NAPF Longevity Model

Supporting technical appendices

November 2014

Club Vita LLP

Introduction

Welcome to the technical appendices supporting the NAPF's longevity model report.

This document describes the key data, assumptions and analysis behind Club Vita and NAPF's collaborative research into longevity trends. As such, it is deliberately technical in nature as it is designed to provide confidence in the rigour of the research and the necessary supporting documentation to enable Scheme Actuaries, corporate actuaries and longevity consultants to be able to comfortable in referring to this work as part of forming their advice[1].

We start (**section 1**) by providing an overview of the data we have used in the research – including its origin, how we verified it, the types of data available to us, and crucially the data volumes used in our analyses. In order to maximise the insights we can gain from this data we have designed a practical approach to handle missing data (**section 2**). We also create a measure of deprivation that is comparable across all of our data (**section 3**).

In **section 4** we provide a brief history of longevity projections, before moving on to describe the model ('CMI projections model',) widely used by the pensions industry (**section 5)**.

We then take our first look at improvements for different groups of pensioners. To do this we need a consistent method for calculating life expectancies and associated statistical confidence intervals (**section 6**). We start with the improvements in life expectancy seen between different schemes (**section 7**) before digging deeper, segmenting defined benefit (DB) pensioners by such factors as pension size or socio-economics (**section 8**). Many of these factors are closely related – and so we need to identify which factors are most important to allow for (**section 9**). We describe how we grouped the data by these factors into a manageable number of groups (**section 10**).

Having established our DB pensioner groups – 'hard-pressed', 'making-do' and 'comfortable' – we then smooth the historical data and embed it into the approach widely used for the industry (**sections 11, 12 & 13**). This provides a starting point for trustees and sponsors seeking to reflect DB pension scheme data in their longevity improvement assumptions.

We conclude by looking at the financial impact of our results for different schemes, and for different scenarios for longevity improvements within each of our DB pensioner groups. The profile of members used for these assessments, and for our four example schemes are detailed in **section 14**; whilst the method, assumptions and approximations underlying for the impact assessments are set out in **section 15**. In **section 16** we describe the scenarios themselves.

On behalf of all the team we thank you for your interest in this research and we would be delighted to respond to any questions you may have.

Steven Baxter

steven.baxter@clubvita.co.uk

Matt Fletcher

matthew.fletcher@clubvita.co.uk

Steve Hood

steve.hood@clubvita.co.uk

[1] This document also complies with Technical Actuarial Standards on Data and Modelling

Reliances and Limitations

The National Association of Pension Funds (“NAPF”) and Club Vita LLP (“CV LLP”) have provided, to the pensions industry as a whole, both: an understanding of how differently longevity has been improving for different groups of DB pensioners (such as those at different ends of the deprivation spectrum); and materials that pension schemes, and their advisors, can use in practice to better inform the assumptions that are adopted for longevity trends (together, the “Research”).

The Research is based upon NAPF and CV LLP’s actuarial understanding of legislation and events as at November 2014 and therefore may be subject to change. The Research is NAPF and CV LLP’s understanding of how differently longevity has been improving for different groups of DB pensioners and is not, nor is it intended to be, specific to the circumstances of any particular pension scheme.

The information contained herein is therefore not to be construed as advice and should not be considered a substitute for specific advice in relation to individual circumstances. Where the subject of the Research refers to legal matters please note that neither NAPF nor CV LLP are qualified to give legal advice therefore we recommend that you seek legal advice. Neither NAPF or CV LLP (nor their respective licensors) accept liability for errors or omissions in the Research and neither NAPF or CV LLP (nor their respective licensors) owe nor shall accept any duty, liability or responsibility in regards the use of the Research except where we have agreed to do so in writing.

The Research contains copyright and other intellectual property rights of NAPF and CV LLP and their respective licensors. You shall not do anything to infringe NAPF or CV LLP’s or their licensors’ copyright or intellectual property rights.

If you are seeking to use the information contained in the Research after the date it was produced then please be aware that the information may be out of date and therefore inaccurate.

We recommend that you speak with your appointed longevity consultant and/or other professional advisers should you have any queries in relation to the Research.

Contents

Longevity Trends: Technical Appendices PAGE

It all starts with the data

1 Data underpinning our analysis

1.1 Club Vita dataset

The Club Vita database (VitaBank) is a pool of data of individual pension scheme member records, submitted by the participating schemes. This database (as at September 2014) consists of nearly 6 million member records; including:

- **Over 2.5 million** pensioners and widow(er)s;
- **1 million** deaths.

The records collected include personal, but non-sensitive information recorded by pension scheme administrators. This includes information relevant to predicting longevity, such as date of birth, sex, postcode, pension, final salary and retirement health.

1.2 Data pre-processing

Only data which has been through our initial quality control process enters the statistical analysis. The data quality control process is designed to ensure the data for each pension scheme is as reliable as possible. However it also recognises that the quality of the data is often dependent on historic record keeping processes and so may have some inherent shortcomings.

A suite of checks are carried out on the data received to ensure it is correct and reliable, and where necessary corrections are made if possible. Where a member record has a predictor which our checks suggest is unreliable it is excluded from analysis. We also check for concentrations of unreliable records within schemes and biases in exclusions between living and deceased records and limit a scheme's inclusion in our analysis where there is a risk of bias.

Ensuring a complete history of deaths

We recognise that some schemes may not have a complete record of deceased pensioners prior to some point in time. For each scheme we have determined an "*earliest useable date*" (EUD) – the date from which we believe we have a complete history of deaths.

The mortality data we receive includes experience data up to a date shortly before it was extracted from the pension scheme's administration system. As such it is liable to 'incurred but not reported deaths' i.e. an understatement of deaths in the most recent weeks of the extract as a result in the delay in reporting deaths.

In order to ensure that mortality rates are not underestimated we carry out similar analysis to that described above to verify the point up to which we believe we have full and complete death data. This leads to a "*latest useable date"* (LUD) for each scheme, which is used to right censor the data (i.e. no observations of survival beyond this date are included in our analysis). Typically the latest useable date excludes between 1 and 2 months' worth of data.

Since we are analysing mortality by calendar years, we need to take care to avoid seasonal biases resulting from including part years therefore we have for these purposes restricted our analysis for each scheme to the period from the first 1 January on or after the EUD to the last 31 December on or before the LUD for each scheme.

When analysing the patterns in longevity by specific factors, for example pension amount, we also check whether we have complete information on that factor from the EUD onwards. Where this is not the case we use a factor-specific EUD for that scheme.

1.3 Data extract used in this analysis

Through the help of the NAPF we obtained access to an additional 500,000 UK pensioner records, provided by schemes who wished to participate in the research project. These pensioners (and associated deaths) are included in the numbers quoted on the previous page.

This additional data was subject to the same checks as were performed on the data supplied by existing Club Vita participants. Data was only taken forward to the final dataset if it was found to be of sufficient quality.

Exposed to risk & deaths

Club Vita collects data annually from each of its subscribers, with these data feeds spread over the calendar year. As such it is regularly refreshed with the latest longevity data.

For the purposes of our analysis we have focussed on an extract of the database as at July 2014 throughout. The charts (right) shows the pattern of (pensioner and dependant) 'exposed to risk'[2] and deaths over time for men (blue bars) and women (pink bars) within the data analysed in this report.

We can see how:

- The exposures increase over time reflecting
  - schemes within the Club having reliable data starting at different points in time due to historical administration practices;
  - the maturation of pension schemes leading to larger numbers of pensioners
  - There is a step-up in 2001 – the point at which a number of the larger schemes first have reliable data
- The deaths follow a similar pattern to the exposed to risk.

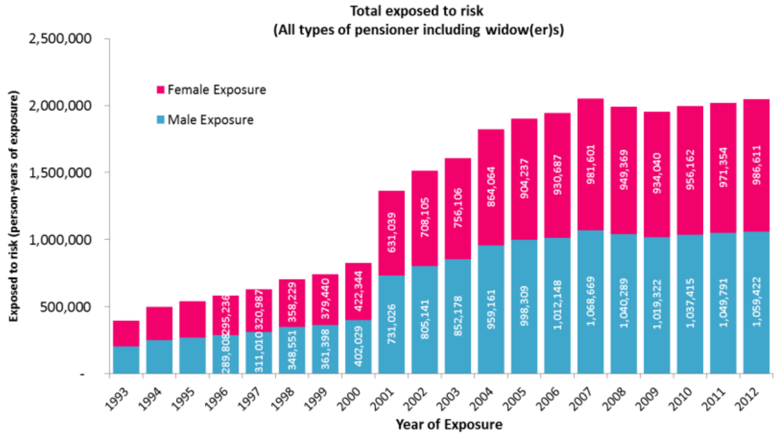

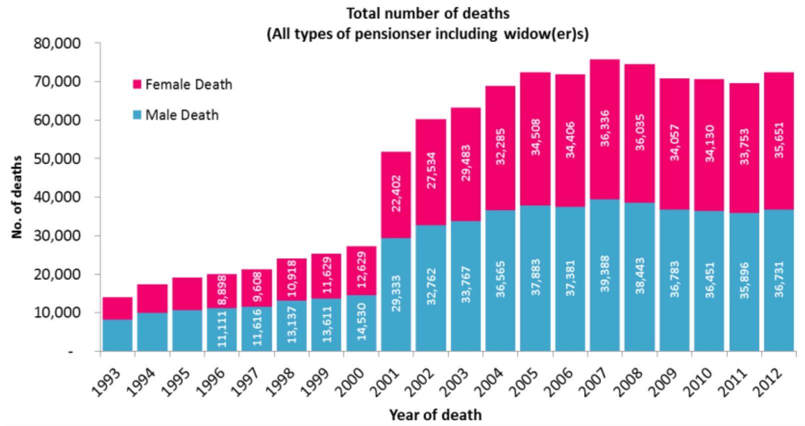

---

[2] Broadly speaking a measure of the number of lives in each year but adjusted to allow for the fact some individuals were only in the analysis for part of that year.

1.4 Key rating factors

By collecting information at the individual level, VitaBank contains a wide range of rating factors potentially relevant to both baseline mortality and improvements coming through over time. These rating factors include gender, retirement health, pensioner type (pensioner or dependant), postcode based socio-economic measures (such as Index of Multiple Deprivation), affluence (pension and salary), age and occupation (manual and non-manual)[3]. We briefly discuss some of the rating factors used in our analysis below.

Postcode based measures

Club Vita's lifestyle measure

The Club Vita lifestyle measure uses an individual's full postcode to assign him or her to one of seven different groups, labelled A-G.

These groups have been calculated using geo-demographic data provided by a specialist third party provider (CACI) which maps each residential UK postcode onto a demographic type. These different types have then been condensed using statistical clustering methods into 7 different lifestyle categories which are predictive of material differences in longevity. Our group A relates to those with the 'worst' lifestyles in the sense of having the shortest life expectancy, whilst group G relates to those having lifestyles linked to the longest life expectancies.

[3] See Madrigal et al (2012) for more detail on how Club Vita have determined the key ratings factors for mortality levels.

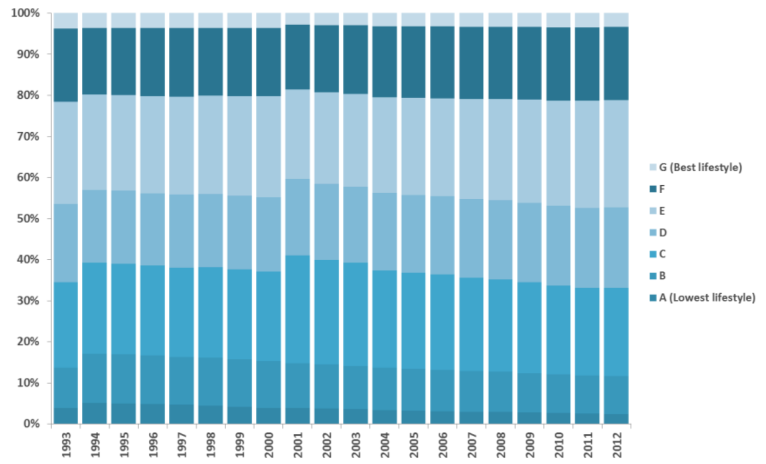

The chart above shows the split of our data (men and women combined) between the 7 lifestyle categories. Notice how a small proportion of the data belongs to the two most extreme lifestyle groups i.e. group A and G.

The Club Vita lifestyle categorisation has the advantage of using the full individual postcode. However our research is aimed at providing a widely applicable analysis of longevity trends. As such we have also considered a variety of publicly available postcode-based demographic measures; accepting that this results in a trade-off between availability and granularity – as publicly available measures typically cover much broader geographical areas than an individual postcode (circa 15-20 houses).

Index of Multiple Deprivation

The statistics agencies of each of the nations within the UK measure the deprivation of local areas via an index which captures multiple indicators, typically including such factors as income, employment and crime.

The scores are publicly available at a fairly detailed local level. For example within England they are available for regions known as 'Lower-layer Super Output Areas' (LSOAs) which typically cover around 35,000 houses. However, they are not directly comparable across countries within the UK. Accordingly we have used a method to generate an index which spans all of the UK. This method is detailed in **section 3** and the resulting UK-wide index of deprivation is available from our website (www.clubvita.co.uk).

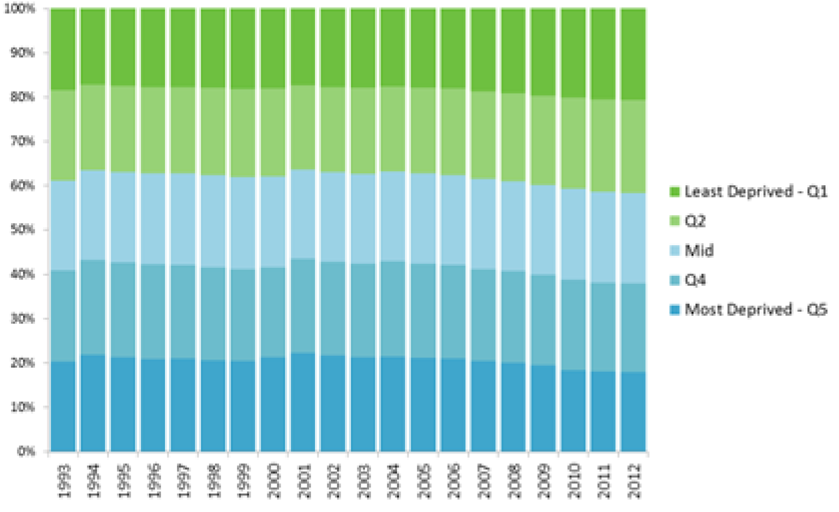

We can see from the chart of the split of our data (men and women combined) between the quintiles of these scores that, despite the groups being determined at the UK population level, we find broadly 20% of our data in each group.

Low income families

We have also considered two alternative methods to measure deprivation.

The first of these, 'low income families' score reflects the proportion of children living in families in receipt of out-of-work benefits or tax credits and where reported income is less than 60 per cent of UK median income. This measure is easily accessible[4] and covers the whole of the UK.

For the purposes of our analysis we have split the UK areas into five groups based upon their rankings for this score, running from the 20% of areas with the lowest scores (Q1) to the 20% with the highest scores (Q5).

Area classification score

The second alternative deprivation measure we have used is the Area Classification Score (ACS). This is a publicly available[5] form of geo-demographic profiling produced by the ONS based on 2001 census data which splits LSOAs in England & Wales – and their equivalent, Datazones, in Scotland – into 7 super-groups, 20 groups and 53 subgroups.

For the purposes of our research we have focussed on the super-groups:

1 Countryside
2 Professional City Life
3 Urban Fringe
4 White Collar Urban
5 Multicultural City Life
6 Disadvantaged Urban Communities
7 Miscellaneous built-up areas

4 https://www.gov.uk/government/publications/personal-tax-credits-children-in-low-income-families-local-measure

5 http://www.ons.gov.uk/ons/guide-method/geography/products/area-classifications/ns-area-classifications/index/index.html

The chart below shows how the data (men and women combined) splits into these 7 super-groups.

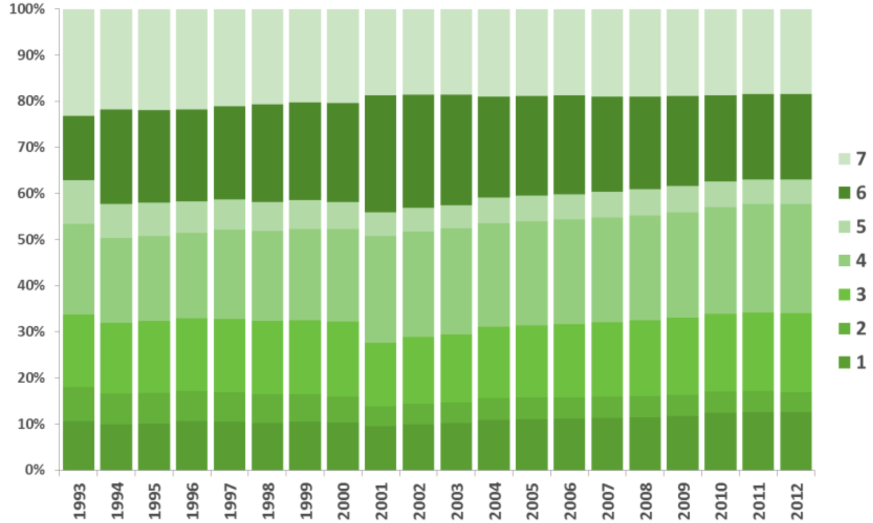

Affluence measures: Pension and salary

The Club Vita data contains two measures of affluence: pension and last known salary.

Pension size can be a poor proxy for overall affluence as it depends not only on earnings but length of service in the pension scheme – a modest pension could arise from long service on low pay, or very short service on high pay. However, whilst salary is a better measure of affluence, pension will almost always be available, whereas salary may be harder to extract from some pension scheme records.

To allow for inflation both pension and salary are revalued from their as at date to a common date (1 July 2013) in line with RPI. For deceased pensioners the revaluation of pension amounts are performed using a proportion of RPI (below 100%) for broad consistency with the pension increases paid historically to surviving pensioners which will typically be a mix of full RPI, limited price inflation and nil increases. (The same approach is not required for live pensioners as their pensions will usually be recorded at a recent date.)

The charts below show the distribution of pensions and salary amounts within our data (men only).

Pension

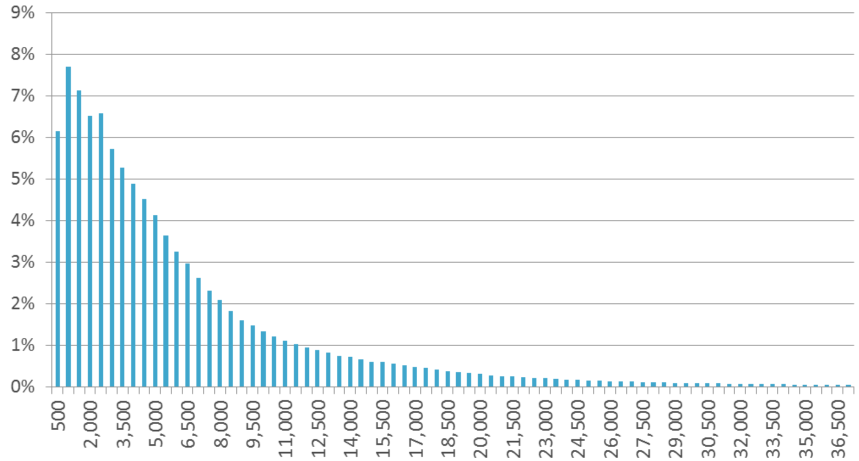

Salary

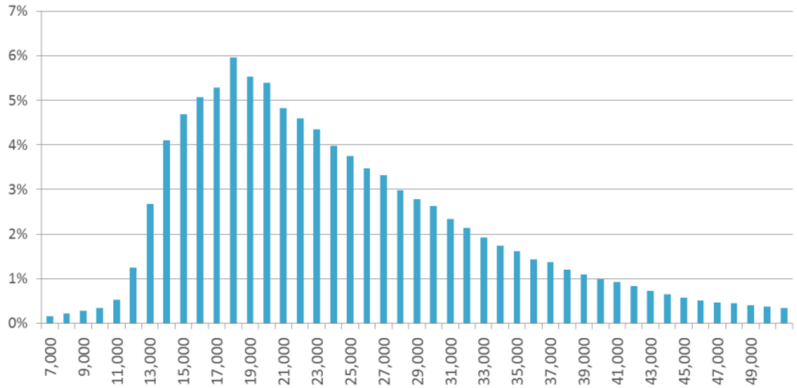

2 Making maximum use of available data

In section 1 we discussed how the scheme data used in our analysis has undergone a thorough data quality control process, to determine what data will be used in the onward analyses and ensure reliability of data. This is done both at the scheme level and at the covariate level (so for example a particular scheme may have reliable postcode data but suspect pension amounts in a particular year).

It is important to maximise the data used in our analysis. This will help minimise volatility and also reduce the sensitivity of the analysis to the experience of individual schemes (although this is not a material issue given the overall data volumes).

In this section we set out the process that we have adopted to ensure that we have maximised the available data without compromising on overall data quality.

2.1 Exploring 'missing' data

As part of our analysis we look at the impact of different covariates on life expectancy. In our analysis, we need to exclude data which falls below the required quality threshold on the covariates used to divide the data. As such, we are reducing the available data as a result of excluding members who fail the quality checks that are applied.

Levels of unknown covariates can be expected to increase as we go further back in time (due to having less stringent administration standards historically, records not being updated, etc.). In particular these issues are more likely to affect deaths (i.e. higher levels of unknowns), so there is the possibility that we could be biasing the results by excluding more deaths relative to living pensioners in a given calendar year.

At a scheme level, the proportions of 'unknowns' is again likely to increase as we go back in time, until, in some cases, reaching the 'trigger' level – the point in time before which no exposures are included (the EUD discussed in section 1).

There is, therefore, a growing risk of understating rates of mortality historically (if we exclude more deaths than lives, we are reducing the mortality rate). This will have a knock on effect on mortality improvements, which will again be lower than their 'true' level, due to historical mortality rates being lower.

The chart below illustrates, for a sample scheme, how the proportion of 'unknown' pension amounts for male lives and deaths varies over time.

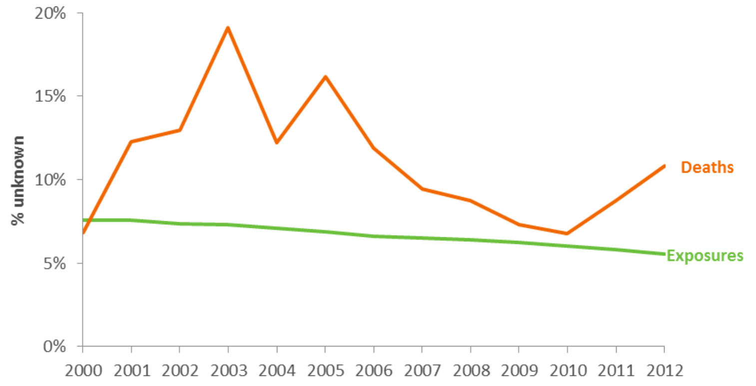

We can see that, in this case, the proportion of unknown deaths is more volatile than exposures, but generally the proportion of unknown deaths and lives are both increasing as we go further back in time.

We have sought to overcome this issue by reallocating 'unknown' data – the section below illustrates the method using males as an example.

2.2 Adjusting for missing male pensioner data

In our analysis for men we have divided the data using both pension and adjusted IMD (see section 10). We have sought to maximise the amount of data used by re-allocating lives and deaths with 'unknown' covariates across the covariate groups, as follows.

We initially take the (cleaned) submitted data, and allocate individual members (lives and deaths) to the appropriate pension bands and IMD quintiles (including 'unknowns' for each covariate as appropriate).

For each age and calendar year we therefore have four distinct categories of member:

- where both pension and IMD are known;
- where pension is known and IMD is unknown;
- where pension is unknown and IMD is known; and
- where both pension and IMD are unknown.

The following tables show the relative levels of exposures and deaths in each of these four groups for male pensioners (for 1993 to 2012 and ages 60 to 95).

