI Working Paper 80: "The release dates of future updates to the CMI Mortality Projections Model" (2015)

CMI Working Paper 83: "Recent mortality in England and Wales" (2015)

CMI Working Paper 84: "The CMI Mortality Projections Model, CMI_2015" (2015)

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

CMI Working Paper 91: "CMI Mortality Projections Model consultation – technical paper" (2016)

CMI: "User Guide for The CMI Mortality Projections Model: 'CMI_2015'" (2015)

All CMI Working Papers can be accessed via the CMI website: https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/numeric-listing-working-papers.

Please note that while this Working Paper is being made publically available, versions of the Model and some other outputs issued since March 2013 are restricted to those organisations and individuals who register as CMI users. Information on how to register is available on the CMI's website: http://www.actuaries.org.uk/research-and-resources/pages/how-access-cmi-outputs.

Non-CMI documents

Cairns AJG, Blake D, Dowd K, Coughlan GD, Epstein D, Ong A, and Balevich I (2009) "A Quantitative Comparison of Stochastic Mortality Models Using Data From England and Wales and the United States". North American Actuarial Journal 13:1 pp1-35.
http://dx.doi.org/10.1080/10920277.2009.10597538

Cairns AJG, Blake D, Dowd K, and Kessler A (2014) "Phantoms Never Die: Living with Unreliable Mortality Data". Working Paper.
http://www.macs.hw.ac.uk/~andrewc/papers/ajgc71.pdf

Currie ID (2013) "Smoothing constrained generalized linear models with an application to the Lee-Carter model". Statistical Modelling 13:1 pp 69-93.
http://smj.sagepub.com/content/13/1/69

Delwarde A, Denuit M, and Eilers P (2007) "Smoothing the Lee-Carter and Poisson log-bilinear models for mortality forecasting: a penalized log-likelihood approach". Statistical Modelling 7 pp 29-48.
http://smj.sagepub.com/content/7/1/29

Whittaker ET (1923) "On a new method of graduation". Proceedings of the Edinburgh Mathematical Society 41 pp3-73.
http://journals.cambridge.org/download.php?file=%2FPEM%2FPEM41%2FS001309150000359Xa.pdf

Disclaimer: This Document has been prepared by and/or on behalf of Continuous Mortality Investigation Limited (CMI). The CMI does not accept any responsibility and/or liability whatsoever for the content or use of this Document. Whilst care has been taken during the development of the Document, CMI does not (i) warrant its accuracy; or (ii) guarantee any outcome or result from the application of this Document or of any of CMI's work (whether contained in or arising from the application of this Document or otherwise). You assume sole responsibility for your use of this Document, and for any and all conclusions drawn from its use. CMI hereby excludes all warranties, representations, conditions and all other terms of any kind whatsoever implied by statute or common law in relation to this Document, to the fullest extent permitted by applicable law. If you are in any doubt as to using anything produced by CMI, please seek independent advice.

Copyright: You may reproduce the contents of this Document free of charge in any format or medium provided it is:

1. reproduced accurately and is unaltered;
2. not used in a misleading context; and
3. correctly referenced and includes both CMI's Disclaimer notice set out above and CMI's copyright notice, as follows:

© Continuous Mortality Investigation Limited.

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