Exposure

		IMD	
		Known	Unknown
Pension	Known	96.3%	3.2%
	Unknown	0.5%	0.0%

Deaths

		IMD	
		Known	Unknown
Pension	Known	95.4%	4.1%
	Unknown	0.5%	0.0%

Where one of the covariates is unknown, then the exposures and deaths for the group are assumed to be spread across the unknown covariate in the same proportions to where the covariate is known – e.g.

- where pension band is unknown (but IMD is known), the spread of "unknown pension band, known IMD" exposure and deaths across the different pension bands matches the spread of "known pension band, known IMD" (for the given IMD) for exposure and deaths respectively
- where pension band and IMD are both unknown the spread of "unknown pension band, unknown IMD" exposure and deaths across the different IMDs and pension bands matches the spread of "known pension band, known IMD" for exposure and deaths respectively.

This minimises the risk that the mortality rates as measured over time are polluted by any imbalances in data coverage between lives and deaths.

2.3 Smoothing by age

The proportions to use for spreading data across the unknown covariate are volatile from one age to the next. To smooth this out, we average across the 5 year age bracket centred on each age when determining the ratios.

Example

If we consider the example of the group with unknown pension band, where IMD is 1, then we have the following (for each age and year):

IMD 1, Pen band unknown

IMD 1, Pen band 1-5

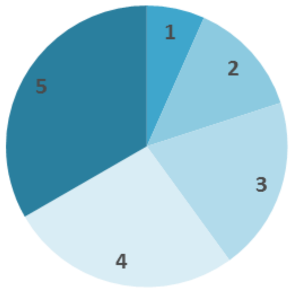

IMD 1, Pen band unknown - reallocated

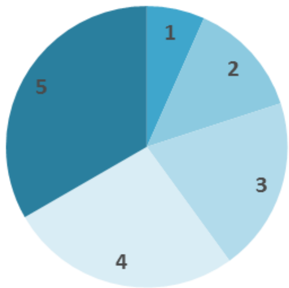

So the exposure in group “IMD 1 Pension Band 1” following reallocation is:

$$IMD_1, PB_1 + IMD_1, PB_{unknown} * \sum_{x-2}^{x+2} \frac{IMD_1, PB_1}{\sum_j IMD_1, PB_j}$$
$$+ IMD_{unknown}, PB_1 * \sum_{x-2}^{x+2} \frac{IMD_1, PB_1}{\sum_i IMD_i, PB_1}$$
$$+ IMD_{unknown}, PB_{unknown} * \sum_{x-2}^{x+2} \frac{IMD_1, PB_1}{\sum_{i,j} IMD_i, PB_j}$$

2.4 Adjusting for missing female pensioner data

The same approach of reallocating unknowns was also applied for women. However as we only use one covariate – IMD – for women, the calculations are less complex than for men, although the levels of unknowns are similar.

	IMD	
	Known	Unknown
Exposure	96.5%	3.5%

	IMD	
	Known	Unknown
Deaths	95.1%	4.9%

2.5 Impact of reallocation

The following charts show the increase in exposures and deaths, for men and women, as a result of the re-allocation process set out above.

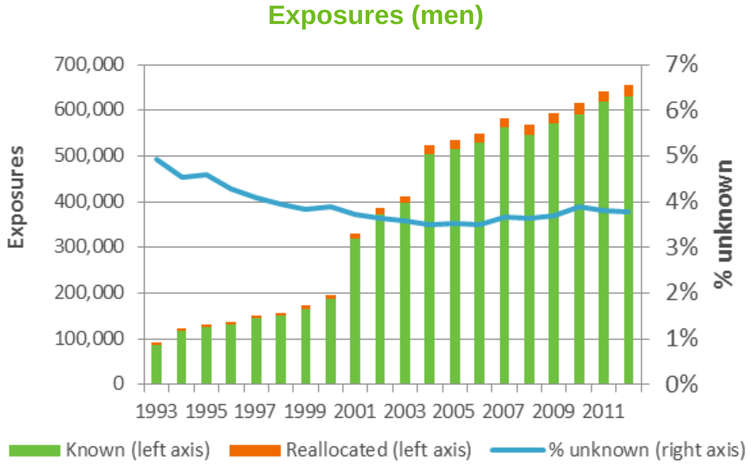

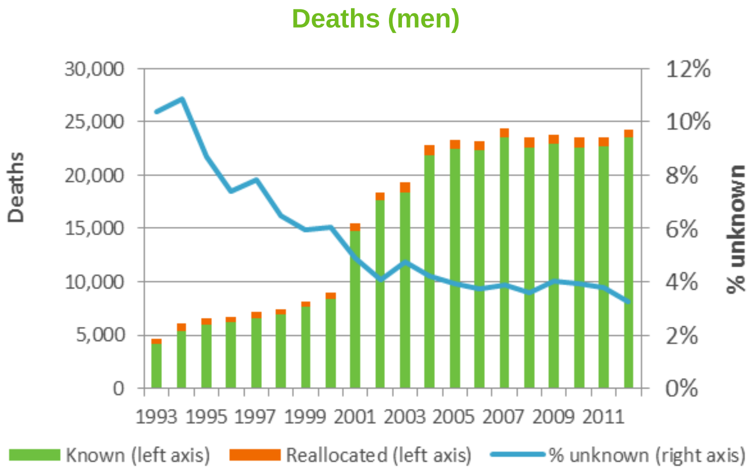

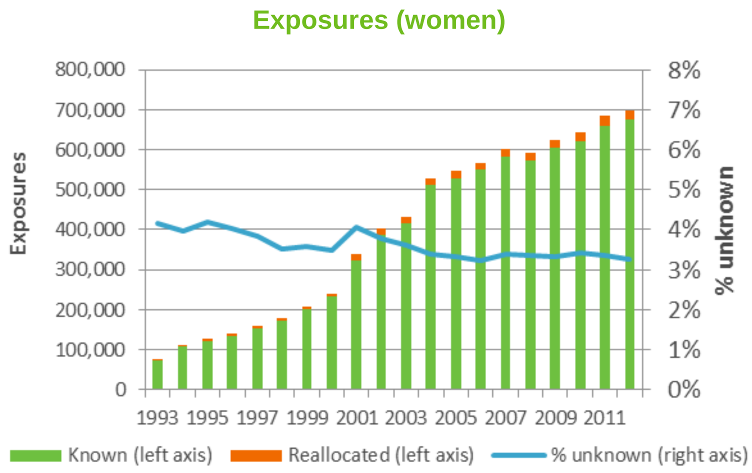

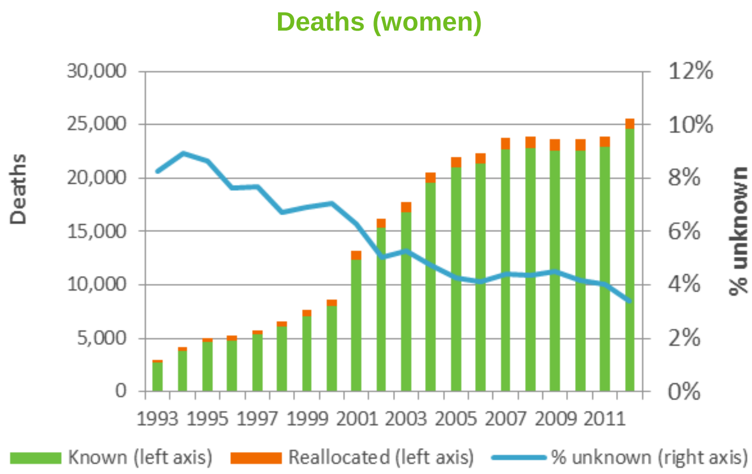

We can see from these charts that the levels of reallocation are relatively low. Prior to 2001 the levels of reallocation increase in percentage terms, particularly for deaths.

The impact on our results of this reallocation is relatively small. However we can be confident that we have removed a possible area of bias in our analysis of historic improvements.

3 Ensuring we can compare IMD data across all of UK

3.1 The challenge

In order to analyse how life expectancy has changed for different groups of lives we need ways of segmenting our data by different socio-economic and demographic measures.

The statistics agencies of each of the nations within the UK measure the deprivation of local areas via an index which captures multiple indicators, typically including such factors as income, employment, crime, etc. This Index of Multiple Deprivation (IMD) is publicly available. However, one challenge of this measure is that it cannot be used with pension schemes which have membership living in more than one of the UK's constituent countries. This is because the index is country specific – with the weighting to the different factors varying from country to country (and indeed in some countries factors are included which are not included in other countries). Further, many of the factors are measured relative to the country-specific average value.

3.2 Our solution – a UK-wide measure of deprivation

We have followed a recognised method[6] to calculate an Index of Multiple Deprivation which can be used across the UK.

The method works on the following principles:

1 Choose a small number of factors used in the indices for multiple deprivation of the constituent countries which:

  1.1 Are used in the calculation in all countries

  1.2 Have significant weighting in the calculation of each country's IMD

  1.3 Have the underlying data used in generating the index published

  The published work focusses on Income & Employment which we will also restrict our attention to.

2 Carry out a linear regression of the IMD score for each country against these factors

3 Rebase the values for each country by:

  3.1 Choosing a base country

  We have used England as this dominates the UK data.

  3.2 Using the regression coefficients for that base country, along with the observed values of the factors, to calculate revised IMD for each area in each of the other countries

  3.3 Adjusting these calculations by the residuals from the individual country fit (i.e. the extent to which there is a component not captured by the chosen factors), standardised to the variability seen in the residuals of the base country.

Formulaically

We can express the above process formulaically in two stages.

Linear regression

Fit, for each country C, a two factor linear regression model:

$$IMD_{Ci} = \alpha_C + \beta_C I_{Ci} + \gamma_C E_{Ci} + \varepsilon_{Ci}$$

Where:

[6] The method was published by the ONS in Health Statistics Quarterly 53, Spring 2012, under: "UK indices of multiple deprivation – a way to make comparisons across constituent countries easier"

- I_{ci} is the value for the income factor in area i of country C
- E_{ci} is the value for the employment factor in area i of country C
- α_C, β_C and γ_C are the country specific regression coefficients
- ε_{Ci} is the residual value for area i of country C (i.e. the difference between the fitted and actual IMD value)

Generate 'adjusted' IMD scores

Calculate the scores for each country C, other than England, as:

$$IMD_{Ci} = \alpha_{England} + \beta_{England} I_{Ci} + \gamma_{England} E_{Ci} + \varepsilon_{Ci} \frac{\sigma_{England}}{\sigma_C}$$

Where σ_C is the estimated standard deviation of the residuals for country C. (For England the existing scores are used.)

Domain scores or values?

The different factors used in calculating the index of multiple deprivation are known as domains. In each case a *score* is calculated for the domain (typically, but not always, on a scale of 0 to 1). An exponential transform is applied to these scores to create domain *values* – converting the rankings of the scores to values on a scale from 0 to 100. It is these domain values which are weighted to get the IMD score for each area.

We have the option to use domain scores or values in our regression (as both can be sourced from the individual statistical authorities). Ideally we would want to use whichever of these values which:

- Best mirrors the distribution of the IMD statistic
- The assumptions underlying linear regression (e.g. normally distributed residuals) holds best for

By visual examination of plots of the respective distributions and Q-Q plots for residuals, domain values appear to meet the above criteria only slightly better. However, these also require more stages to the calculations for those seeking to replicate our work with future data publications. As such we have elected to use domain scores in our work.

3.3 Results of fitting the model

Regression coefficients

The fitted coefficients from our regression analysis (along with R^2) scores are presented in the table below.

Country	α_C	β_C	γ_C	σ_C	R^2
England	-0.190	0.849	0.930	3.585	0.95
Scotland	-1.491	0.831	0.865	2.779	0.97
Wales	-4.333	0.972	0.583	3.363	0.95
Northern Ireland	-6.601	0.720	0.761	2.855	0.97

Reassuringly we see that each of the countries has a high R^2 value – this is a statistical measure of goodness of fit and indicates a very modest proportion (2% in the case of Scotland) of the variation in IMD values is explained by factors other than income and employment.

Adjusted IMD values for each country

A natural comparison to make is how the revised IMD values compare to the original published values, and the extent to which we have changed the ordering of areas within different countries. On the following pages we do this in two ways – firstly by comparing a plot of the adjusted and original IMD values (which should be clustered around the diagonal), and secondly by showing the movements between quintiles ***within*** each country. *Since the adjusted index for England[7] is identical (by construction) to the original index we focus on Scotland, Wales and Northern Ireland below.*

[7] https://www.gov.uk/government/statistics/english-indices-of-deprivation-2010

Scotland

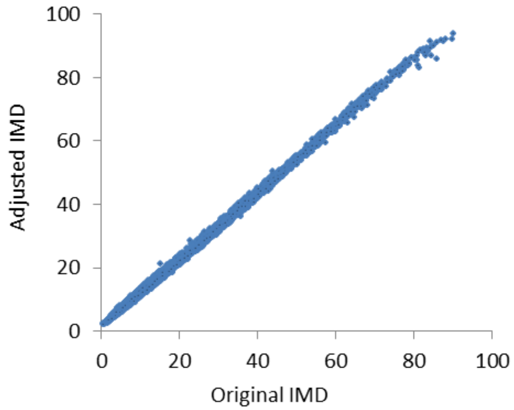

Original IMD	Adjusted IMD – quintile (1=lowest) 1	2	3	4	5
1	1,250	51	0	0	0
2	51	1,185	65	0	0
3	0	65	1,192	44	0
4	0	0	44	1,238	19
5	0	0	0	19	1,282

Whilst there is some change to the values it is reassuring to see that very few areas have moved between quintiles.

Data sources:

http://www.scotland.gov.uk/Topics/Statistics/SIMD/background2simd2009

Wales

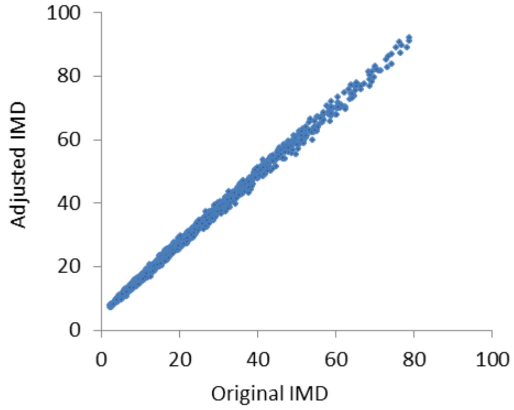

Original IMD	Adjusted IMD – quintile (1=lowest) 1	2	3	4	5
1	363	17	0	0	0
2	17	342	20	0	0
3	0	20	343	16	0
4	0	0	16	351	12
5	0	0	0	12	367

Reassuring the adjusted IMD is very consistent with the original IMD values.

Data sources:

http://wales.gov.uk/statistics-and-research/welsh-index-multiple-deprivation/?lang=en

Northern Ireland

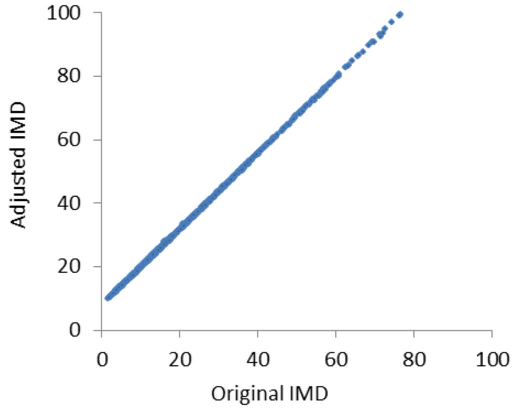

Original IMD	Adjusted IMD – quintile (1=lowest) 1	2	3	4	5
1	176	2	0	0	0
2	2	175	2	0	0
3	0	1	173	3	0
4	0	0	3	175	0
5	0	0	0	0	178

Again, reassuring the adjusted IMD is very consistent with the original IMD values.

Data sources:

http://www.nisra.gov.uk/deprivation/nimdm_2010.htm

The industry approach to modelling improvements

4 A brief history of longevity projections

Before considering how longevity trends may evolve in the future it is useful to consider how they developed in the past, and how this has compared to the longevity projections made at various points.

The chart below (focussing on the England and Wales population) shows the significant increase in observed lifespans since the 1950s.

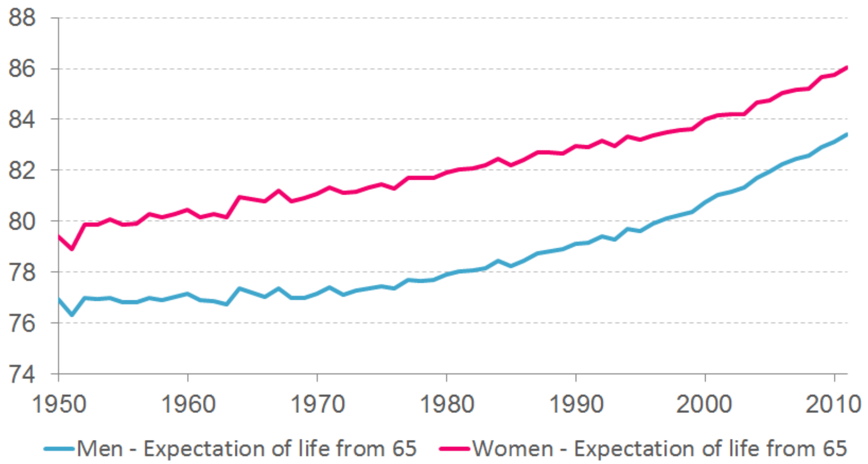

Period life expectancies from age 65 increased by over 6 years for both men and women, with particularly rapid increases in recent decades. Indeed, the majority of those increases occurred in the last 25 years for women and last 15 years for men.

Throughout this time pension schemes (and their actuaries) have needed to make assumptions about how the trends would evolve. The evolution of typical assumptions (for male pensioners) from the early 1990s through to the start of this decade is shown below.

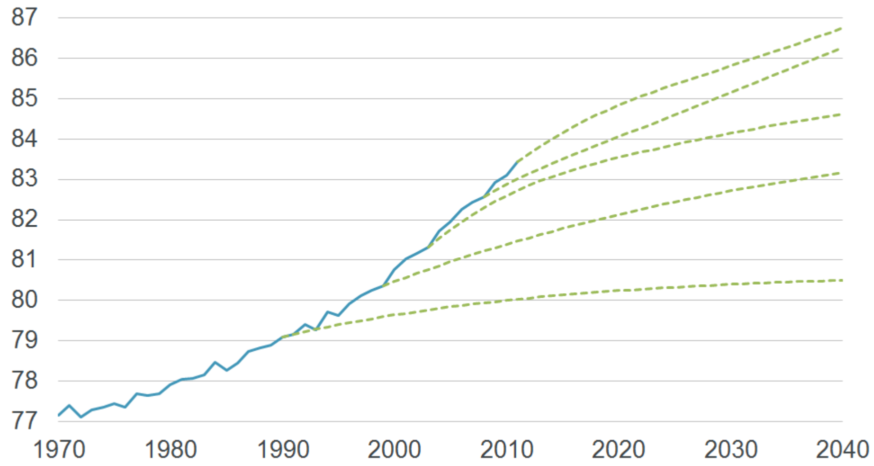

The dotted lines illustrate typical trend assumptions adopted from 1990 (specifically the “80 series” projections), through to those adopted in the late 1990s and 2000s (the 92 series, medium cohort projection, and medium cohort with 1% underpin) up to those typically used in recent years (the 2011 CMI projection model with a 1.5% long term rate, discussed in more detail in Section 5).

The Office for National Statistics (and formerly the Government Actuaries’ Department) also produces longevity projections roughly biennially, for use in population projection. Comparing these projections to actual improvements in longevity has also shown a consistent under-estimation of improvements over time (although more recently, projections have been much more in line with actual improvements).

It is clear that lifespans have repeatedly increased more rapidly than projected, leading to the successive revisions to assumptions.

The changes over the last decade alone have reduced funding levels by 10%, and by around 20% if we look back 2 decades. So the risk is clear; schemes could see further strains emerge from longevity in the future.

We can also draw the following conclusions:

- Longevity projections are a necessary tool to help set the pace of funding for pension schemes, but they cannot provide certainty of outcome.
- Improving the accuracy of projections would be of benefit to pension schemes and their sponsors. Using the most relevant data – that from defined benefit schemes – and understanding the impact of factors such as socio-economic group on trends – would inform trustees' understanding of the risks and issues.
- Whatever projections are made, the outcome is likely to differ from any best estimate assumption. Hence it is valuable to consider a range of possible future scenarios, to gain an understanding of the risks associated with longevity trends.

5 The ‘industry standard’ – the CMI model

5.1 Introduction

Many of the well adopted models discussed in Section 4 have been provided by the Continuous Mortality Investigation (CMI), part of the Institute and Faculty of Actuaries (IFoA). The CMI seek to produce an industry wide starting point from which professionals can easily understand and communicate the approach adopted.

5.2 The CMI Model

Currently the most widely used model for mortality improvements in the actuarial industry is the CMI mortality projections model (the ‘CMI Model’), first published in November 2009. With this model the CMI sought to improve the realism and flexibility of projection models, reflecting more closely the improvements experienced to date without reducing the flexibility or simplicity of the model.

We have illustrated our results using various calibrations of the CMI model of mortality improvements because it is widely used, flexible and well understood within the industry.

5.3 Description

The CMI Model is a deterministic model driven by user inputs, based on the assumption that current rates of mortality improvements converge over time to a single[8] long-term rate (which is typically lower than the current, historically high, improvement rates).

The model has been updated roughly annually to reflect emerging experience, with the version including 2013 (and partial 2014) experience data (CMI_2014) published November 2014. For the purposes of this project, we have used the version of the model published in September 2013 (CMI_2013) because the start-point for projecting improvements (in 2010) is the same as the start-point for improvements in our dataset.

There are broadly three parts to the longevity improvement model:

- Initial rates of improvement
- Long-term rate of improvement
- The “pathway” connecting the short term and long term

We consider each of these separately below.

Each input to the model can be adjusted by the user to better reflect their views on the current level and likely future path of longevity improvement. The flexibility offered in the model allows the user to change many aspects of the improvements, although in the pension scheme context it has typically been used in ‘Core Parameters’ mode, where all factors bar the long-term rate are pre-defined.

Within the main report, we illustrate the impact on life expectancy and scheme liabilities of changing each of these features.

[8] Technically the model reduces the user input long term rate to 0 at the oldest ages

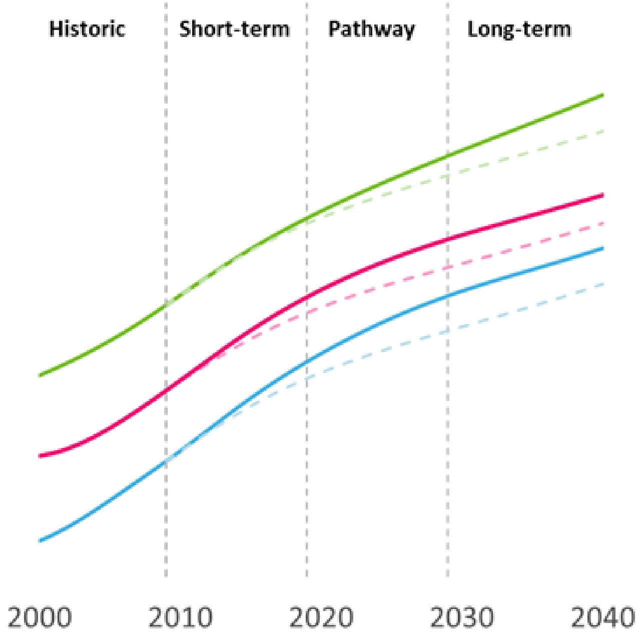

Initial rates of improvement

The CMI_2013 model takes in crude mortality rate data for the England and Wales population for ages 18 to 102 over the period 1961 to 2012. From this, raw rates of improvement are derived, which are then smoothed by fitting a statistical model.

England and Wales population data was used to produce the default initial rates of mortality improvements, principally due to the lack of a single alternative dataset that would otherwise meet the needs of all users.

The resulting smoothed rates are the default tables for the Initial Rates of Mortality Improvement within the model; these rates are then split into two component parts (Age/Period and Cohort).

Splitting the Initial Rate of Mortality Improvements into two component parts was well-supported by research, paralleled the structure of the Interim Cohort Projections and the approach adopted by the GAD and ONS within the National Population Projections

We have been able to collect data on DB pensioners that can be used to produce starting rates that are more relevant to pension schemes. We have followed a similar smoothing and splitting approach to that used by the CMI on this data. Further details are set out in section 11.

Long-term rate of improvement

Within the model, the user is required to define a rate of improvement that will apply in the long term. This is typically lower than the current historically high rates.

The Long-Term Rate ("LTR") assumption is considered to be the single most important parameter for users of the CMI model to set and this is the only input to have no default proposed assumption. A higher long term rate implies that mortality rates are reducing more year on year, this in turn means that life expectancies are higher.

The Pension Regulator's recent publication on annual funding statistics (www.thepensionsregulator.gov.uk/docs/scheme-funding-2014.pdf) shows that 81% of schemes with valuations in Tranche 7 (effective dates between 22 September 2011 – 21 September 2012) use the CMI Projections Model and of these, 62% use a long-term rate of improvement of 1.5% per annum. It is worth noting that a significant minority of schemes using the interim cohort projections also use a 1.5% per annum underpin to the rate of improvement.

This suggests that there is a de facto 'standard' future mortality improvement assumption within the industry, namely the CMI model with a 1.5% long term rate.

We do not comment on the suitability of a 1.5% long-term rate; however, we want to show some alternative futures as 'food for thought' to illustrate the range of outcomes.

Convergence

The 'pathway' between the Initial Rates of Mortality Improvement and the LTR is controlled by two sets of parameters:

- the period of convergence (how long it takes to get to the LTR); and
- the proportion of convergence remaining at the mid-point of the convergence period (the rate at which the LTR is achieved).

Behind the Core Parameters setup of the model are default assumptions for the length and shape of convergence.

The period to convergence varies by age and by cohort, and is capped at 40 years.

In some of our scenarios, we consider the impact of lengthening or shortening the period.

The proportion remaining at midpoint is set by default at 50%. If this default assumption is adopted, and the long-term rate is lower than current rates, then mortality improvements will start to fall immediately. This can be seen from the chart below (the green line illustrates the default assumption), extracted from CMI Working Paper 39.

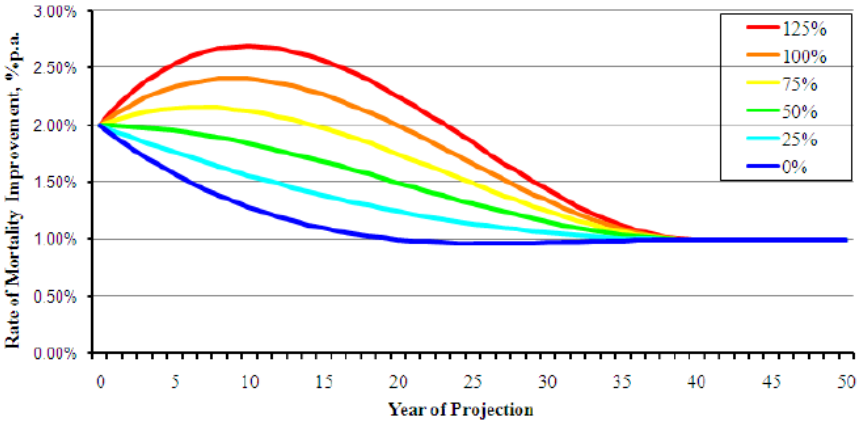

Figure 5.2: Illustration of the operation of the convergence formula over a 40-year period, with various Proportions of Convergence Remaining at Mid-Point

It can be seen that a higher proportion means that the initial rate of improvement increases before falling toward the assumed long term rate (assuming that the long term rate is higher than the initial rate). Therefore, having a higher proportion will result in higher life expectancies.

To make some allowance for the 'direction of travel' of improvements over the short to medium term, in some of our illustrative scenarios we use higher and lower proportions remaining – see for example scenario 1.1 "Health Cascade".

5.4 Other comments

There are some potential weaknesses in the CMI model which include

- the underlying population,
- issues with population data

For example, it is well known that the 1919 birth cohort (and, to a lesser extent, the cohort born shortly after the Second World War) has a particularly idiosyncratic pattern of births, which causes problems with calculating mortality rates and improvements (Phantoms Never Die: Living with Unreliable Mortality Data www.macs.hw.ac.uk/~andrewc/papers/ajgc71.pdf). There have also been historic concerns relating to unreliable population projections from census data for those aged over 90.

The CMI have recently undertaken a consultation and will be investigating some of these known issues.

Initial results

6 Calculating life expectancy for our initial analyses

6.1 Introduction

There are a number of different factors which are known to impact on mortality rates, based on analysis of historical mortality rates (e.g. age, pension amount, salary, etc). In section 8 we explore a number of these variables, and examine the impact that each can, in isolation, have on life expectancy (known as 'univariate analysis' as we only consider one variable). In section 9 we then go on to consider how best to model mortality improvements using combinations of these factors (known as 'multivariate analysis').

6.2 Calculating life expectancies

In order to consider the impact of a particular factor on life expectancy (sections 7-9) we need a method for calculating both the relevant life expectancies over time and the appropriate confidence intervals to place on those life expectancy calculations (so we can consider whether any differences in life expectancy are statistically significant)[9].

The method we use is a slightly modified version of the approach known as the Chiang method[10], which is widely used by the ONS (who provide a template spreadsheet model[11] which estimates life expectancy at birth from population data) and others.

Please note that the rest of this section provides details of the formalities of the Chiang method. It is relatively heavy on formulae.

Details of the 'Chiang' calculation

The standard Chiang method, as implemented in the excel template available from the ONS, can be summarised as follows:

- Group ages into 'buckets' of age intervals.

  The template uses 19 distinct age intervals to cover the full age spectrum (<1, 1-4, 5-9, 10-14, …, 80-84, 85+).

  Where:

  x = youngest age in age interval (e.g. 65 for 65-69 interval)

  n = number of ages in age interval (e.g. 5 for 65-69 interval – the 85+ interval is assumed to have 11 years)

  a_x = fraction of age interval at which deaths in the interval are assumed to die (0.5, i.e. half way through the interval, except the <1 interval, where it is 0.1)

- For each age interval, obtain details of the population and deaths in that interval.
- Calculate M_x (the age specific death rate) as deaths/population in the interval.
- Calculate q_x (the probability of dying in the interval) as

$$q_x = \frac{n*M_x}{(1+n*M_x*(1-a_x))} \qquad \text{for } x < 85$$

$$q_{85} = 1$$

- Calculate p_x (the probability of surviving the interval) as

$$p_x = 1 - q_x$$

[9] Note that in Section 16 the life expectancies are based upon fitted mortality rates and so are calculated directly from these values using the usual method i.e. $e_x = 0.5q_x + p_x(1 + e_{x+1})$

[10] Chiang C L *The Life Table and its Applications* (1984)

[11] http://www.ons.gov.uk/ons/rel/subnational-health4/life-expec-at-birth-age-65/2004-06-to-2008-10/ref-life-table-template.xls

- Calculate l_x (the life table, starting from 100,000) as

$$l_x = l_{x-1} * p_{x-1}$$

- Calculate d_x (the deaths in the life table for the age interval) as

$$d_x = l_x - l_{x+1}$$

- Calculate L_x (the number of years lived in the age interval) as

$$L_x = \mathrm{n} * \left(l_{x+1} + (a_x * d_x)\right) \qquad \text{for } x < 85$$

$$L_{85} = \frac{l_{85}}{M_{85}}$$

- Calculate T_x (the cumulative number of years lived in the age interval and subsequent intervals) as

$$T_x = T_{x+1} + L_x \qquad \text{for } x < 85$$

$$T_{85} = L_{85}$$

- Calculate e_x (the life expectancy at the start of the interval) as

$$e_x = \frac{T_x}{l_x}$$

How we used the Chiang method

We applied the method described above with:

- 8 age intervals were used, starting from 60-64, with the highest interval being 95+ (as supported by our data)

  Thus the specific calculations for the 85+ period in the description above now apply to 95+, e.g. $q_{95} = 1$

- n was set to 5, except for the top band where it was set to 10 for men and 12 for women

Amendments to the standard Chiang calculations

The Chiang method relies on the user having mid-year population estimates in order to calculate 'central' death rates, and from these deduces annual probabilities of death. In contrast we are working with start year population numbers (and 'initial exposed to risk'). This is a technical distinction which requires some changes to the method.

Specifically

- For the 95+ bucket

$$q_{95} = 1$$

$$M_{95} = \frac{q_{95}}{\left(\mathrm{n} - n * q_{95} * (1 - a_{95})\right)}$$

- For buckets apart from 95+ (assuming sufficient exposure):
  - Calculate M_x as

$$M_x = \frac{deaths}{exposures - \frac{deaths}{2}}$$

  - Calculate q_x as

$$q_x = \frac{n * M_x}{\left(1 + n * M_x * (1 - a_x)\right)}$$

In addition we modified the method to be appropriate for our circumstances where we are wishing to be able to compare mortality between different groups of lives:

- The a_x values were solved to be appropriate for the intervals, and the curvature of mortality rates over the interval meaning that using the ONS implementation of $a_x = 0.5$ provided numbers inconsistent with accurately calculated life expectancies

- Calculation of q_x direct from the underlying data is subject to the age group having minimum exposure levels

Each of these is detailed further below.

Setting the a_x values

Within each age interval deaths are assumed to occur at a particular point in the age interval (a_x in the formula above). Therefore the values chosen for a_x can have a material impact on the resultant life expectancy calculations. The ONS implementation assumes that the deaths occur on average halfway through the interval, however this is unlikely to hold when mortality rates rapidly increase over the age interval as is the case at older ages. We therefore assessed the appropriate a_x to use in each age interval.

This was done by constructing a life table using the q_x rates from age 60, and in each age interval working out a suitable value for a_x by equating the standard Chiang method for calculating the number of years lived by the deaths with the same calculation using a one year approach, which assumes that the average survival period for deaths in a one year period is half a year. This reduces the potential for distortions and ensures that we achieve the desirable feature that the resultant life expectancies are similar to those that would result from calculating life expectancy using the 'standard' life expectancy calculation based on the full mortality table[12].

This reduces to equating the following calculations for the average number of years lived by the deaths within each age interval (where the age interval runs from x_0 to x_1:

- Chiang method:

$$number\ of\ years\ lived\ by\ deceaseds = d_i n_i a_i = a_i n_i \sum_{x_0}^{x_1} d_x$$

- One year approach:

$$number\ of\ years\ lived\ by\ deceaseds = \sum_{x_0}^{x_1} d_x \left(x - x_0 + \frac{1}{2}\right)$$

This leads to us using:

$$a_x = \frac{\sum_{x_0}^{x_1} d_x \left(x - x_0 + \frac{1}{2}\right)}{n_i \sum_{x_0}^{x_1} d_x}$$

Note: In our calculations we used the q_x values from the most recent calibration of mortality tables to the Club Vita data and then used the same values of a_x throughout our calculations.

Allowing for sparse data

As we are looking at various ways of subdividing the data into a number of subgroups, and then grouping into age buckets, we can in some (very rare) instances have cells in the calculation which have very low levels of exposure, particularly at older ages. This can lead to misleading mortality rates and false conclusions.

In such (very rare) cases, it is necessary to adjust the calculation method above to use a sensible mortality rate for cells. Which we take to be the average rates from the CV dataset. Formally we make the following adjustments in these cases:

- Take q_x^{single} to be based on a mortality table produced by Club Vita
- Calculate q_x as

$$q_x = 1 - \left(1 - q_x^{single}\right)^n$$

- Calculate $\mathrm{M_x}$ as

$$M_x = \frac{q_x}{\left(\mathrm{n} - n * q_x * (1 - a_x)\right)}$$

[12] i.e. $e_x = 0.5q_x + p_x(1 + e_{x+1})$

6.3 Confidence intervals

Life expectancies

The confidence interval around the life expectancy value calculated by the Chiang method is found as follows:

- The variance of q_x is

$$Var(q_x) = \frac{n^2 * M_x * (1 - a_x * n * M_x)}{Population * (1 + (1 - a_x) * n * M_x)^3}$$

- Then the variance of e_x is

$$Var(e_x) = \frac{\sum\left(l_x{}^2 * \left((1 - a_x) * n_x + e_{x+1}\right)^2 * var(q_x)^2\right)}{l_x{}^2}$$

Further details of the derivation of these formulae can be found in Chiang's monograph *The Life Table and its Applications*. Using the variance derived above, and an assumption that the variability in life expectancy is normally distributed, we can readily calculate confidence intervals for the life expectancies.

Changes in life expectancy

In order to estimate the confidence interval of the difference in life expectancy we introduce an additional assumption that the life expectancies involved are independent.

The standard deviation of the change in life expectancy, can then be calculated by taking the square root of the sum of the variances of the life expectancies at the start and end of the period.

$$sd(e_x - e_y) = \sqrt{var(e_x) + var(e_y)}$$

7 How life expectancy has varied between different schemes

We have used the methods described in section 6 to explore the variation in life expectancy (from age 65) for each of the schemes in the combined Club Vita and NAPF dataset.

The chart below shows, for over 40 of the largest schemes[13] in the dataset the period life expectancy based upon observed mortality over the period 2008-2012 (and so in effect a smoothed value for life expectancy in 2010).

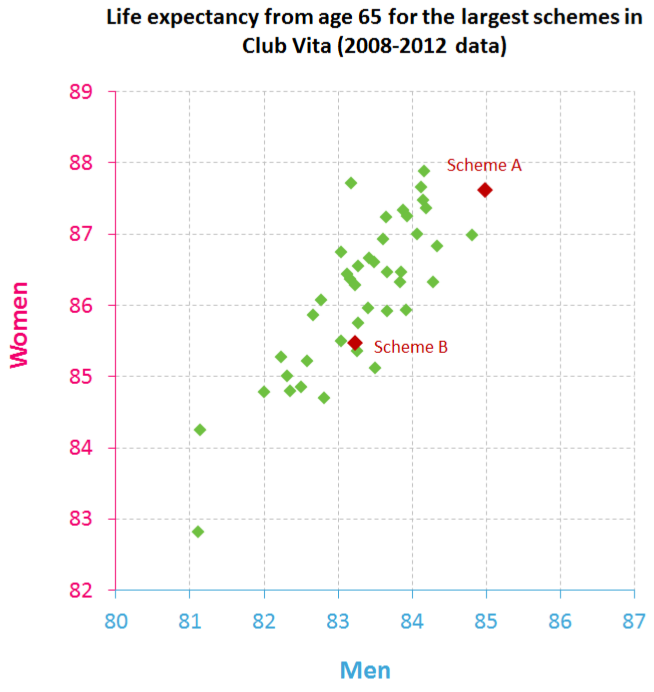

We can see how there is considerable variation in life expectancy – spanning 4 years for men, and 5 years for women. These variations though are largely well understood by pension scheme trustees, sponsors and their advisors, and as such are routinely incorporated into funding valuations.

In contrast variations in *improvements* in life expectancy are less well understood. We have calculated the corresponding period life expectancy (and confidence intervals therein) for the period 1998-2002. By comparing the change in life expectancy between these two points in time we have an estimate for the increase in life expectancy (in effect) between 2000 and 2010, as illustrated in the chart below.

[13] Selected to have sufficient data volumes for men and women, and across the older age spectrum, so that have a tight degree of confidence in the calculated life expectancies.

When looking at this chart please note fewer schemes are shown as not all the schemes have a reliable back history (or have sufficient data at all ages in early years) to enable us to plot the increase in life expectancy with confidence.

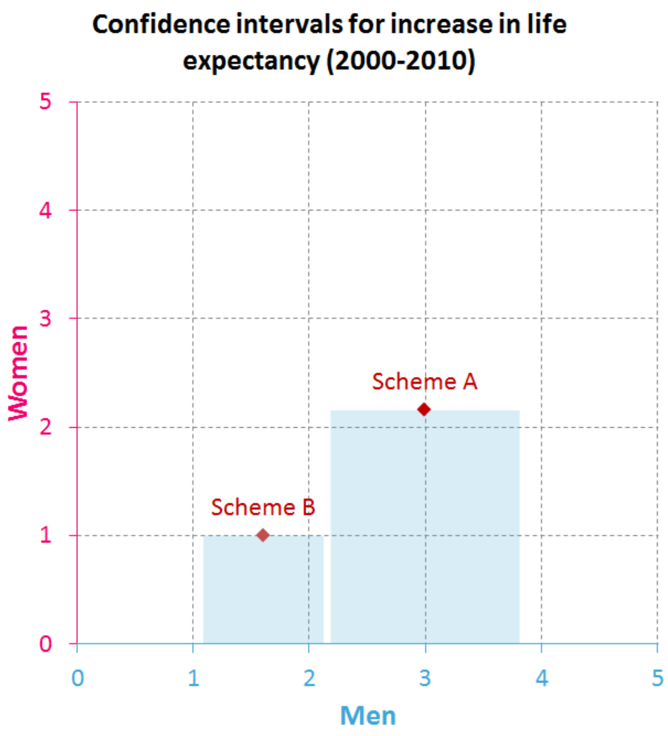

In both of the charts on the preceding pages we have highlighted two schemes, Scheme A & Scheme B. These represent two schemes which have had, for men, some of the highest and lowest increases in life expectancy over 2000-2010. Over this period Scheme A saw life expectancy for men rise by 3 years, compared to 1.6 years for Scheme B. Both are material rises in life expectancy, but for Scheme A the extra improvements in life expectancy equate to more than 5% extra in liabilities.

Of course there is some uncertainty in the measurement of life expectancy and some of the differences seen for Scheme A and Scheme B could be due to noise. The chart to the right shows – by the width of the blue shaded bars – the 95% confidence intervals[14] for the change in life expectancy for men, for each scheme.

If Scheme A and Scheme B had been selected at random from the wider dataset, the fact that these bars are non-overlapping would indicate that there is a statistically significant difference between the two sets of improvement. However, because Scheme A and Scheme B were picked specifically from the wider dataset as schemes whose improvements were different from each other, there is still a possibility that this difference is due to statistical 'noise'. The magnitude of the differential liability impact on the two schemes, though, encourages us to proceed further with our investigation into how life expectancy may have changed differently for different types of individuals within DB pension schemes.

[14] Calculated as per methods in 6.3

8 Exploring historical improvements in life expectancy

In this section we explore how life expectancy has increased between 2000 and 2010 for a variety of different groups of individuals within our data.

In each case the life expectancies and the differences therein, have been calculated using the approach described in section 6. The life expectancies are based upon data averaged over 3 years i.e. the 2000 figure uses data spanning 1999-2001 and the 2010 figure data spanning 2009-2011.

8.1 An introduction to our charts

The chart below illustrates as the increase in period life expectancy at age 65 over the decade from 2000 to 2010 for four different industries

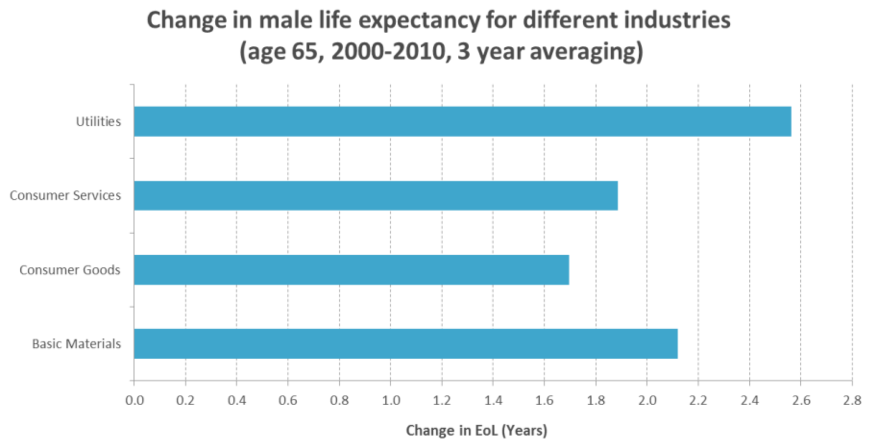

We can see there is a 0.9 year difference between the lowest (consumer goods) and highest (utilities) increases in life expectancy – however is this a significant difference? Or is it simply down to the uncertainty in measurement of the life expectancies in 2000 and 2010 for each industry?

In order to answer this we need to calculate confidence intervals for the individual life expectancies, and the differences therein.

The chart below illustrates, for the same industry types, the 95% confidence intervals around the increase in life expectancies in the previous chart.

We can see that in this case Consumer Services have seen lower increases than Utilities. The non-overlapping confidence intervals provide us with considerable confidence that the different increases in life expectancy are not simply due to random variations.

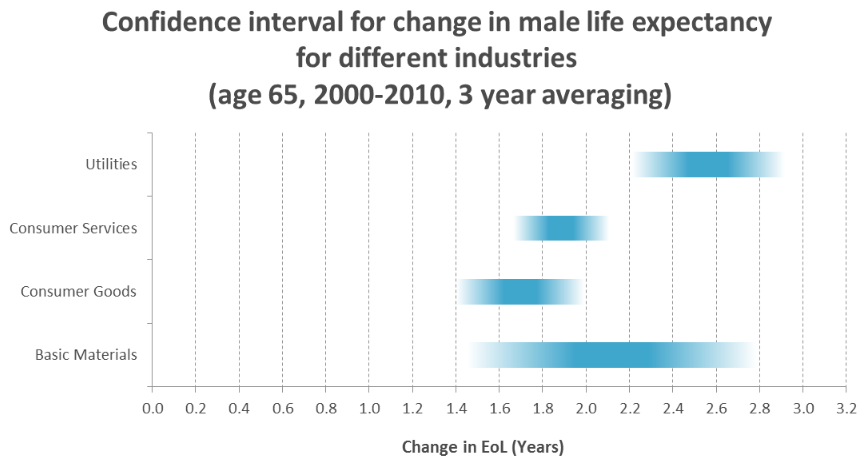

8.2 Results for a selection of variables

We have set out below the results of univariate analysis on a range of covariates (pension amount, salary and deprivation).

Pension amount

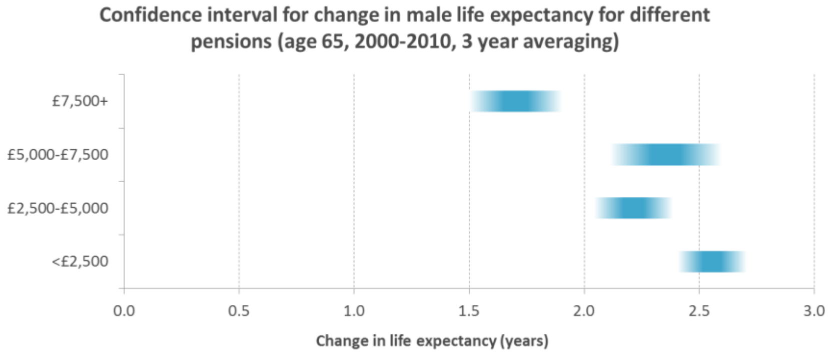

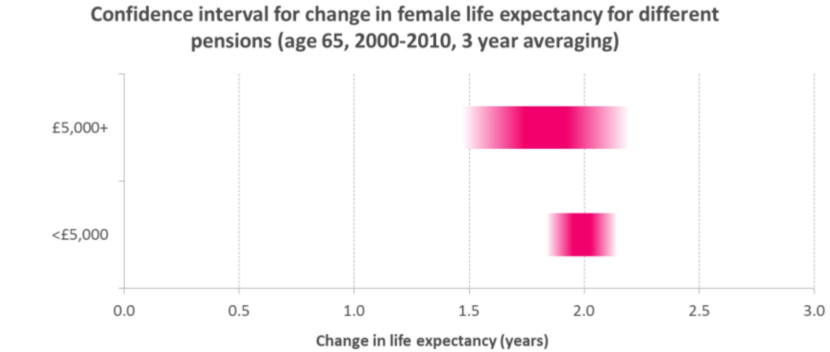

We can see:

- clear differences in life expectancy improvements for men by pension income.
- no clear differences for women

Salary

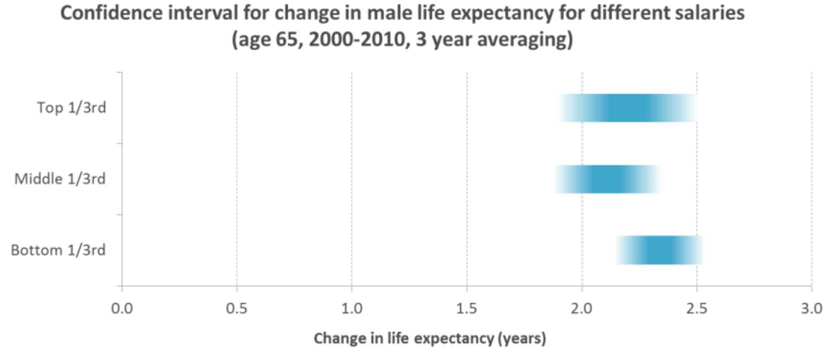

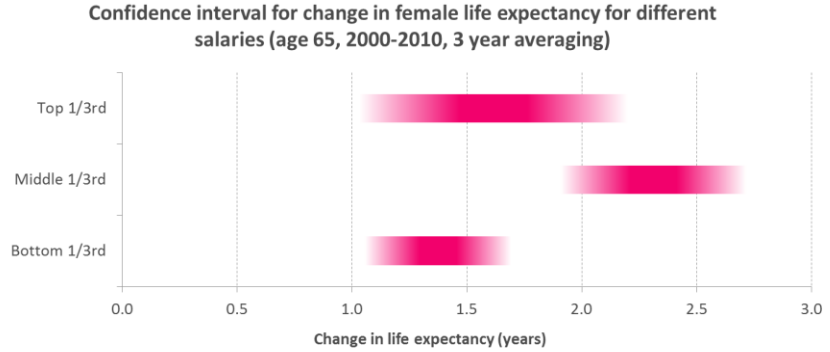

For both men and women we see some indications of an affluence effect when looking at salary (based on splitting the data into equally sized groups).

Club Vita LLP

Deprivation (England)

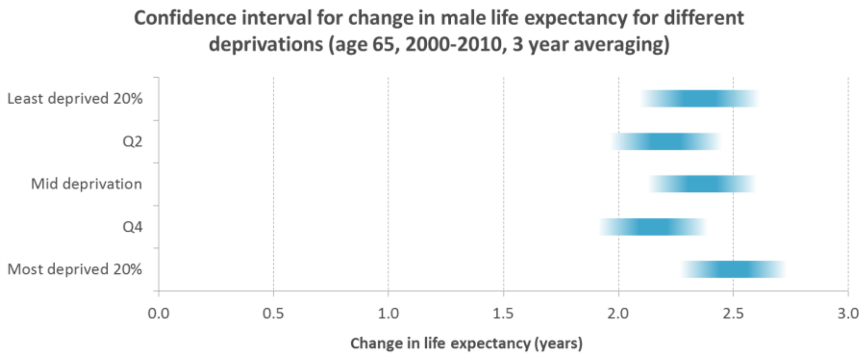

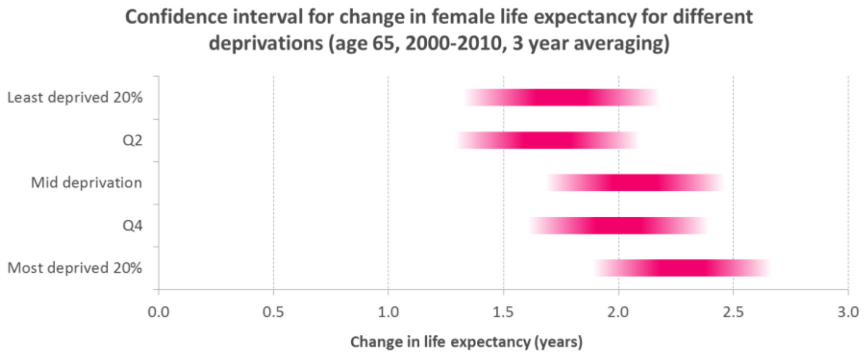

For men the biggest improvements in life expectancy have been seen in the most deprived areas. For women we see a clear gradient with the biggest improvements again seen in most deprived areas.

Public v Private Sector

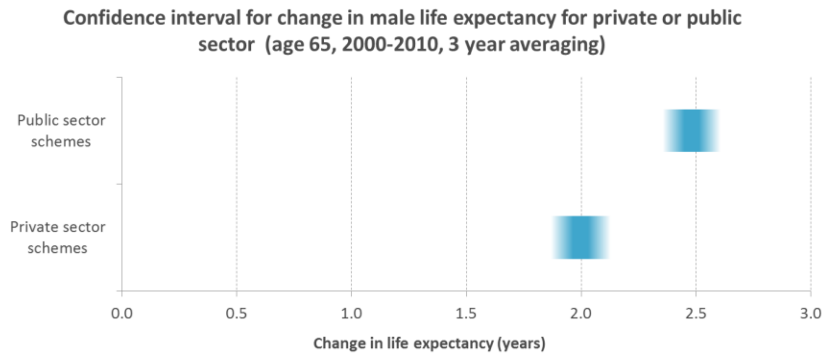

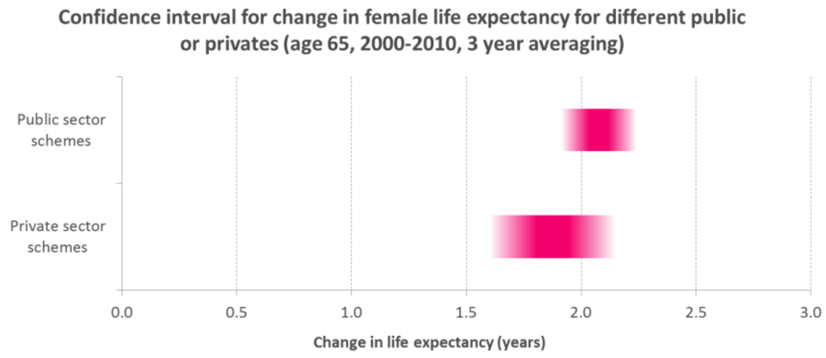

Over the last decade public sector schemes have seen larger improvements in life expectancy than private sector schemes, particularly for men.

9 What factors best capture historic improvements?

9.1 Introduction

We saw in section 8 how there have historically been clear differences in improvements between different groups of individuals within pension schemes. For the purposes of creating projections we want to identify a manageable number of groups which have had clearly different trends historically, and for which we reasonably believe will have different trends in the future. A natural starting point is to ask ourselves if there is a single factor which clearly segments individuals such as pension or postcode, or whether we need to group people by a combination of factors such as pension *and* postcode.

9.2 Narrowing down the range

From the Club Vita dataset we have identified a wide range of possible variables we could use to identify groups of individuals. These are illustrated in the schematic below.

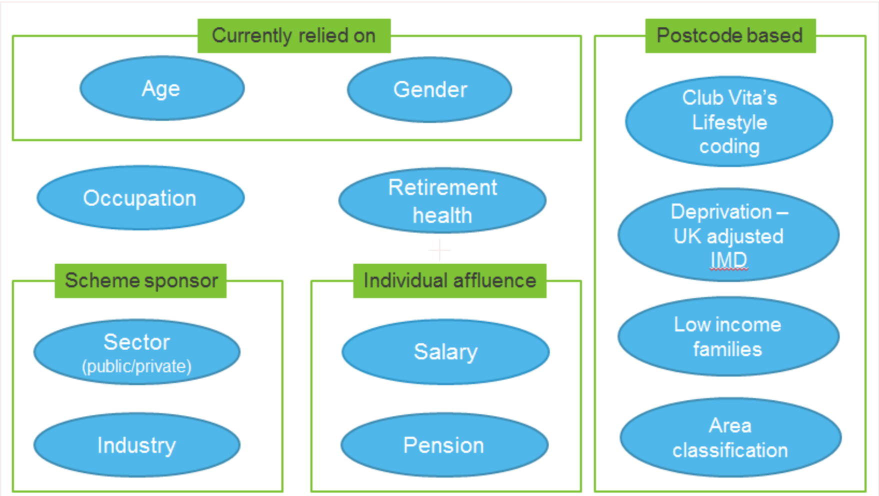

It is helpful to narrow this range down a little first.

One key consideration is that we want the results of our research to be widely applicable across pension schemes. Accordingly we can remove some factors from the list:

- **Occupation:** Whether an individual is performing / performed a primarily manual or non-manual role is one traditional socio-economic measure. However this has historically only been recorded in certain sectors and many schemes are not able to segment their membership in this way.
- **Salary:** A modest pension can be achieved through short service / high income, or long service / modest income. Comparisons of pension between schemes are also 'noisy' owing to different accrual rates, salary offsets etc… Accordingly last known salary tends to be a better measure of affluence than accrued pension – and one we generally prefer to use for measuring baseline longevity. However while salary is in our experience readily accessible for around 70% of pension schemes, we prefer to focus on pension for modelling longevity improvements as it is (almost) always available.
- **Retirement health:** Not all schemes record this – and where they do, ill health retirement can have very different meanings between schemes (from a 'bad back' preventing ability to do *current* role; to severe illnesses preventing performance of *any* role).
- **Club Vita's lifestyle:** Socio-economic groupings developed by Club Vita (based upon ACORN classification) which is not publicly available - and will therefore not be considered in the analysis.

Further, we want our ultimate projections to be useful to individual pension schemes. We therefore need to consider whether it makes sense to use the scheme level variables such as industry or public/private sector.

These factors will capture broad differences between schemes, but arguably are less helpful to individual schemes, and less helpful for projecting improvements. When considering drivers of improvements such as access to new medical treatments, or uptake of new health behaviours, it is important we capture proxies to affluence and educational attainment at the individual level rather than at the macro level of the ‘typical employee’.

For example consider a medium sized manufacturing company. The average level of improvements in life expectancy may broadly track those of other manufacturing companies, but the liabilities of the scheme are likely to be concentrated in the executive management team. The factors which drive longevity improvements for these individuals are likely to be more closely linked to their personal lifestyle, affluence and educational attainment than the fact the particular industry within which they are carrying out their role.

For this reason we have chosen to focus on the affluence and postcode based metrics in our analysis.

This leaves us age and gender (both of which we automatically include), affluence (as measured by pension) and a range of socio-economic factors related to an individual’s postcode, as illustrated by the schematic top-right. We then need to decide:

- which of the postcode based items, marked in blue, is most important to model longevity.

  The high level of correlation between the postcode variables means that we wish to include at most one.

- whether it suffices to use a single variable (such as pension or IMD), to group people, or whether it is beneficial to use several of these variables.

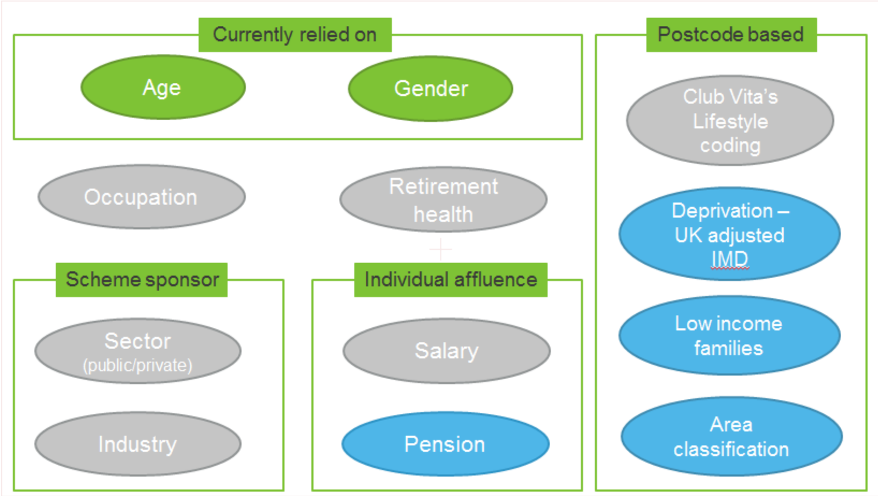

9.3 A statistical model of historical improvements - men

The univariate analysis in Section 8 revealed clear differences in men’s life expectancy improvements by pension income. Those belonging to the lowest pension band observed the highest historical improvements and those belonging to the highest band observed the lowest. However, allowing for the impact of a socio-economic measure as well might help to control for some of the noise observed when grouping data by pension.

The dataset includes three postcode based variables which can be used as a proxy for pensioners’ socio-economic status. In order to identify which of these three variables best explains differentials in mortality improvements, we carried out statistical analysis of observed historical improvements within the Club Vita dataset.

Using a framework of multivariate analysis and binomial generalised linear models we carried out a two-step process to fit to the observed improvement data:

1 Fit a baseline model as a linear function of the key mortality predictors identified by Club Vita; age, retirement health, pension amount, postcode based lifestyle[15] and IMD deprivation quintile i.e.

$$logit(q) \equiv log\left(\frac{q}{1-q}\right) = \text{Baseline(Age, Postcode, Affluence, ...)}$$

This provides a proxy to the general industry approach of using a granular model to capture the baseline for a portfolio, and was fitted to data spanning 1993-2011.

2 The resulting baseline model was then extended by adding mortality improvements to the models, conditional on the already fitted baseline parameters i.e.

$$\begin{aligned} logit(q) = \ & \text{Baseline(Age, Postcode,} \\ & \text{Affluence, ...)} \\ & + \text{ Time *Age} \\ & + \text{ Time*Imps}(\textit{Postcode variables only}) \end{aligned}$$

The aim here is to investigate the relative importance of the variables with respect to mortality improvements in a simple way, rather than coming up with the 'perfect' model for historic improvements.

Step 2 was repeated varying which postcode predictors were included in the "Imps" component above. This iterative process revealed that adjusted IMD was the most predictive postcode derived variable in relation to historical improvements. It proved to give the best balance between the simplicity and fit[16] when allowed for in the improvements.

Our final step is to check whether adding this socio-economic measure (IMD) to an 'Imps' model which already allows for pension, will improve the fit in a significant way – which it did[17].

Having identified pension and deprivation (IMD) as key variables in relation to model past improvements the next step is to segment the data further according to different combinations of these two variables (see section 10). This will allow us to explore differences in historical improvements for various groups of people with similar characteristics (in terms of affluence and postcode).

9.4 Historical improvements - women

It would be natural to translate the results for men across to women i.e. to seek to use adjusted IMD and pension amount as the main characteristics to allow for when modelling differences in improvements.

However, accrued pension is generally both a very modest and a misleading affluence measure for women; and especially amongst the generation of current pensioners, many of whom will have had fragmented careers and part time service. As such it makes less sense to use pension amount. This is supported by the analysis in Section 8 where we see no clear differences in improvements in life expectancy when we split women by pension amount.

[15] Using Club Vita's proprietary postcode based lifestyle rating factors – see Appendix A

[16] Reference AIC (before adding in any improvement variable) is 1,238,953. By adding adjusted IMD to the model the AIC reduced by 13 units and improved the reference model in a significant way. The other postcode derived variables showed less decrease when added to the model.

[17] The drop in the total deviance of the model, by adding IMD to the improvement part of it, was significant enough to pass the AIC criteria.

In contrast though, we observed clear differences in life expectancy improvements for women by deprivation (IMD). For the rest of our analysis we therefore focus on deprivation quintiles only when investigating differences in historical improvements for women. (This also has the advantage that we can combine the data on female pensioners with widows to maximise the pool of data being used to analyse improvements.)

NB: In theory we could have repeated the analysis done for men to identify the ‘best’ postcode based variable. However, using the same measure of socio-economic status for men and women seems more practical.

10 Generating our socio-economic groups

In order to create scenarios for future improvements in longevity we first need to identify groups of pension scheme members to use in those scenarios. The analysis so far suggests grouping men by some combination of pension and deprivation; whilst for women this can be simplified to deprivation.

10.1 Divide and group

Starting with men, we need to split the data into cells of individuals who are similar in pension and deprivation terms. We can then look to combine these 'cells' into groups of individuals alike in terms of longevity and improvements thereof.

In splitting the men, we need the number of cells to be sufficiently large that they are useful in helping identify which combinations of pension and deprivation can be considered alike for modelling longevity trends; yet sufficiently small that we have enough data in each cell to draw conclusions on, for example, observed changes in longevity for that cell historically.

In section 8 we saw very clear differences in life expectancy improvements for those with pensions below £2.5k p.a. to those with pensions above £7.5k p.a.. The £2.5k-£7.5k category is a relatively large group – covering over 40% of the data. We have therefore split it into two groups; £2.5k-£5k p.a. and £5k-£7.5k p.a..

For deprivation we have used the quintiles of the UK-wide IMD score, reflecting the observation in Section 8 of some differences in historical improvements between the quintiles.

This results in 20 'cells' which we can then cluster into a small number of groups. The table below illustrates the proportion of our data in each 'cell'. We can see how:

- The cells generally have around 3-9% of the data
- The cells closest to the 'diagonal' from top left to bottom right tend to have slightly more data, consistent with the correlation between affluence and living in an area of low deprivation.

		Deprivation Quintile				
		Most (5)	(4)	**Mid** (3)	(2)	**Least** (1)
Pension	<£2.5k p.a.	7%	7%	7%	6%	5%
	£2.5k-5k p.a.	6%	6%	6%	5%	4%
	£5k-£7.5k p.a.	3%	3%	3%	3%	3%
	>£7.5k p.a.	2%	4%	5%	7%	9%

10.2 Principles for creating the groups

Our next stage is to cluster our 'cells' into groups. A natural inclination would be to group those 'cells' which have seen similar levels of improvement historically. However, there are a number of other desirable features for our groups. We have sought to balance six core principles in creating our groups:

1. **Credible size:** Each resulting group needs to be sufficiently large that we can confidently use it for projections, but not so large that it dominates the projections for all schemes.
2. **Separate clear differences in improvements:** We wish to ensure that the groups capture the major differences seen in historical improvements.

3 **Group where similar improvements:** Where particular parts of the characterising population have experienced similar levels of improvement we would generally keep these together.

4 **Separate clear differences in mortality levels:** Where different groups of the characterising population have very different current levels of mortality we would wish to keep these separate as they are liable to be subject to different major causes of death and so respond differently to future longevity improvements (even if they have exhibited similar trends in the past).

5 **Interpretable:** The resulting groups should contain like individuals (i.e. similar in terms of real world features such as affluence) and thus have some interpretable and intuitive meaning. This enables the user to apply their broader understanding of the drivers of improvements in exercising judgement within the modelling of these groups.

6 **Manageable number:** The resulting number of groups should be a manageable number so as to ensure that they can be readily used by the industry.

10.3 A statistical method for applying these principles

We can apply a commonly used statistical clustering method to group the cells into clusters, designed to adhere to the principles above (and principles 2-5 in particular).

The essence of the method used is to:

1 Identify a distance metric which measures the level of dissimilarity between these cells striking a balance between the competing principles.

2 Use statistical techniques to cluster these cells into our desired number of groups.

3 Interpret the results of the clustering and consider whether it is appropriate to adjust the allocation of cells to ensure groups are both interpretable and credible in size.

Measuring 'dissimilarity'

We need a measure of 'distance' between the cells in order to apply standard statistical algorithms to group the cells which are 'closest' together. Our core principles provide us with three natural dimensions across which to measure the distance:

1 **Characteristics:** The similarity of cells in terms of the underlying variables which define the cell e.g. pension and deprivation. *(Our interpretability principle)*

We measure the distance as the weighted average of the number of rows apart in the table (as a proportion of the maximum number apart, 3) and the number of columns apart (as a proportion of maximum number apart, 4). Greater weight is given to moving between rows than between the columns commensurate with being moves of roughly 25% and 20% through the distribution of these variables respectively.

2 **Recent mortality:** The similarity of cells in terms of the levels of recently observed mortality. *(Our mortality levels principle)*

Observed mortality for the period 2008 to 2010 is calculated for each cell. Distance is measured as the absolute value of the difference between values of observed mortality, as a proportion of the maximum difference seen across all cells. For example if the observed mortality ranges from 1% to 5% and two cells have observed mortality of 4% and 4.2% then the cells are (4.2%-4%)/(5%-1%)=0.05 apart i.e. close together in the context of the range of mortality rates seen across the grid of cells.

To (broadly) control for the possibility of different age profiles in the different cells we do this for three age bands 65-74, 75-84 and 85-94 and take an average of the three distances.

3 **Mortality improvements:** The similarity or otherwise of cells in terms of observed mortality improvements (*Our principles of grouping similar improvements, but separating clear differences.)*

This is calculated using the same approach as recent mortality, but applied to smoothed annual mortality improvements for each cell. The smoothed improvements are calculated as the gradient of a linear regression fitted to the observed mortality rates by calendar year (for 1993-2011).

We then combine the distances in each of these dimensions as a simple weighted average. In general we wish to strike a balance between placing considerable weight on observed historic mortality improvements, whilst providing sufficient weight to the other dimensions to both achieve our interpretability principle and to avoid groups which are over-fitted to a specific dataset / time-period.

These considerations would tend to favour restricting the weight given to the improvements dimension to around the 50% level. The weightings we have chosen are those which both performed best in empirical testing, with 20% given to the characteristics, 30% to the recent mortality levels and 50% to the improvements dimensions.

Statistical techniques to identify groupings

We apply two clustering techniques to help identify possible groupings.

1 **Partitioning about medoids (PAM):** Under this method a single cell ('medoid') is picked to represent each of the desired number of groups. The remaining cells are grouped with whichever of these representative cells it is closest to. The distance between each cell and the representative cell it is grouped with are then totalled, providing a measure of how good the grouping is. By varying the initial choice of cells the algorithm seeks to minimise the total distance to find the best grouping.

2 **Fuzzy Analysis:** Fuzzy Analysis also seeks to minimise a (weighted) sum of the distances between the cells within each group. However, rather than allocating each cell to a cluster, it instead considers that each cell could be split between groups i.e. belong, in part, to one or more group. This provides a 'probability' of a cell being in each group and a natural grouping by allocating each cell in accordance with the highest probability.

10.4 Groups for men

For men we have chosen to split the data into 3 groups. Whilst more groups may be supported by the data, we are conscious that to use more would increase the risk that the groups will not be of credible size.

Both of the statistical approaches gave the same suggested split of the data into 3 groups, bar one cell:

		Deprivation Quintile				
		Most (5)	(4)	Mid (3)	(2)	Least (1)
Pension	<£2.5k p.a.					?
	£2.5k-5k p.a.					
	£5k-£7.5k p.a.					
	>£7.5k p.a.					

For the cell marked with a question mark there was some disagreement between the two methods, with PAM allocating to the pink group and Fuzzy Analysis opting for the green group. The underlying probabilities from the Fuzzy Analysis highlight that the allocation to pink or green is not 'clear cut'. We therefore, follow the PAM allocation which is more intuitive and avoids fragmented groups.

We then have three groups which have are easy to interpret; a group living in deprived areas who have the shortest life expectancies ('**hard-pressed**'), a group covering those individuals with the highest pensions and longest life expectancies ('**comfortable**') and a group in between ('**making do**').

		Deprivation Quintile				
		Most (5)	(4)	Mid (3)	(2)	Least (1)
Pension	<£2.5k p.a.					
	£2.5k-5k p.a.					
	£5k-£7.5k p.a.					
	>£7.5k p.a.					

It is reassuring that each group has very different life expectancies and that the comfortable group in particular has experienced noticeably different improvements in life expectancy over the last decade (charts below).

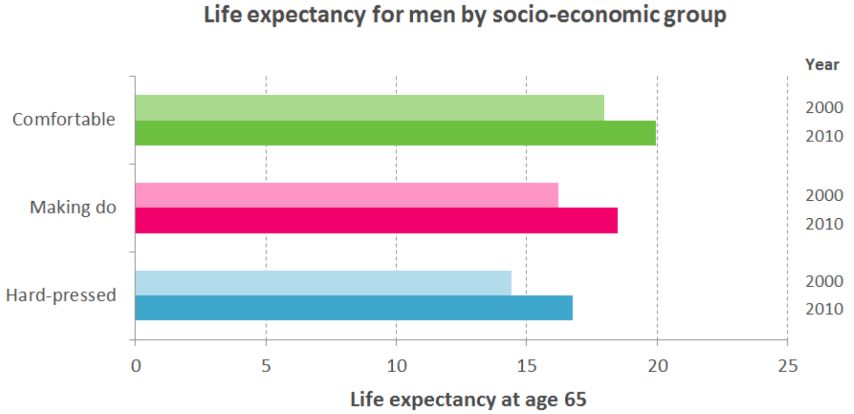

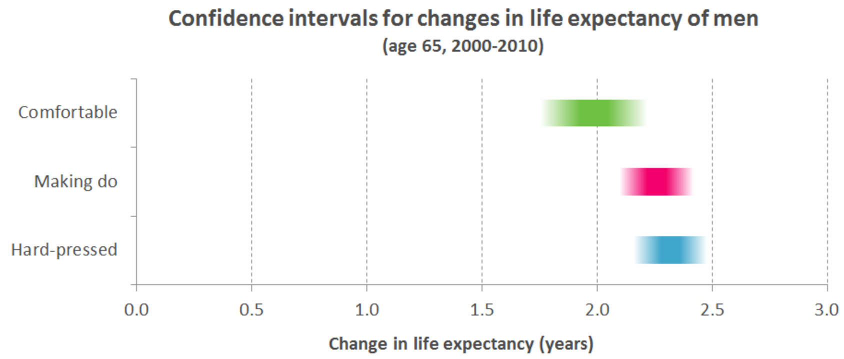

Individual schemes have considerable variation in their exposure to each of these three groups as can be seen from the chart below, with some schemes barely having any members in the 'comfortable' group and others being dominated by this group.

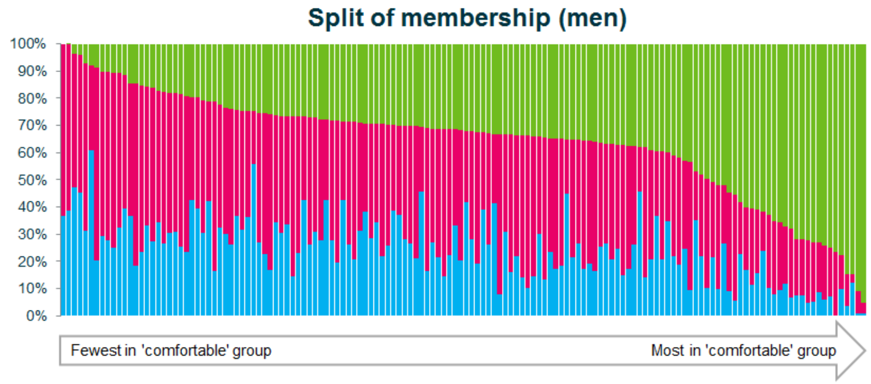

Further, drilling down to amount of pensions in payment (a proxy to liabilities) we find that the comfortable group represents under 10% of pensions in payment for some schemes, increasing to over 95% for other schemes; whilst the hard-pressed group can be over 50% of pensions in payment.

10.5 Groups for women

We decided in section 9 to split our data on women by the UK-wide deprivation measure only. Using quintiles this gives five possible cells, each with broadly similar data volumes.

Using 3 groups we would have at least one group with just one of the quintiles in it. Empirical testing verified this was liable to create a group where random noise was too great to reliably model improvements. Consequently we have opted for 2 groups, and the statistical clustering[18] suggests the following split:

Deprivation Quintile				
Most (5)	(4)	**Mid** (3)	(2)	**Least** (1)

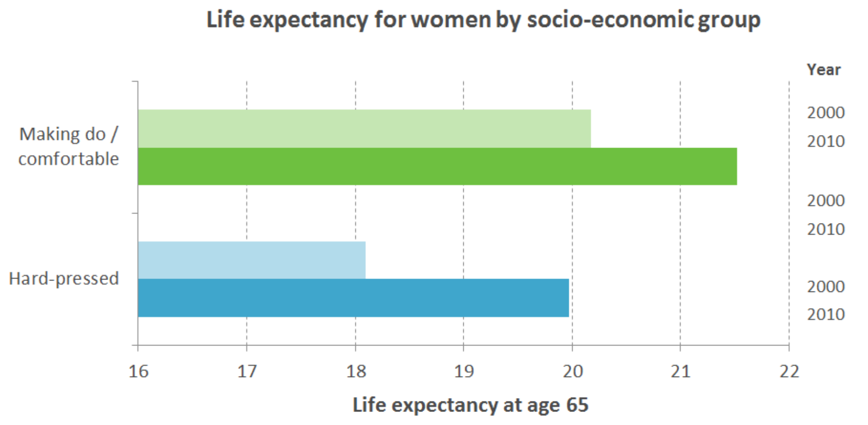

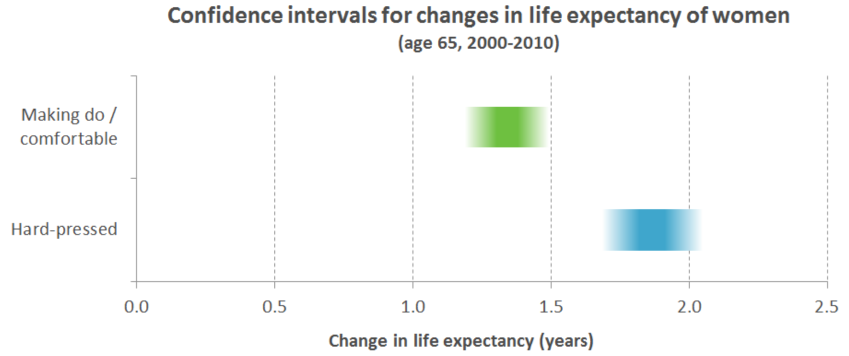

This split:

- Has a clear interpretation: areas of 'below average' deprivation (a '**hard-pressed**' group) and 'average & above' deprivation (which we refer to as the '**making do / comfortable**' group)
- Seems 'intuitive' given the results for men which – in the IMD dimension - are broadly between the two most deprived and the three least deprived quintiles
- Creates groups with a clear differences in historical improvements in life expectancy

[18] PAM method only; Standard implementation of fuzzy Analysis requires more 'cells' to fit groups.

Further different schemes have different exposures to the two different groups.

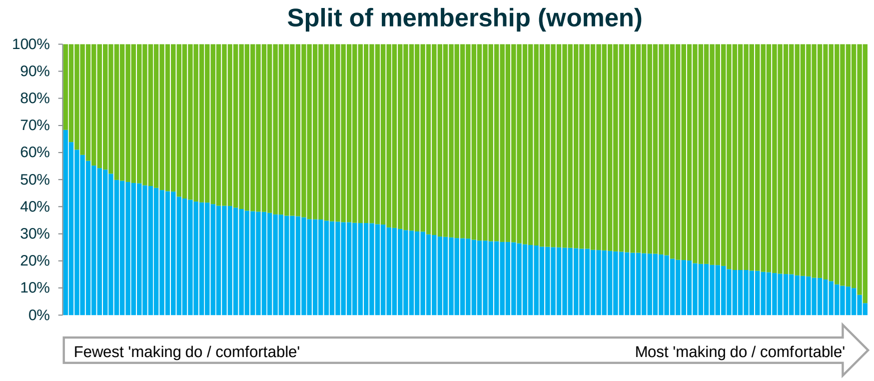

Embedding our data into the CMI model

11 Smoothing DB pensioner data

11.1 Introduction

Any mortality experience will include an element of random noise, both from year to year and from age to age. This is true even when considering population data. Therefore, we need to smooth our data to reduce the volatility from this random noise before we seek to use our data to calibrate longevity projections.

When considering the level of smoothing to apply, there is a balance between on the one hand applying no smoothing, and so having extremely volatile experience, and on the other applying excessive smoothing, and so missing out key features of the data.

11.2 Approach taken to smoothing (aggregate male data)

Data used

We carried out smoothing of the aggregate data at the outset (i.e. before subdividing the population into the clusters set out in section 9). We started with member data covering ages 60 to 95, and years 1993 to 2012, as set out in section 1. The reallocation approach discussed in section 2 was then applied to maximise the available data, and so minimise volatility.

Smoothing method

The software and method that we use to carry out the smoothing is based on, and consistent with, that provided by the Continuous Mortality Investigation (CMI), as set out in their Working Paper[19]. The model seeks to minimise the penalised log-likelihood.

We have applied smoothing in both the age/period and age/cohort dimensions for comparison, and have used knot spacings of both 4 and 5 years (and combinations thereof). These spacings have been focussed on as 6 years is likely to over-smooth, while 3 years is likely to over-fit to the data.

The goodness of fit of the different methods and knot spacings above is assessed by examining the Bayesian Information Criterion ('BIC'), where the best fit is provided by the smoothing with the lowest BIC. Age-period smoothing with knot spacings of 5 years in both age and period dimensions gives the lowest BIC.

11.3 Generating heat maps (aggregate)

We can use heat maps to graphically illustrate annual improvements in mortality rates. These heat maps can be a useful tool to highlight trends and patterns in annual improvement rates.

When generating heat maps, 'warm' colours (e.g. red/orange) show strong improvements (i.e. high annual mortality improvements), whereas 'cool' colours (e.g. green/blue) show lower improvements.

The heat maps on the following page show the results (for men) of a number of different smoothing methods, including that described above:

- On the left, crude smoothing (using 5 year averaging across age and calendar year).

  So for age x and year y, we have the following smoothed crude mortality rate:

  $$q_{x,y} = \frac{\sum_{i=x-2}^{x+2} \sum_{j=y-2}^{y+2} e_{i,j}}{\sum_{i=x-2}^{x+2} \sum_{j=y-2}^{y+2} d_{i,j}}$$

  The heat map shows the year-by-year improvements in these smoothed rates for each age.

[19] See CMI Working paper 20

- In the middle, smoothing using the age-period method, as set out above. We can see that the smoothed improvements exhibit similar features to the crude smoothing:
  - Warmer colours at younger ages, with improvements cooling for older ages;
  - Evidence of cohorts, shown by diagonal patterns as particular years of birth age over time; and
  - Some limited evidence of a ‘golden generation’ born around the late 1940s, so currently in their late 60s, although this effect seems less pronounced in recent years.

It is reassuring that the smoothing method adopted has not overly smoothed the underlying features.

The corresponding heat map for ONS data, as adopted by the CMI in CMI2013, is also shown below, on the right. We can see how there are similarities in the heat map to the data used in our analysis, albeit that there are a number of distinct cohorts showing here, and at different years of birth.

Crude smoothing (men)

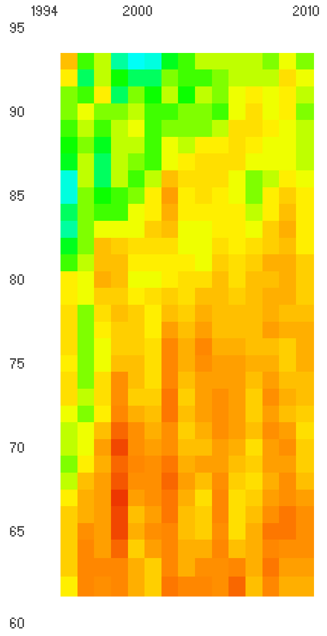

Age-period (5x5) smoothing (men)

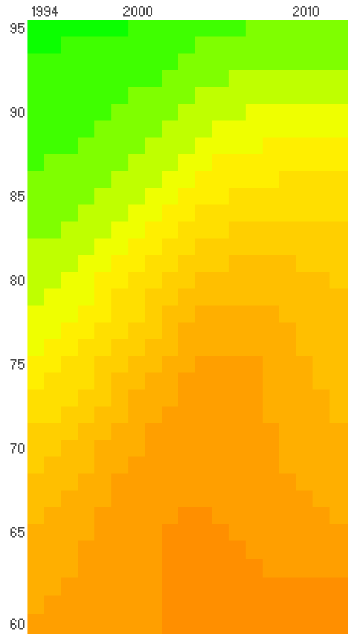

ONS data (men)

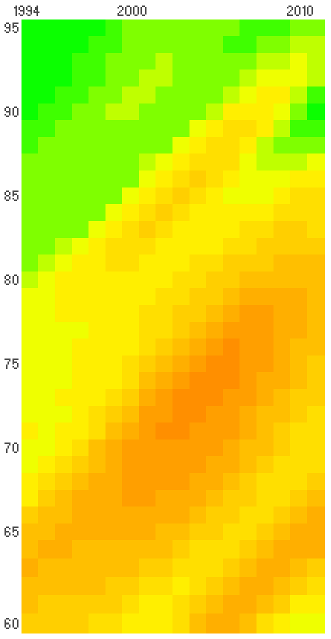

Legend

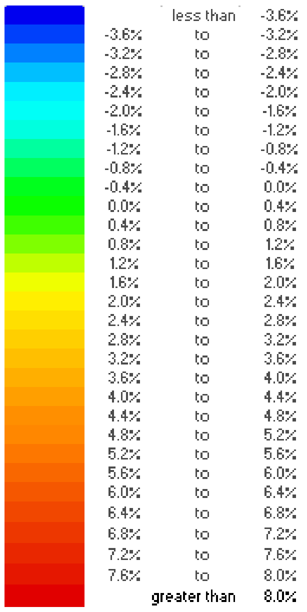

11.4 Disaggregating improvements

The CMI model splits the initial mortality improvements for the calendar year of projection for the model into age/period and cohort components (see section 5).

We have used a disaggregation model which is consistent with the approach adopted by the CMI when populating the core settings of the CMI model. This model solves for the age, period and cohort components, minimising the sum of the squares of the residuals.

We have applied this disaggregation model to the smoothed annual improvements above (i.e. improvements that have been smoothed through the age-period method with 5 year knot spacings). The charts below show the results of this disaggregation (for calendar year 2010).

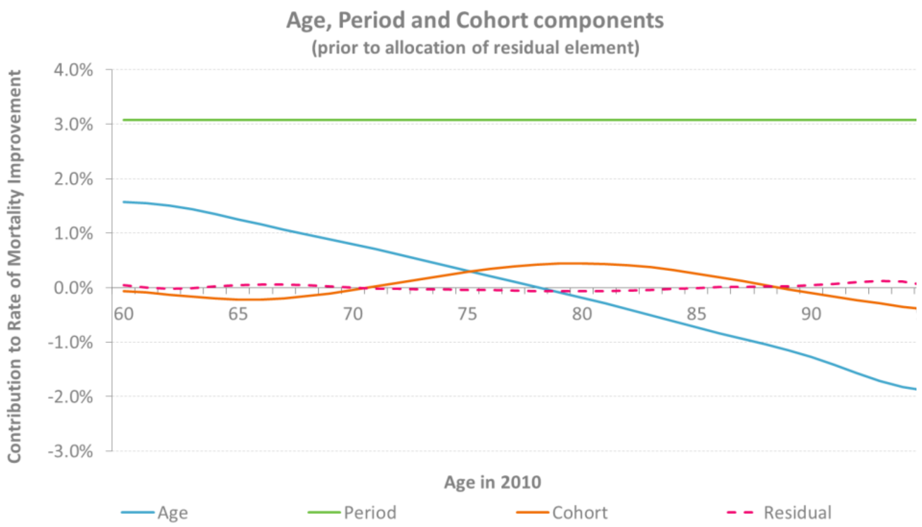

Comparison with CMI_2013

We have shown below the equivalent disaggregation used by the CMI. Note that this is based on England & Wales population data, from age 18 to 102 and years 1961 to 2012, although we have only shown the range from 60 to 95 for ease of comparison.

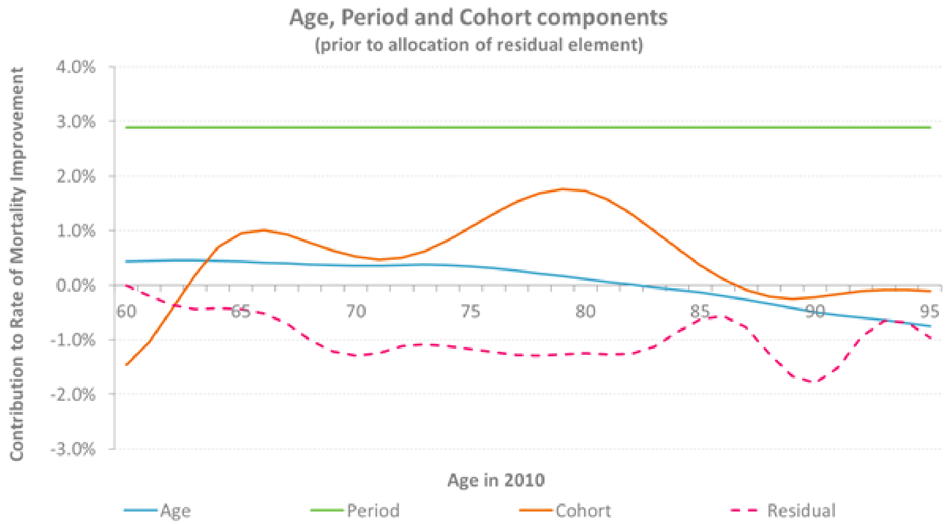

In comparing the disaggregation with that used by the CMI, we can see that, for the ages covered by the DB pensioner population:

- The age component has a different pattern at younger ages (60-75) where it falls with age in the DB pensioner population, however after age 75 it decreases for both;
- The period component is higher in the DB pensioner population in 2010;
- The cohort component has a similar shape, although the peak is less pronounced in the DB pensioner population, and is much less significant below age 70; and

- The difference between the fitted improvements (age *plus* period *plus* cohort components) and the smoothed improvements being used by the disaggregation model are considerably lower for all ages within the DB pensioner population.

For the purposes of projection the age and period components are combined, whilst (for these ages) the residuals are added to the cohort component. This has the impact of creating a more complex cohort pattern within the CMI model. In contrast – as can be seen from the following – the cohort component is much smoother for the DB pensioner population, omitting some of the known 'false' cohorts arising from approximations within the ONS data[20], and avoids the very 'wavy' line from the CMI model which is symptomatic of over-fitting at the initial p-spline smoothing stage.

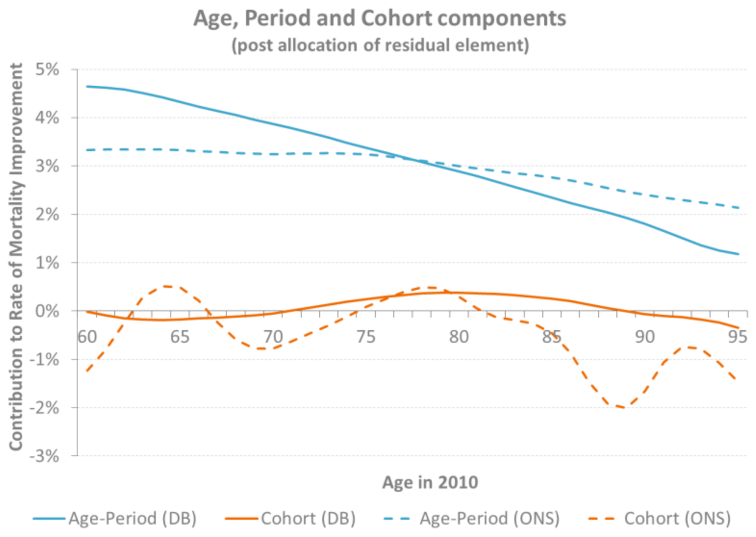

[20] For further details of the approximations adopted by the ONS, and the issues these can cause, see our blog on "The mystery of the vanishing nonagenarians".

There are a number of reasons for the differences between ONS E&W population and DB pensioner population seen in these charts:

- There is an inherent selection effect from using pension data rather than general population data, in that pension data will, for example, exclude those who have never been fit for work.
- The DB pensioner data has a more restricted age range (60 to 95) than the population data (18 to 102) – as a result of lower levels of exposures at younger ages restricting the credible data.
- The time period using is also more restricted (population data goes back to 1961 while the DB pensioner data is only back to 1993) – as a result of both the lack of credible historical data held by pension schemes and scheme maturity.

Extending the cohort component to younger ages

The Club Vita pensioner data for men used for this smoothing and disaggregation starts at age 60 i.e. year of birth 1950. In order to be able to use for projections for younger generations i.e. the post-retirement mortality of the deferred pensioners and active members of pension schemes, it is necessary to extend the cohort component to younger ages.

On grounds of materiality we have applied a pragmatic approach of blending the solved cohort component for our DB pensioner data into the cohort component resulting from the England & Wales population data over 7 years from 1950 to 1957, as shown in the chart below. The initial rates of improvement below age 60 are assumed to be in line with the England & Wales population data, so the age/period component is the resultant difference between this initial rate and the blended cohort component.

11.5 Extending mortality rates to older ages

We also need to extend the mortality rates above age 95 up to age 120, as we only have smoothed improvements up to age 95. Again we have adopted a pragmatic approach, as follows:

- Assume that the force of mortality at age x, μ_x, is equal to 1 at age 120.
- Apply extrapolation of the smoothed μ_x (which covers age range 60 to 95) above age 95 up to 120.
- Converted the extended μ_x into q_x and so annual improvements can be found.

This enables the construction of a complete mortality improvement table up to age 120 for the aggregate male population, and is consistent with the way the CMI extends the Self Administered Pension Schemes (SAPS) dataset above the modelled range.

12 Are DB pensioners different to the general population?

12.1 High level summary

Before starting this project, we had an expectation that defined benefit (DB) pensioners would be slightly longer-lived than the England & Wales (E&W) population as a whole, and that they may have experienced a different level of improvement in life expectancy to the wider population.

This is because DB pensioners represent a specific subgroup of the population – they have been in formal employment for at least a proportion of their adult lives which means that they might be expected to be:

- less likely than the wider population to be severely disabled or have a life-limiting condition;
- wealthier on average than the wider population.

Also, previous analysis by Club Vita and the Self-Administered Pension Schemes survey by the Continuous Mortality Investigation (CMI) both suggest that life expectancy is slightly higher for pension scheme members than for the population as a whole.

Surprisingly, the data gathered for this project suggests that life expectancy from age 65 for the subset of the DB population analysed is actually very similar to that in the E&W population as a whole over the period under consideration (1993-2010).

	Men		Women	
Year	Population	NAPF	Population	NAPF
1993	14.2	14.3	17.8	17.7
2000	15.7	15.6	18.9	18.8
2005	16.9	16.9	19.6	19.7
2010	18.1	18.1	20.6	20.6

(UK population data is from the Office for National Statistics)

12.2 What does this mean for projections?

With very similar increases in life expectancy between 1993 and 2010 for the two sets of populations, it might be expected that projecting life expectancy forwards would give roughly the same result regardless of which population was chosen as the starting point.

However, the surface similarity obscures some significant differences that can be found by digging deeper into the data. Because life expectancies are produced by summing the likelihood of survival to each future age, it is possible that two very different patterns of mortality improvement at different ages could produce the same overall improvement in life expectancy.

For example, very strong improvements in mortality rates at the oldest ages in one group could offset weaker improvements at younger ages, compared to a second group experiencing moderate improvements at all ages.

The following charts compare the average annual improvement rate at each age between 65 and 95 across 2000 to 2010, for the NAPF data set and the population data set as a whole (population mortality rates are taken from the CMI model (CMI_2013)).

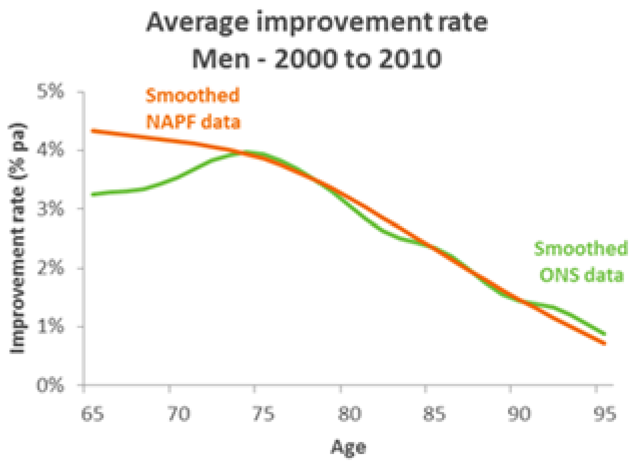

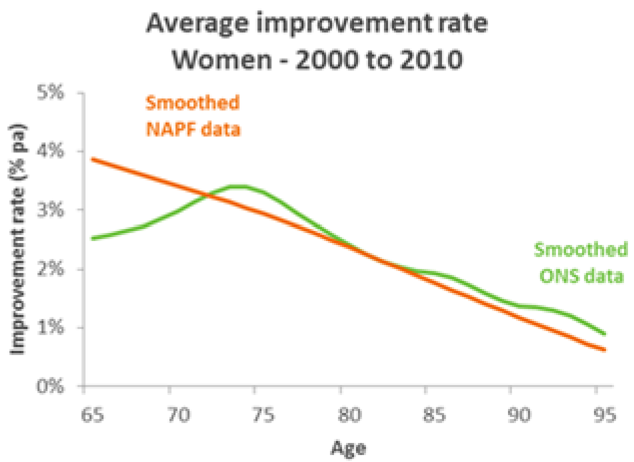

We can see that, for both men and women, the average improvement rate is very similar at ages above 75, with NAPF data typically having slightly higher improvements at younger ages than ONS data.

It is also clear that the average improvement rate typically reduces with age; this is consistent with widely-made observations on ‘rectangularisation’ of mortality at the oldest ages (that is, more people are living to the oldest ages, but mortality rates at these ages are improving less quickly), as well as the assumption within the CMI’s Core Parameters which allows for improvement rates to move towards zero after age 90.

There are various possible reasons for higher rates at younger ages – it may be due to the ‘select effect’ of DB pensioners noted above; alternatively, it could be due to the shape of the data set. The data used in the CMI model covers a wider range of ages and years than the NAPF dataset and hence the results at age 65 and above in the CMI model depend partly on results for younger people, who are not present (or arguably relevant) to the DB pensioner dataset.

Re-smoothing the ONS data but using the same range of ages and years to the NAPF dataset gives a set of improvement rates which are much closer to the DB pensioner result, implying that the range of ages and years chosen has an impact on the result that can be as material as changing the dataset modelled.

However, the projections are not directly based on the average improvement rates over time; the starting point is actually the 2010 improvement rates. Looking at these rates in isolation for the two populations, we produce the following charts:

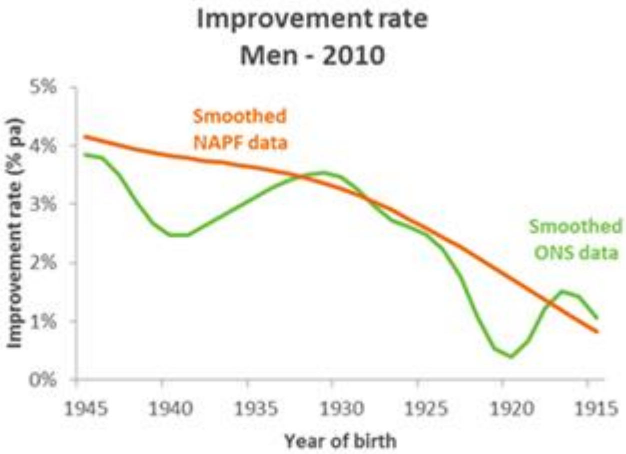

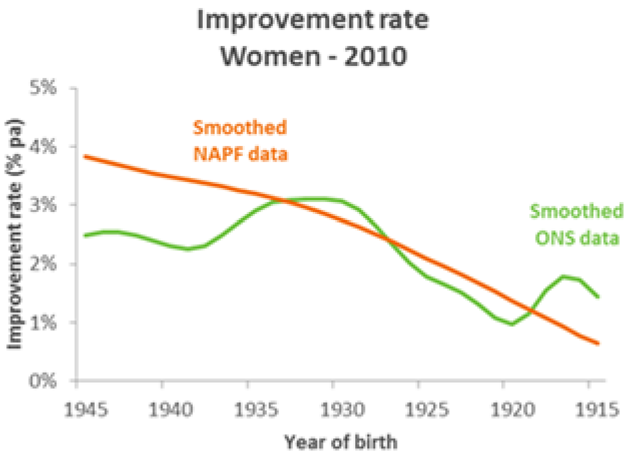

We can see that the ONS data, even when smoothed, is more variable than the NAPF data. Because of the larger size of the ONS dataset, the modelling process is more likely than the process used for the DB pensioners to pick up cohort effects, whether real (one year of birth has experience significantly different to others) or spurious (e.g. due to overfitting or data issues).

It is also clear that, particularly for males, there are two cohorts where ONS improvements are significantly lower than NAPF improvements. These relate to:

- individuals who are born around 1920 – research[21] shows that this is at least in part due to inconsistent patterns of birth in 1919 and 1920 as a result of World War I and Spanish flu.
- individuals who are born either side of 1940 – it is slightly less clear where this effect comes from, but it could be that this is again a result of patterns of birth in the lead-up to and aftermath of World War II

Because we have access to individuals' dates of birth (which is not available from ONS data), we do not see the inconsistencies due to birth patterns in our analysis.

Because the 1940 cohort in particular can be expected to form a material proportion of the overall membership of a typical pension scheme, substituting NAPF data for ONS data in the CMI model (without making other changes) would be expected to lead to a higher value being placed on liabilities.

[21] see "Phantoms Never Die: Living with Unreliable Mortality Data" (Cairns et al, 2014) *www.macs.hw.ac.uk/~andrewc/papers/ajgc71.pdf*

13 Smoothing historical improvements for each group

Having carried out the smoothing processes at the aggregate level (section 11) we then need to consider the smoothing approach to adopt for each of our socio-economic groups (hard-pressed, making do and comfortable).

13.1 Inheriting the cohort effect

We have carried out research[22] which suggests that cohort effects seen in large populations tend to persist in subgroups of the population. We have adopted this approach in our analysis, effectively assuming that the cohort component seen in the overall aggregate population is carried through to each of the subgroups.

It is worth noting however that there is a cascade effect in population mortality rates, where improvements in life expectancy as a result of lifestyle changes such as ceasing smoking typically initially occur in higher socio-economic groups, as they are generally better educated etc, before working down to lower socio-economic groups as benefits become more generally known. There may be some cohort component to this cascade effect, which we have not allowed for in our analysis. We hope to investigate this further in the future.

[22] To be published in December 2014

13.2 Heat maps (cohorts)

We can generate heatmaps of the annual mortality improvements in respect of each individual cluster, following the same p-spline smoothing process as used for the aggregate DB pensioner data. These heat maps are set out below, along with the heat map for the all DB pensioner men for comparison.

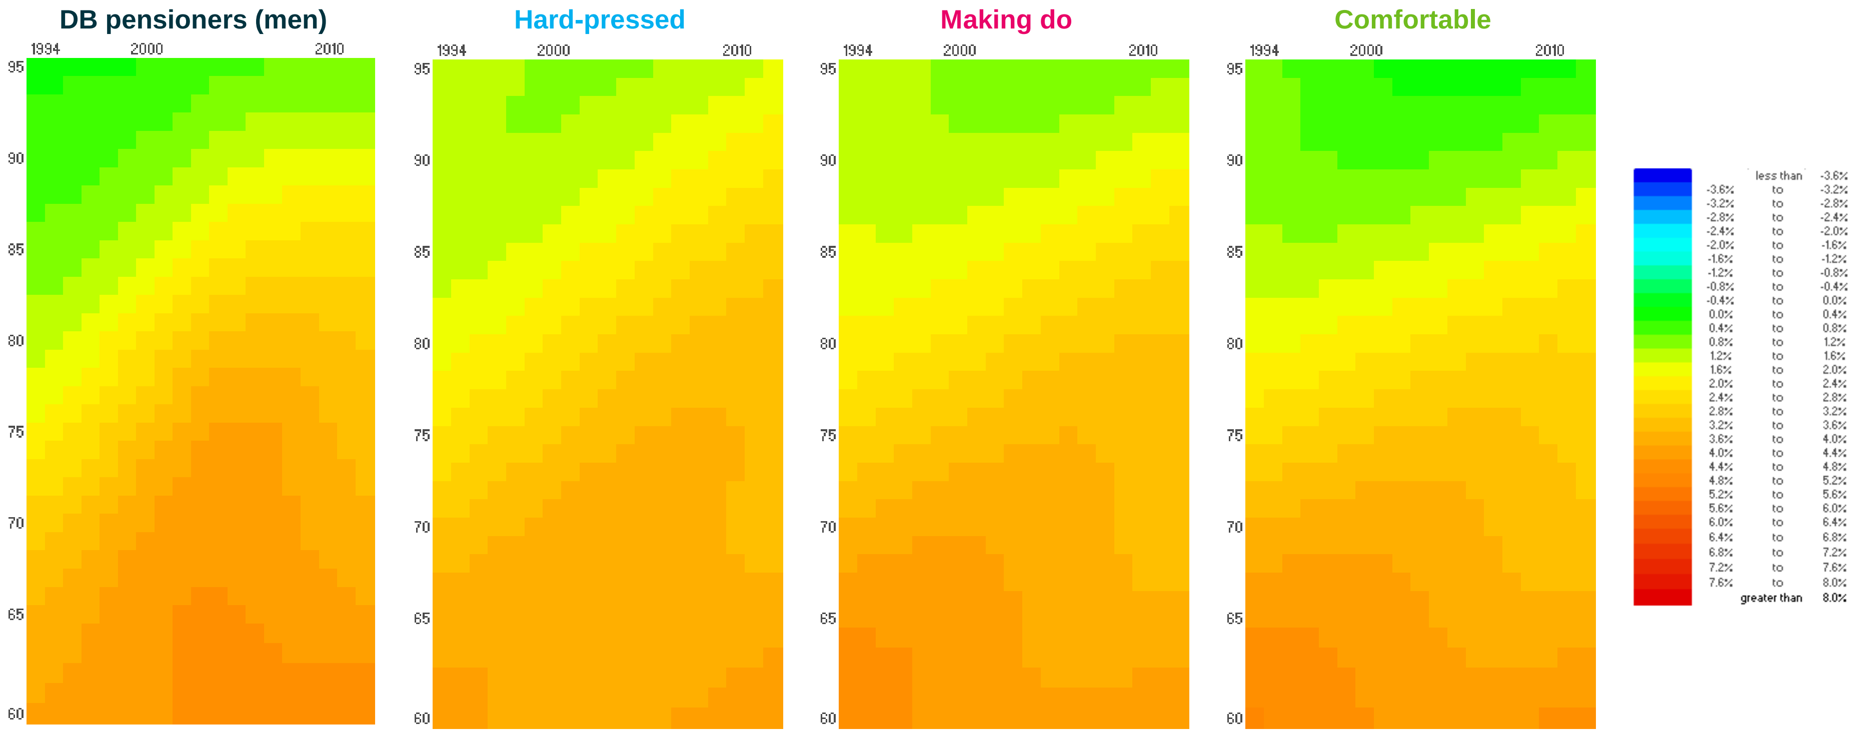

We can see that the heat maps are reasonably similar across each of the clusters, with evidence of diagonal cohorts, and declining improvements with age.

There are some slightly different patterns for the different groups – the strong reds appear at different ages/years, and the comfortable group seem to be cooling more than the other groups.

13.3 Smoothing for women

The analysis above has focussed on the treatment of men. An identical approach to smoothing, including assuming that the cohort component observed at the level of all women is common to the sub-groups, was also adopted in respect of women (although over the age range from 65 to 95, rather than 60 to 95 as used for men).

The resultant heat maps in respect of the two subgroups, along with the heat map in respect of the aggregate data, as well as the corresponding heat map for ONS data, as adopted by the CMI in CMI2013, are shown below:

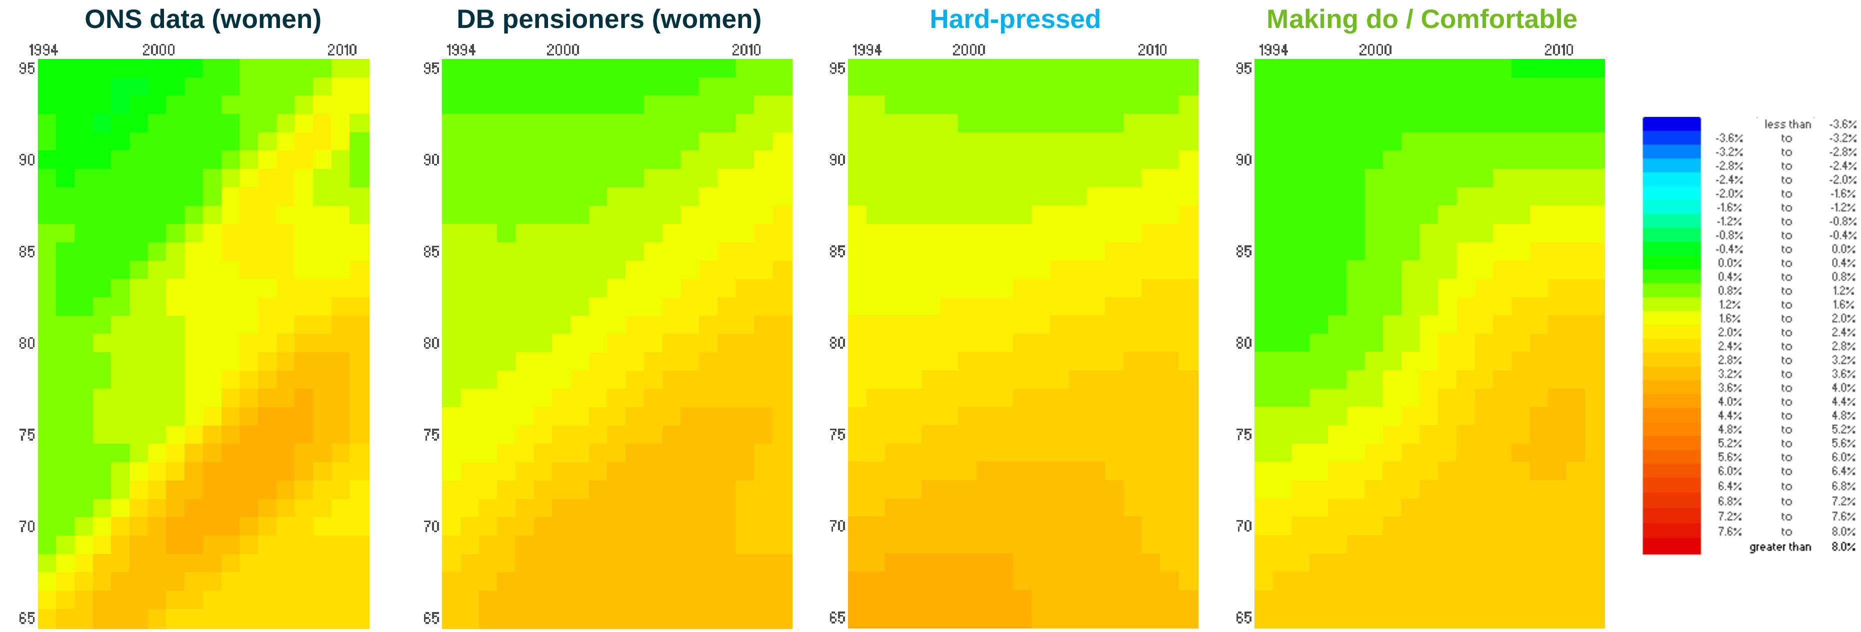

Again the aggregate heat map is similar to that used by the ONS, although the ONS data shows some evidence of an older cohort with high improvements, currently in their early 90s.

The heat maps for the subgroups show some notable differences from the aggregate level, with cohorts of varying degrees in different places.

13.4 Historic improvements in life expectancy

Having smoothed the improvements for each subgroup, for men and women, as set out above, we can then examine the trends in period life expectancy over time for each subgroup. The charts below set out how period life expectancies have evolved over the decade to 2010, for both the subgroups and the overall data set.

13.4.1 Improvements for men

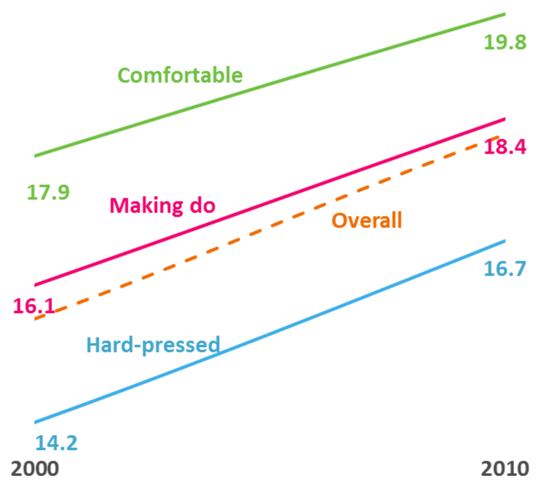

We can see from the above that:

- The increase in life expectancy is most significant for the 'Hard-pressed' group, both in number of years and as a proportion of starting life expectancy.
- While the 'Hard-pressed' and 'Making do' groups have seen life expectancy increase by 2.5 and 2.3 years respectively, these increases represent different relative increases in life expectancy.
- The 'Comfortable' group began the period under review with life expectancy of 17.9 years at age 65 and has seen life expectancy increase by around 1.9 years over the decade; a slower rate of increase than other pensioners.
- The result of this is a narrowing of the life expectancy gap between the 'Comfortable' and 'Hard-pressed' groups of just over half a year over the period.

13.4.2 Improvements for women

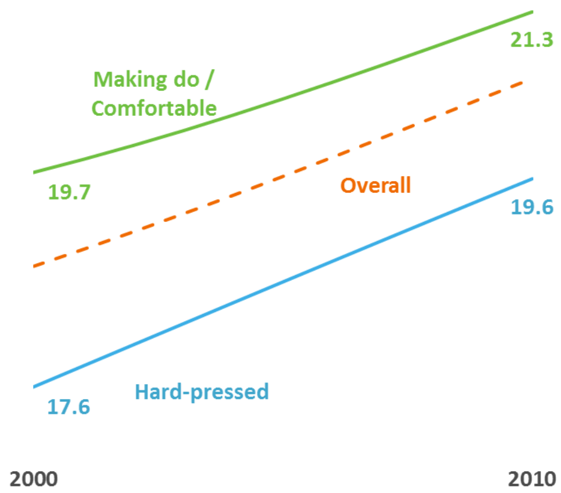

We can see from the above that:

- The 'Hard-pressed' group has seen life expectancy increase by 2 years from 17.6 to 19.6 years at age 65. This represents a more significant proportion of starting life expectancy than the more affluent group.

- The ‘Making do / Comfortable’ group have seen life expectancy increase by around 1.6 years, from 19.7 years at age 65, over the decade
- The results once again highlight a narrowing of the life expectancy gap between the groups. For women the gap narrowed by just over a third of a year.

13.4.3 Other published sources

A number of other studies have been carried out into the differing rates of improvement in mortality rates for different socio-economic groups.

For example, a discussion paper was presented in September 2012[23] which concluded that, based on population data for England over the period from 1981 to 2007, the least deprived IMD quintiles saw the greatest reductions in ***mortality*** of this period with the difference most notable in the most recent years. However, greater percentage reductions in mortality amongst those with lower mortality to start with are needed in order to have the same increase in life expectancy. As such these results need not be contradictory. Further, within DB pension schemes we are looking at a select group of lives. Those who have been fit and active enough to earn a DB pension may be less representative of the general population in areas of high deprivation (by definition associated with higher levels of unemployment) than the population in areas of low deprivation.

In contrast analysis by the Office for National Statistics[24] shows that, for men, whilst the gap in life expectancy at age 65 had *widened* between those in most and least affluent socio-economic circumstances (as proxied by the occupation-based NS-SEC measure) between the mid 1980s and the mid 2000s, it had ***narrowed*** between the late 1990s and mid 2000s, consistent with our findings here.

[23] See Lu JLC, Wong W and Bajekal M Mortality improvements by socio-economic circumstances, presented at a sessional research meeting on 24 September 2012

[24] See ONS Statistical Bulletin Trends in Life Expectancy by the National Statistics Socio-economic Classification 1982-2006

Interpreting the results

Club Vita LLP

14 Our example schemes

Within the NAPF's report we have illustrated the financial impact of changing the longevity improvement assumption for four different schemes. These schemes have been constructed by considering the range of schemes seen within the Club Vita database and selecting schemes which:

- Reflect the spectrum of splits between our different socio-economic groups (hard-pressed, making do and comfortable)

  Specifically we have picked a scheme with one of the larger biases towards the hard-pressed group, and a scheme with one of the larger biases towards the comfortable group. Our other two schemes are more 'central' in terms of their concentrations in each of our socio-economic groups.

- Have differing maturity profiles.

For each scheme we have smoothed the age profile of accrued pensions, applied a multiplier to all pension amounts and removed any other features which could identify the scheme. In each case we then assume the split of pensions in payment between our different socio-economic groups applies to the pensions in payment at each age, and to the accrued pensions for non-pensioners. Doing so creates four example schemes genuinely reflective of the industry, whilst respecting the confidentiality of individual scheme information.

The rest of this section provides details of the profile of each of these four example schemes, in terms of pension amount and socio-economic mix. We also include a further illustrative scheme based upon our knowledge of pension schemes. This scheme is used when assessing the impact of specific scenarios on each of the socio-economic groups within the scenario pages of the NAPF longevity model report.

Example A: Mature, skewed towards lower socio-economic groups

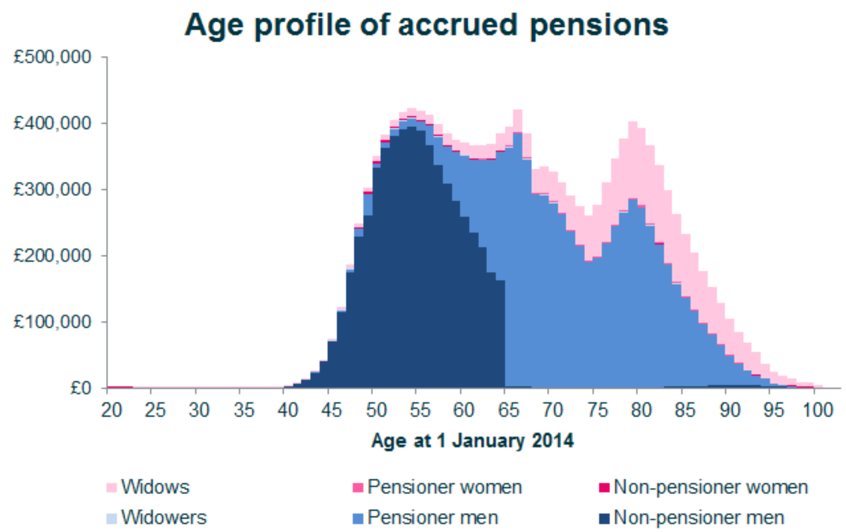

Split of pensions for men

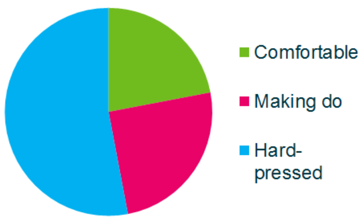

Split of pensions for women

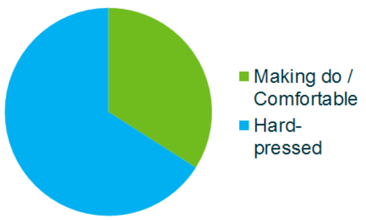

- Mature and closed to new-entrants
- Skewed towards lower-socio economic groups
- Likely to be similar to schemes from heavy manufacturing industries

Club Vita LLP

Example B: Open scheme, reasonably typical mix

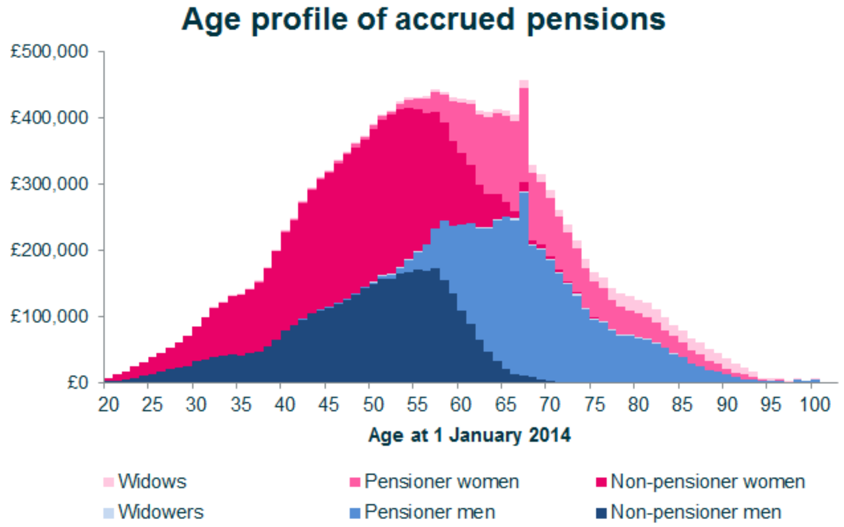

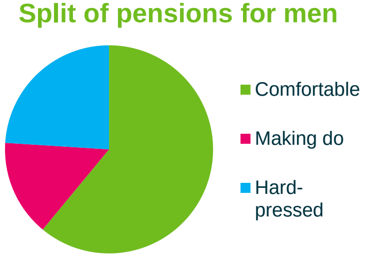

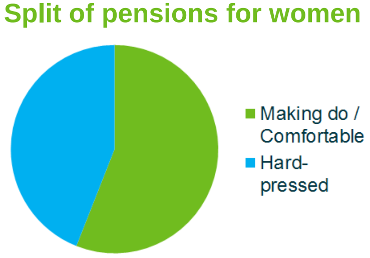

- Long standing scheme, open to new-entrants
- Broad mix of socio-economic groups
- Likely to be similar to schemes from consumer services or cyclicals and also local government schemes

Example C: Closed scheme, reasonably typical mix

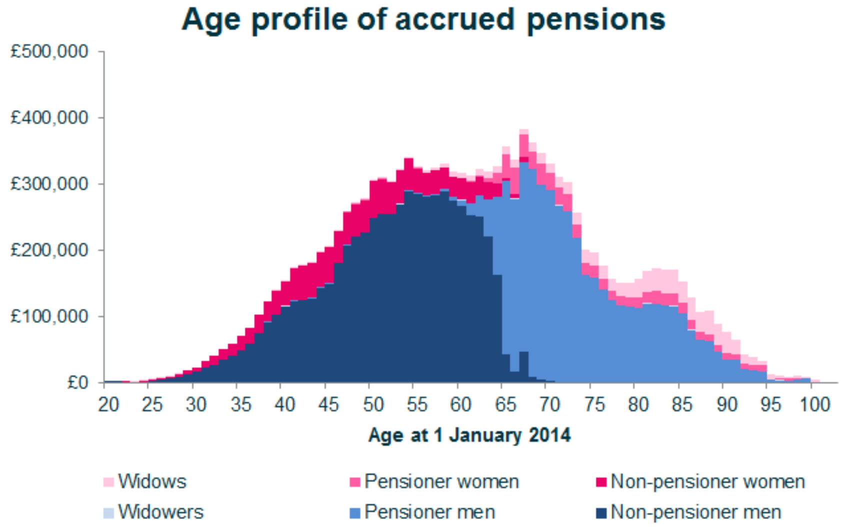

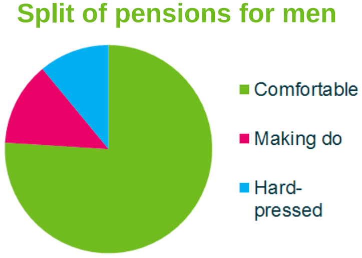

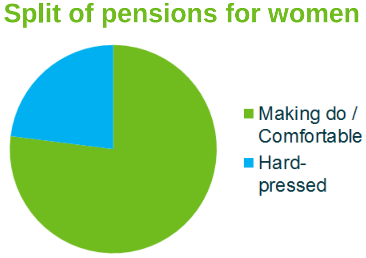

- Long standing scheme, closed to new-entrants
- Mix of socio-economic groups, although biased towards higher socio-economic groups
- Likely to be similar to schemes from technology, pharma and skilled engineering industries

Example D: Scheme skewed towards higher socio-economic groups

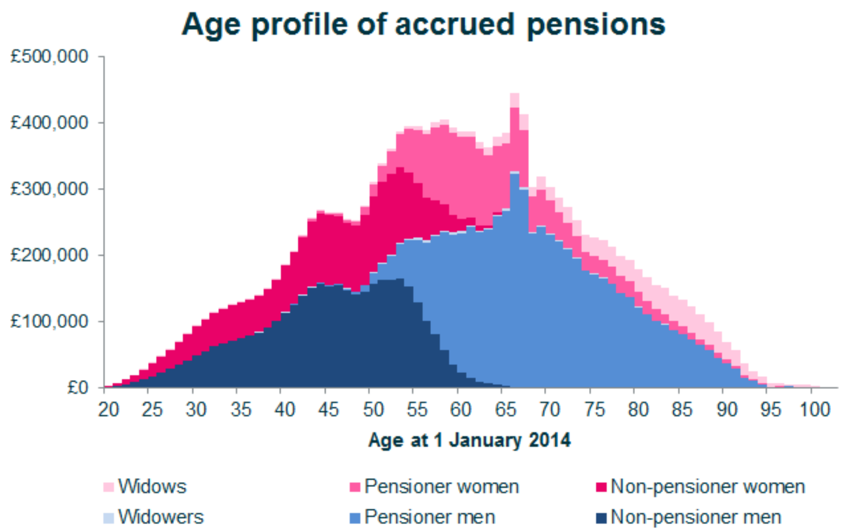

Split of pensions for men

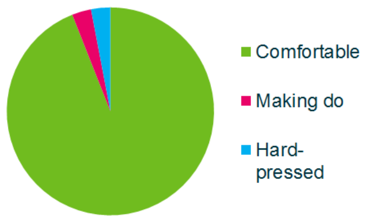

Split of pensions for women

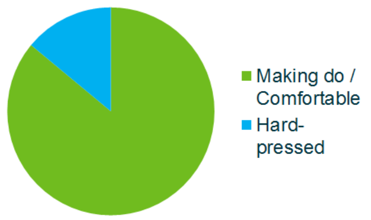

- Long standing scheme, open to new-entrants
- Propensity to allow early retirements
- Skewed towards higher-socio economic groups
- Likely to be similar profile to schemes from financial services sector

Illustrative scheme

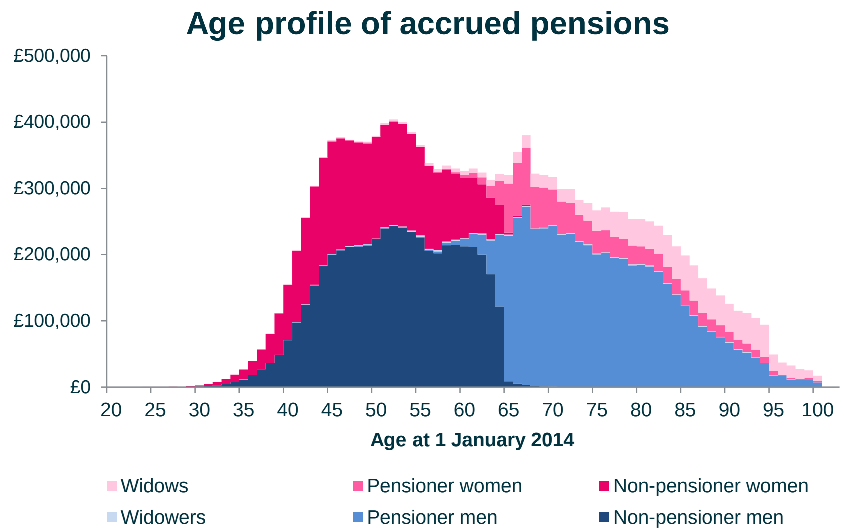

- Designed to be broadly 'average' age profile for use in assessing the impact of scenarios for different socio-economic groups in our individual scenario pages

15 Assessing the impact on pension scheme liabilities

Within the main report we have assessed the impact that adopting an improvement assumption specific to different socio-economic groups of pension scheme members would have on liability assessments. In this section we outline our approach to those assessments.

15.1 A simple valuation

Our approach is to carry out a 'simple' valuation of an illustrative pension scheme. We use age-grouped data of total amounts of pension in payment (for pensioners and in payment dependents) and accrued pensions (for active members and deferred pensioners).

The liabilities are valued assuming pensions are payable continuously[25] and, where we include spouse's benefits (see below), we assume that there is an attaching 50% spouse's pension based upon the level of the member's pension prior to any commutation for cash. For simplicity we have assumed that all pensions get the same level of pension increase.

For non-pensioners we ignore death in service / death in deferment and assume the member will survive to retirement. An allowance is made for the member to commute some of his/her pension at retirement. For the purposes of these calculations we assume a single commutation factor independent of the scenario used for future improvements. This reduces the sensitivity of the liabilities to the longevity improvement scenarios.

[25] So for example a single life annuity is calculated by the recursive formula (using the usual actuarial notation) of: $\bar{a}_x = p_x(v^{1/2} + \bar{a}_{x+1}) + q_x v^{1/4}$

Contingent life annuities are calculated as the difference between a single life annuity for the spouse, less a joint life annuity for the member and his/her spouse i.e. $\mathbf{max}(\bar{a}_y - \bar{a}_{xy})$ where $\bar{a}_{xy}$ is calculated in the same manner as single life annuities but using the joint mortality of the two individuals (member and spouse).

Longevity improvement scenarios

For each scenario we have assessed an impact on liabilities for each kind of individual. In order to do this we have focussed on the men or women in our illustrative membership profile and assumed that they are entirely one kind of individual e.g. men in our 'comfortable' group. We have also assessed the impact on a single life basis i.e. assumed the scheme pays annuities to those individuals without any payments to a contingent spouse. This avoids any complications arising from the different groupings used for the scenarios for men and women.

Impact for example schemes

When considering the financial impact for our example schemes, we have taken into account spouse's pensions payable upon death after retirement. To do this we have needed to model the mortality of members and their spouses. For men we know that:

- The ''hard-pressed' group reside in areas amongst the two most deprived quintiles. It is reasonable to assume that their spouses / financial dependents also live in these quintiles and so are in the 'hard-pressed' category of women.
- The 'making do' group reside in areas amongst the three less deprived quintiles. We therefore assume their spouses / financial are in the 'making do / comfortable' category of women.
- The 'comfortable' group include people who live in areas of reasonably high deprivation (Q4), but the vast majority live in the top three quintiles (Q1-Q3). As a pragmatic simplification we have therefore treated their wives as being in the 'making do / comfortable' group.

For women the benefits payable upon death to financial dependents are less material (as wives are more likely to outlive their husbands). Therefore we take a pragmatic, simplified approach. Specifically:

- We know that women in the 'hard-pressed' category will have husbands also living in Q4 or Q5 and so we treat all such women as having husbands in the 'hard-pressed' group for men
- Women in the 'making do / comfortable' category will be a mix of those with husbands in the 'making do' and the 'comfortable' categories. We assume that their husbands are broadly evenly split between the 'making do' and the 'comfortable' group.

Schematically this means:

Scheme man	Wife
Hard-pressed	Hard-pressed
Making do	Making do / comfortable
Comfortable	Making do / comfortable

Scheme woman	Husband
Hard-pressed	Hard-pressed
Making do / comfortable	Making do and Comfortable split 50:50

15.2 Longevity assumptions

For each group (hard-pressed, making do and comfortable) we need to make an assumption about current levels of longevity (baseline) and how this will change in future (future improvements)

Baseline longevity

For each group we have smoothed historic mortality experience for each calendar year from 1993 to 2012 using the p-spline smoothing software produced by the CMI. We have then used the 2012 calendar year mortality as the starting point for projection. To extend these mortality rates to the older ages (95+) outside of the range of available data we have applied the same approach as the CMI SAPS committee have used in S1 and S2 series of mortality tables, i.e. a smooth progression from the fitted force of mortality at age 95 to a force of mortality of 1 at age 120 (and a gradient of 1 at 120)[26].

Future improvements

In each scenario we apply the appropriate set of improvements (from 2012 onward) to these base rates to calculate the projected mortality rates for each group.

15.3 Other demographic assumptions

In order to assess the liabilities we also need to make some additional demographic assumptions. Specifically:

- Non-pensioners retire age 65 (or immediately if older)
- No separate cash lump sum payable upon retirement by right
- 20% of pension commuted for cash at retirement at 18:1
- 90% of men are married (at retirement or current age for pensioners)
- 80% of women are married (at retirement or current age for pensioners)
- Husbands are 3 years older than wives

15.4 Financial assumptions

- Net discount rate prior to retirement = 2%
- Net discount rate post retirement = ½%

15.5 Other assumptions

- Effective date of valuation of 1 January 2014

[26] For more details of these calculations see Section 11

15.6 Value of liabilities

We then apply the above assumptions to calculate the value of liabilities for the example pension scheme under each scenario. By taking a ratio of:

- the liabilities under a particular scenario; to
- the liabilities under the reference scenario (i.e. longevity improvements in line with CMI 2013 with a 1.5% long term rate)

:we obtain an indication of the impact of allowing for the particular scenario on each group.

15.7 A final word of caution

The actual impact of each of the longevity scenarios on a particular scheme will depend on a number of factors, including, but not limited to:

- Financial assumptions used

  In general, the lower the net discount rate, the more sensitive liabilities are to a change in longevity. Whether the scenario increases or decreases liabilities will not change, but for lower (higher) net discount rates the magnitude of the change will be bigger (smaller).

- Scheme benefit structures (e.g. pension increase tranches, cash entitlements, early/late retirement, etc)
- Age structure

  For example, younger populations are more sensitive to changes in the view on long term rates of improvement than older populations.

16 Creating scenarios

Our analysis has shown that:

- Different groups have had different longevity improvement trajectories over the recent past
- Reflecting this trajectory in the improvement rates assumed for different groups changes the overall liability

The majority of pension schemes use a limited range of improvement assumptions, typically the same assumption for all members of a given gender. As such the obvious next step is to consider the impact on the liabilities for a range of schemes of:

- overall future development of longevity improvements being different to that expected, and
- different improvement trajectories for each of the proposed groups

Decisions taken up to this stage relate to the past and are broadly objective, for example the allocation of individuals to the different groups, and the calculation of historic longevity improvements.

Conversely, the trajectory for future improvements is subjective as everyone will have a different view as to how life expectancy will change in the future.

This section sets out six diverse scenarios, to illustrate the range of outcomes that could be considered. These scenarios produce liability results in broadly the range -15% to +15% compared to the status quo, which we assume to be the CMI model using Core Parameters and 1.5% pa long-term rate (as set out in section 5, this is the assumption used by around 50% of schemes with recent valuation dates).

We have also ensured that some of the scenarios have similar impact across all subgroups, with others exhibiting convergence or divergence in life expectancy between groups.

It is important to note that this is just an indicative range; it is possible to conceive of scenarios with stronger improvements than the highest liability scenario outlined, or with weaker improvements than the lowest. It is equally important to note that the improvement rates associated with a given scenario are only suggestions that might fit the scenario described; another actuary could take the same narrative scenario and produce a materially different set of improvement rates.

Each of the scenarios outlined involves re-calibrating the CMI Mortality Projections Model. For consistency with the start point of the projection (2010) the CMI_2013 version of the model was used (this version was first published alongside Working Paper 69 in September 2013). This note sets out the core and advanced parameters used in generating scenarios for the NAPF project.

In some cases (for example “Cancer Revolution”) there were features required that cannot readily be modelled in the CMI Mortality Projections Model and were hence carried out separately. We also describe these adjustments. In all cases, we consider the impact of the scenario compared to the status quo projection.

The scenarios considered are:

Description	Type	Converge / diverge
Health Cascade	'Central'	Diverge
Improvement Decline	'Central'	Neither
Cancer Revolution	High	Neither
Challenging Times	Low	Diverge
Extended Youth	Extreme high	Converge
Back to the Fifties	Extreme low	Converge

The 'type' above refers to whether the scenario might be considered more central (i.e. in the realm that some might refer to as 'best estimates'), high or low improvements relative to the central or out towards the wider extremes of the range of the distribution. Note in particular that:

- we have provided an even number of 'central' scenarios

  This is deliberate as we do not want to suggest that there is one single 'central' scenario as, historically, this has tended to encourage a 'herding' of views.

- the extreme high and extreme low do not reflect the ***bounds*** of possible outcomes.

  For example, it is conceivable that life expectancy will go up in the future even faster than recent trends and thus exceed the 'Extended Youth' scenario.

Health cascade

Description

Recent improvements in life expectancy for the ‘golden cohort’ (generation born between two world wars) are believed to be driven by a number of behavioural changes (such as smoking cessation) and medical interventions (including free access to 24/7 medical care via the NHS).

A theory (supported by empirical data from the ONS on smoking cessation) is that uptake of such behaviours and services ‘cascades’ through society with the most educated (proxied by our ‘comfortable’ group) adopting the behaviours first and most fully. As the benefits of these behaviours become more evident so they ‘cascade’ through society.

This ‘health cascade’ is reflected in this scenario. Specifically the pace of longevity improvements for the ‘comfortable’ group is assumed to have ‘peaked’ and hence slows in the short term. In contrast, rapid improvements for the “hard-pressed” group persist in the medium term as we see the delayed impact of the uptake of healthy behaviours (in particular smoking cessation). The “making do” group experiences fast improvements over the short term but these tail off more quickly than the “hard-pressed’” group.

We also reflect that, longer term, new medical therapies / behavioural changes are likely to be accessed by the “comfortable” group, leading to a slightly faster reduction in their mortality. For women the outcome for the “making do / comfortable” group is based on the average of the making do and comfortable scenarios for men.

How we modelled this scenario

As noted in section 5, to allow for current rates of improvement to persist into the medium term, we need to adjust the convergence parameters (the key factor is the proportion of change remaining at mid-point).

Compared to the Core Parameters version of the CMI Mortality Projections Model, we have therefore carried out different adjustments for each group, as follows:

Hard pressed:
- 75% of change remaining at mid-point for age-period and cohort elements (ensuring some continued rise in improvements before tailing off to long term rate)
- 50% addition to convergence period for age-period and cohort elements, retaining maximum of 40 years (to ensure a gradual tail off to long term rate)
- Long-term rate of 1.5% pa

Making do:
- 60% of change remaining at mid-point for age-period and cohort elements
- 25% addition to convergence period for age-period and cohort elements, retaining maximum of 40 years
- Long-term rate of 1.5% pa

Comfortable:
- 50% of change remaining at mid-point for age-period and cohort elements (i.e. core parameters so improvements having peaked’)
- 0% addition to convergence period for age-period and cohort elements, retaining maximum of 40 years (i.e. core parameters)
- Long-term rate of 2.0% pa (i.e. higher than the other groups)

Detailed scenario output

The chart below indicates the pace of longevity improvement implied by the scenario for each group, as well as the indicative impact on liabilities for individuals of various ages. This scenario clearly impacts differently on each group, as well as varying significantly by age.

<table>
<tr><th rowspan="2" colspan="2">What does this scenario mean?</th><th colspan="3">Men</th><th colspan="2">Women</th></tr>
<tr><th>Hard-pressed</th><th>Making do</th><th>Comfortable</th><th>Hard-pressed</th><th>Making do / Comfortable</th></tr>
<tr><td colspan="2">Trends in period life expectancy (LE)*</td><td colspan="3">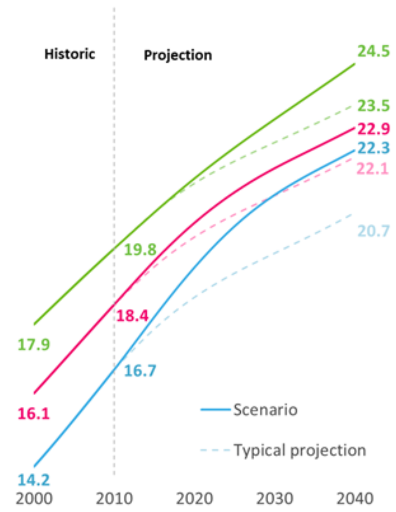
</td><td colspan="2">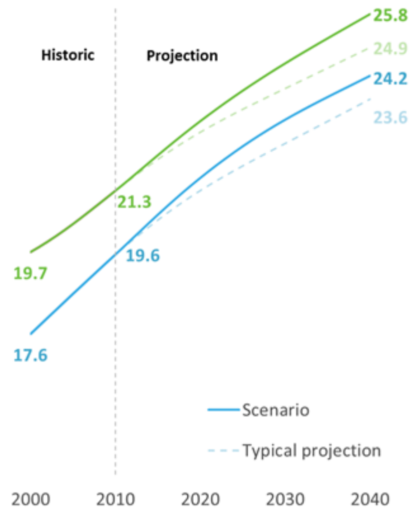
</td></tr>
<tr><td colspan="2">Pace of improvement</td><td>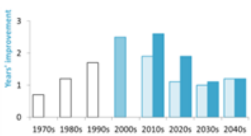</td><td>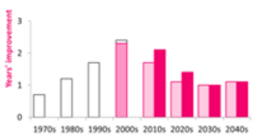</td><td>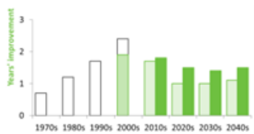</td><td>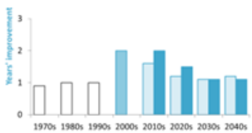</td><td>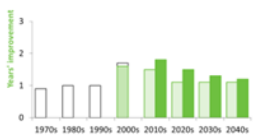</td></tr>
<tr><td rowspan="3">Liability impact* from age:</td><td>50</td><td>8%</td><td>3½%</td><td>5½%</td><td>2½%</td><td>4%</td></tr>
<tr><td>65</td><td>5%</td><td>2%</td><td>2½%</td><td>1½%</td><td>3%</td></tr>
<tr><td>80</td><td>4%</td><td>1½%</td><td>1%</td><td>1%</td><td>2%</td></tr>
</table>

*vs CMI model with CMI starting improvement and 1.5% long-term rate

Improvement Decline

Description

In this scenario we assume improvements will diminish over time, as the frequency and impact of medical advances diminish, coupled with rising obesity and other detrimental lifestyle factors. This means that the “Golden Cohort” of individuals born between the wars continue to exhibit faster improvements in longevity than those born either side.

The benefits of the healthy behaviours (smoking cessation) and introduction of the NHS are inherited by subsequent generations. However you can only give up smoking once. For subsequent generations, medical advances, and benefits of health interventions such as screening provide a driver for some continued improvements, but the behaviours and lifestyle of younger cohorts throughout their life course result in longevity improvements slowing almost to stagnation.

Specifically, long term improvements for the post WWII birth generations drop to around 3/4 year per decade (compared to the long run historic average of 1 year per decade)

How we modelled this scenario

In terms of the CMI Mortality Projections Model, we have adjusted the long-term rate of mortality improvement as follows:

- Age-period rate reduced to 0.75% pa for all ages
- Cohort rate
  - 0% pa for cohorts born before 1920 and after 1954
  - 0.75% pa for cohorts born between 1929 and 1945
  - Moving linearly from 0% to 0.75% for cohorts born between 1920 and 1928 and between 1946 and 1954

Graphically, the shape of the long-term cohort improvement rate can be shown as follows:

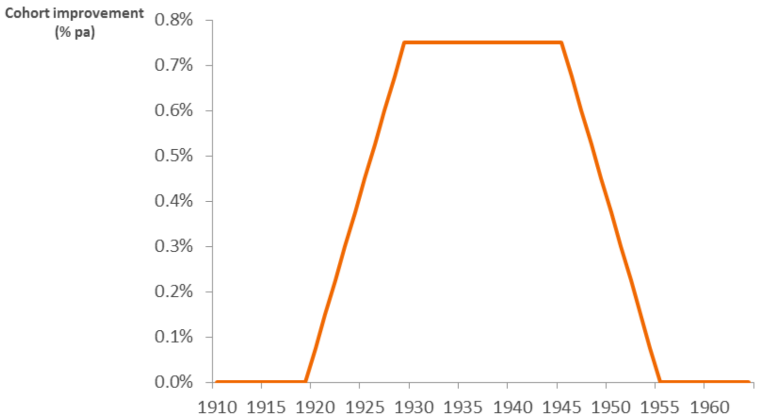

As per the description, this means that mortality improvements for the cohort born between 1930 and 1945 are higher than for the generations either side.

Under this parameterisation, long-term improvements tend towards 1.5% per annum for the cohort born between 1930 and 1945, becoming gradually lower for all other birth cohorts, with a minimum of 0.75% pa for the cohorts born before 1920 and 1954.

Detailed scenario output

The chart below indicates the pace of longevity improvement implied by the scenario for each group, as well as the indicative impact on liabilities for individuals of various ages. This scenario is similar for each group, but varies somewhat by age.

What does this scenario mean?		Men: Hard-pressed	Men: Making do	Men: Comfortable	Women: Hard-pressed	Women: Making do / Comfortable
Trends in period life expectancy (LE)*		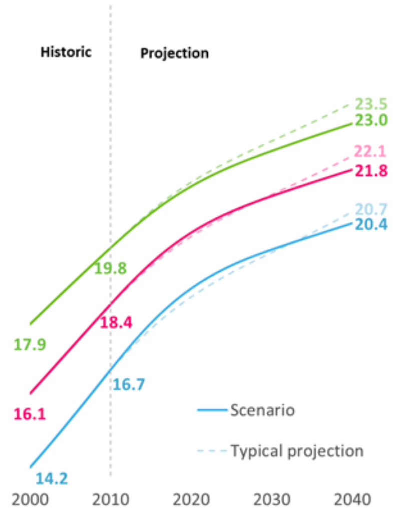			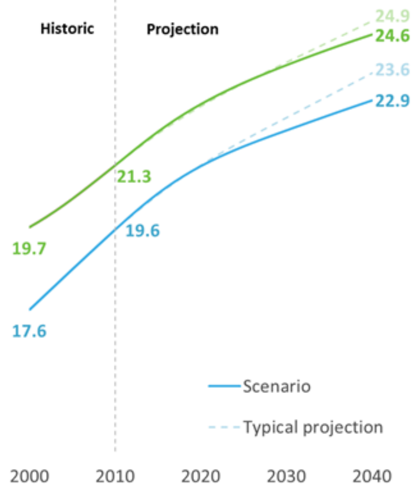	
Pace of improvement		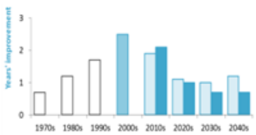	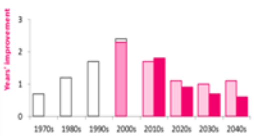	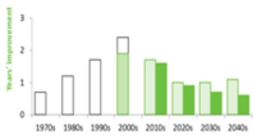	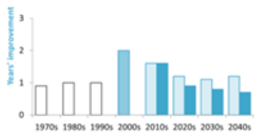	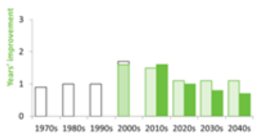
Liability impact* from age:	50	-2%	-2½%	-3%	-3½%	-2%
	65	½%	0%	-1%	-1%	½%
	80	½%	0%	-1½%	-1%	0%

*vs CMI model with CMI starting improvement and 1.5% long-term rate

Cancer Revolution

Description

The eradication of a significant cause of death (typically, but not always, cancer) is a typical suggestion when projecting mortality improvement using scenarios. Whilst a cure for any specific major cause of death is unlikely over any reasonable time horizon, this type of scenario is a useful way of illustrating the impact of any material advance in healthcare.

This scenario would require a step-change in cancer treatment, for example effective national screening (both traditional and genetic) and perhaps a 'pill' being developed to target hard-to-treat cancers. We assume that these changes begin to be available in 2025, and are fully taken up over the following 5 year period – this was chosen as being the fastest reasonable time taken for widespread takeup of such an innovation in healthcare.

Very broadly speaking, in the UK population as a whole cancer accounts for around

- 20% of deaths below age 55,
- 40% between ages 55 and 79, and
- 25% at age 80 and above

To keep the scenario reasonably simple, we assume these proportions are the same for each of the groups. This assumption is very roughly borne out in practice based on research on cancer incidence by IMD.

Note this does not mean that the initial cancer mortality rates are the same for each subgroup – because overall death rates are higher in the "hard-pressed" category, cancer mortality rates are also assumed to be higher.

Because older individuals are more likely to have multiple diseases, we have assumed that, whilst cancer is eradicated as a cause of death, the reduction in mortality is less than implied by the percentages above, as some people who would previously have died of cancer die of another cause relatively soon afterward.

We have also assumed that the long-term rate of improvement 'post-cancer' is slightly lower than it would have been 'pre-cancer' because part of the previously assumed long term rate is likely to have been driven by some gradual improvements via cancer interventions.

How we modelled this scenario

In terms of the CMI Mortality Projections Model, to model the 'post-cure' trajectory of improvements, we have used a long-term rate of mortality improvement structured by age as follows:

- Below 55: 1.2% pa
- 55-79: 1.1% pa
- 80+: 1.5% pa

We have assumed Core Parameters after age 90 (ie falling linearly to 0% at age 120), and have used the same structure for all groups.

To model the 5-year period (2025-2030) over which a cure is assumed to take effect, we apply a 'patch' to improvements (this cannot be readily modelled within the CMI model structure) of:

- +4% p.a. for age 20-54
- +8% p.a. for age 55-79
- +4% p.a. for age 80+

So mortality rates from 2030 will be either 18.5% ($1-0.96^5$) or 34.1% ($1-0.92^5$) lower than the projections produced by the CMI model. This corresponds roughly to the proportion of deaths relating to cancer figures noted above, which means we are modelling a scenario where the significant majority of cancer deaths have been removed.

Detailed scenario output

The chart below indicates the pace of longevity improvement implied by the scenario for each group, as well as the indicative impact on liabilities for individuals of various ages.

This scenario is similar for each group, but varies significantly by age. In particular, there is little impact on current 80 year olds as these individuals will be almost 100 before the benefit of the 'cure' is seen

What does this scenario mean?		Men			Women	
		Hard-pressed	Making do	Comfortable	Hard-pressed	Making do / Comfortable

Trends in period life expectancy (LE)*

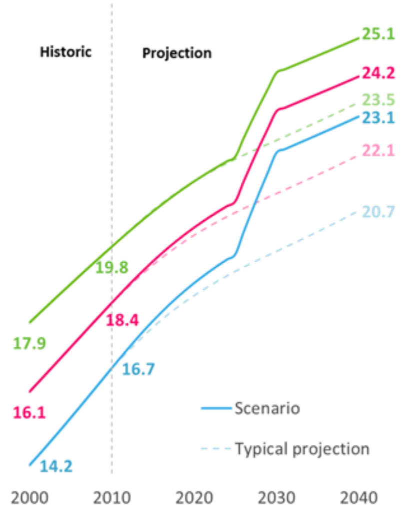

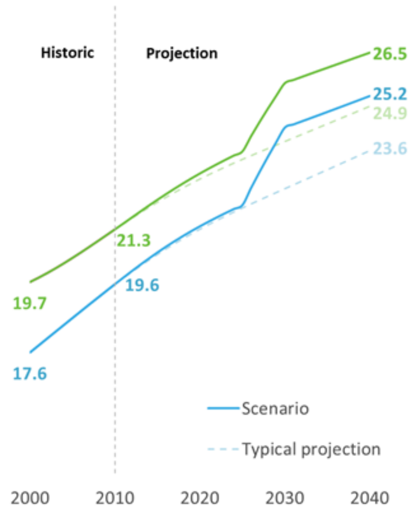

Pace of improvement

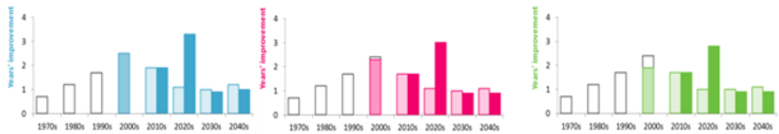

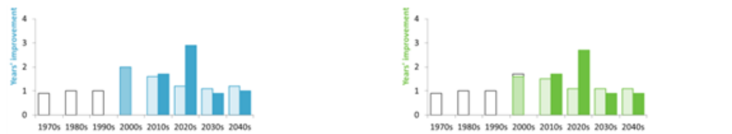

Liability impact* from age:		Men: Hard-pressed	Men: Making do	Men: Comfortable	Women: Hard-pressed	Women: Making do / Comfortable
	50	12%	9%	7%	7%	6%
	65	5%	4%	3%	3½%	4%
	80	1½%	½%	-1%	0%	1%

*vs CMI model with CMI starting improvement and 1.5% long-term rate

Challenging Times

Description

In this scenario we consider the implications of climate change and finite resources, for example fossil fuels. We consider the implications of the possibility that we have reached ‘peak oil flow’ and that the availability of oil will become a constraint to economies in the future.

A consequence of this could be increasing fuel prices, leading to severe constraints in finances and funding of the NHS. Alongside this, reduced access to and increased cost of imported food stocks could have a detrimental impact on health outcomes via for example greater difficulty in maintaining healthy fruit and vegetable rich diets throughout the year.

We reflect this by assuming that a significant proportion of the “hard-pressed” and “making do” groups are unable to afford their basic needs (heating, fuel, medicine) and that this leads to life expectancy ceasing to improve. In contrast we assume that resource constraints impacts are less severe on average for the “comfortable” group, meaning that this scenario leads to longevity improvements that are below the long-term trend, but above zero for this group.

Further, we have included an impact of two consecutive abnormally harsh winters (leading to no overall improvement for two years) earlier in the scenario, with a relatively high improvement in the third year.

For women the outcome for the “making do / comfortable” group is based on the average of the “making do” and “comfortable” scenarios for men.

How we modelled this scenario

In terms of the CMI Mortality Projections Model, we have used a long-term rate of mortality improvement of

- 0% pa for the Hard-pressed and Making-do groups, and
- 1% pa for the Comfortable group.

The only change we have made to the Core Parameters is to increase the convergence parameter to 75% at midpoint for the “hard-pressed” and “making-do” groups. This means that we converge slightly more slowly to the lower long-term rate for these groups.

Core Parameters were used throughout for the “comfortable” group.

The ‘abnormal winter’ assumed to occur in 2012 and 2013 cannot be modelled directly within the CMI Mortality Projections Model. A manual amendment to reduction factors was applied in a separate spreadsheet to produce improvements of:

- 0% pa for 2012 and 2013
- 4% for 2014

This is broadly equivalent to the improvements that were actually seen following the harsh 2012-13 winters, and which have been followed by a relatively strong improvement in 2014 to date.

Detailed scenario output

The chart below indicates the pace of longevity improvement implied by the scenario for each group, as well as the indicative impact on liabilities for individuals of various ages. This scenario impacts more heavily on the 'hard-pressed' and 'making do' groups, and particularly on the younger individuals in these groups.

What does this scenario mean?	Men			Women	
	Hard-pressed	Making do	Comfortable	Hard-pressed	Making do / Comfortable
Trends in period life expectancy (LE)*	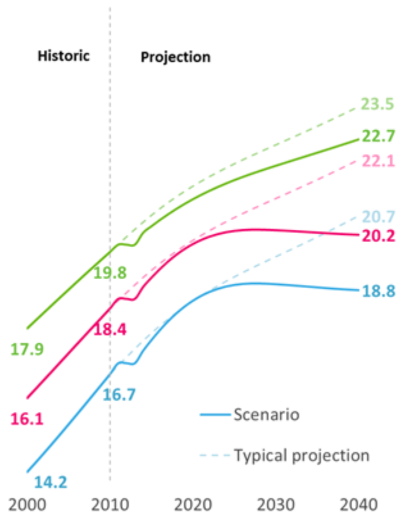			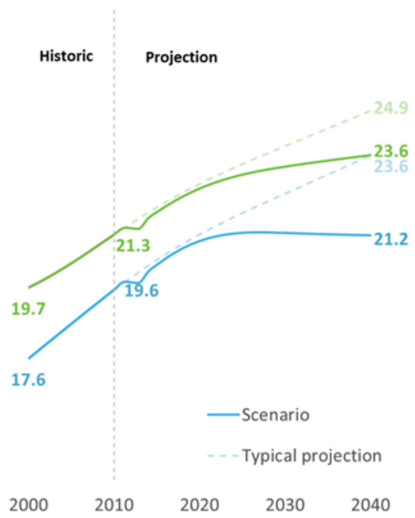	
Pace of improvement	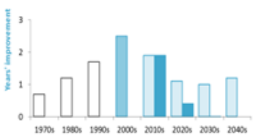	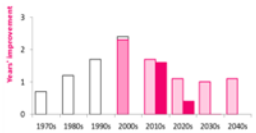	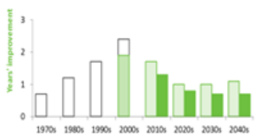	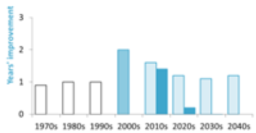	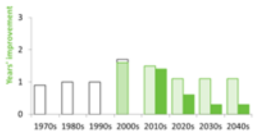
Liability impact* from age: 50	-10½%	-10½%	-4½%	-12%	-7%
65	-4½%	-5%	-3%	-6%	-3%
80	-3½%	-4%	-3%	-4½%	-2½%

*vs CMI model with CMI starting improvement and 1.5% long-term rate

Extended Youth

Description

Across the whole UK population, improvements in life expectancy for a man aged 65 over the 2000s were 2.4 years. This has increased from 1.7 years in the 1990s, 1.2 years in the 1980s and 0.7 years in the 1970s. For women, the increase in life expectancy over the 2000s was 1.7 years, compared to around 1 year per decade over the 1970s-90s.

However, the experience of each of our subgroups has improved in a different way to the population as a whole, with the 'hard-pressed' male group seeing a 2.5 year improvement over the 2000s, the 'making do' group a 2.3 year improvement, and the 'comfortable' group a 1.9 year improvement. For women the 'hard-pressed' group saw an improvement of 2 years, whilst the 'making do / comfortable' subgroup saw a 1.6 year improvement in life expectancy.

In this scenario we consider the possibility that some combination of factors will lead to these improvements being sustainable over the longer term. Just as it would have been hard to predict the last 40 years of strong improvements back in 1970 - let alone the drivers - we do not offer a very specific narrative; however possible contributory factors could be a combination of highly successful screening programs, poly-pills, smart pills aimed to improve drug adherence, ageing medicine breakthroughs increasing survivorship from multiple diseases of later life, increased later life activity and exercise and reduced obesity.

We can translate the life expectancy improvements listed above into long-term rates of improvement for each group:

- For males, a 2.5 year-per-decade improvement for the 'hard-pressed' group represents a long-term rate of around 3.25% pa
- For the 'making-do' male group a 2.3 year-per-decade improvement represents a long-term rate of around 2.75% pa
- For the 'comfortable' male group a 1.9 year-per-decade improvement represents a long-term rate of around 2.25% pa
- For the 'hard-pressed / making-do' female group a 2 year-per-decade improvement represents a long-term rate of around 2.5% pa
- For the 'comfortable' female group a 1.6 year-per-decade improvement represents a long-term rate of around 1.75% pa.

It is worth noting that this scenario leads to convergence in (period) life expectancies - by around 2044 the 'hard-pressed' group has caught the 'making do' group and by 2048 it has caught the 'comfortable' group.

This also implies that the cohort life expectancy will converge – the life expectancy produced by this scenario is higher for a 'hard-pressed' male aged 65 in 2030 than for a 'making do / comfortable' female.

Whilst this may not be a very plausible outcome, it illustrates how life expectancies would change over the very long term if a simple straight-line extrapolation were adopted for each group.

How we have modelled this scenario

In terms of the CMI Mortality Projections Model, we have used long-term rates of mortality improvement between 1.75% and 3.25%, as listed above.

We have assumed that these rates persist at all ages – that is, we have removed the taper between ages 90 and 120 from the Core Parameters of the CMI model.

Detailed scenario output

The chart below indicates the pace of longevity improvement implied by the scenario for each group, as well as the indicative impact on liabilities for individuals of various ages.

This scenario impacts more heavily on the 'hard-pressed' groups, and particularly on the younger individuals in these groups. It also has a much higher impact on male than female groups.

<table>
<tr><th colspan="2" rowspan="2">What does this scenario mean?</th><th colspan="3">Men</th><th colspan="2">Women</th></tr>
<tr><th>Hard-pressed</th><th>Making do</th><th>Comfortable</th><th>Hard-pressed</th><th>Making do / Comfortable</th></tr>
<tr><td colspan="2">Trends in period life expectancy (LE)*</td><td colspan="3">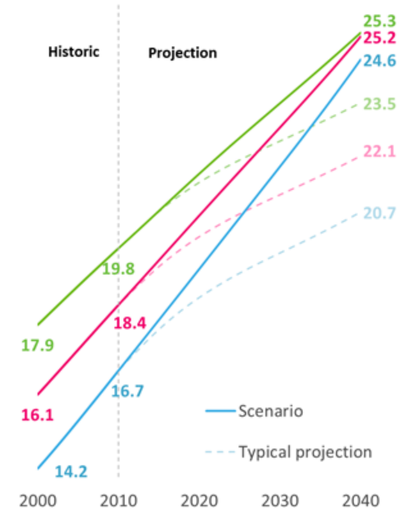
</td><td colspan="2">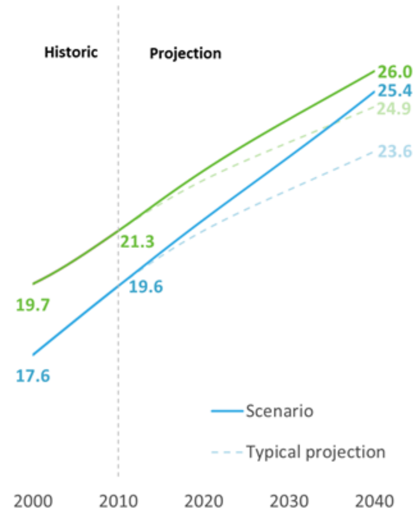
</td></tr>
<tr><td colspan="2">Pace of improvement</td><td>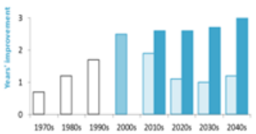</td><td>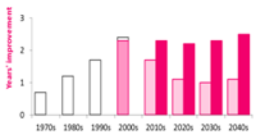</td><td>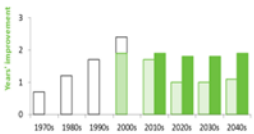</td><td>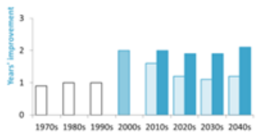</td><td>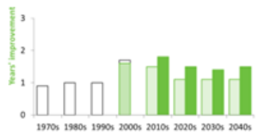</td></tr>
<tr><td rowspan="3">Liability impact* from age:</td><td>50</td><td>27½%</td><td>21%</td><td>12%</td><td>12%</td><td>6½%</td></tr>
<tr><td>65</td><td>12%</td><td>9½%</td><td>5%</td><td>5½%</td><td>4%</td></tr>
<tr><td>80</td><td>9%</td><td>7%</td><td>3%</td><td>4½%</td><td>3%</td></tr>
</table>

*vs CMI model with CMI starting improvement and 1.5% long-term rate

Back to the Fifties

Description

One of the great success stories of the 20th century has been the rapid improvement in health outcomes and commensurate rise in life expectancy. With modern medicine and technology advances we are naturally inclined to assume life expectancy will continue to rise. However this has not always been the case.

For this scenario we have assumed that mortality rates will rise in the future and so life expectancy will fall. As this is a 'doomsday scenario' (though as noted above it is possible to conceive of more extreme outcomes), we assume that this will happen very soon e.g. by the end of this decade.

Like the continuation of trend scenario, we do not offer a very specific narrative for this scenario, instead suggesting it would involve a combination of a number of societal and health changes, possibly including**:**

- widespread antibiotic resistance;
- obesity;
- severe austerity impacting the NHS (possibly to point of dissolution);
- severe resource constraints (particularly oil but also and rare earth metals) impacting heating/ access to imported fruit and veg / medical equipment; and
- severe adverse effects of climate change.

How we have modelled this scenario

Reducing life expectancy improvements to zero **from 2010** would lead to a reduction in liabilities of around 15% compared to the baseline scenario.

We felt that this was a suitable reduction in liability but not a plausible pattern of improvement as there would be an extreme discontinuity in the improvement rates.

Instead, we use a **negative** long-term rate of 'improvement' which is reached more quickly than would typically be assumed within the CMI model.

In terms of the CMI Mortality Projections Model, we have therefore used a long-term rate of mortality improvement of -1% pa for all groups. We have assumed that this rate persists at all ages – that is, we have removed the taper between ages 90 and 120 from the Core Parameters model.

We have also assumed that the proportion of change remaining at the midpoint is 25% for both the period and cohort component, for all groups, which means that improvement rates move more quickly towards the long-term rate than would be assumed in the Core Parameters model.

Detailed scenario output

The chart below indicates the pace of longevity improvement implied by the scenario for each group, as well as the indicative impact on liabilities

This scenario impacts similarly on all groups, but has a particularly high impact on younger individuals in each group.

<table>
<tr><th rowspan="2" colspan="2">What does this scenario mean?</th><th colspan="3">Men</th><th colspan="2">Women</th></tr>
<tr><th>Hard-pressed</th><th>Making do</th><th>Comfortable</th><th>Hard-pressed</th><th>Making do / Comfortable</th></tr>
<tr><td colspan="2">Trends in period life expectancy (LE)*</td><td colspan="3">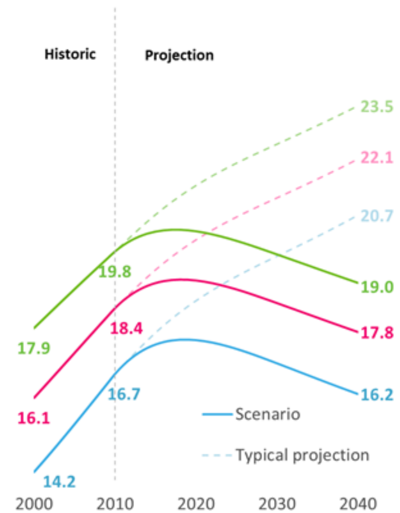
</td><td colspan="2">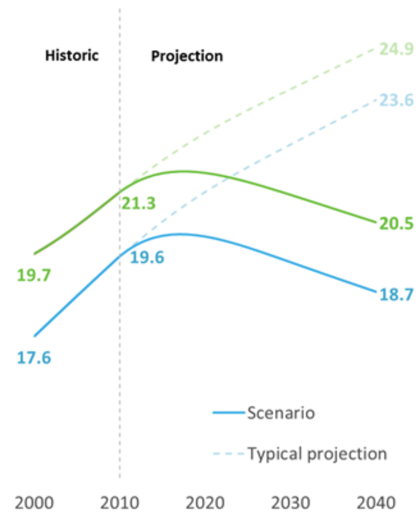
</td></tr>
<tr><td colspan="2">Pace of improvement</td><td colspan="3">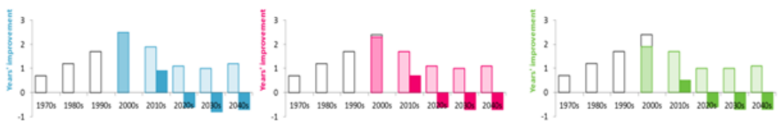</td><td colspan="2"></td></tr>
<tr><td rowspan="3">Liability impact* from age:</td><td>50</td><td>-23%</td><td>-21½%</td><td>-21½%</td><td>-22½%</td><td>-20½%</td></tr>
<tr><td>65</td><td>-12%</td><td>-11½%</td><td>-12%</td><td>-12½%</td><td>-11½%</td></tr>
<tr><td>80</td><td>-9½%</td><td>-10%</td><td>-11%</td><td>-11%</td><td>-10%</td></tr>
</table>

*vs CMI model with CMI starting improvement and 1.5% long-term rate