Longevity Basis Risk Phase 2

Assessing Basis Risk for Longevity Transactions

by Macquarie University

Research Report

Assessing Basis Risk for Longevity Transactions – Phase 2

Research investigation and report by Macquarie University
for the Institute and Faculty of Actuaries and
the Life and Longevity Markets Association

Macquarie University

Jackie Li   PhD, PhD, FIAA
Leonie Tickle   FIAA
Chong It Tan   PhD, FSA, FIAA, CERA

University of Waterloo

Johnny Siu-Hang Li   PhD, FSA

November 2017

Introduction

On behalf of the joint Longevity Basis Risk Working Group, established jointly by the Life and Longevity Markets Association (LLMA) and the Institute and Faculty of Actuaries (IFoA), I am delighted to introduce the results of Phase 2 of our research into the area of longevity basis risk.

This technical report sets out a practical implementation of the methodology developed in Phase 1 of this work for a number of real-life datasets, and shows how index-based longevity hedges might be used to provide a significant reduction in the exposure of insurers, reinsurers, banks, and pension schemes to longevity risk.

This report forms the key output of the second phase of this work, and has been commissioned and funded by the IFoA and the LLMA, and undertaken on our behalf by Macquarie University.

The importance of longevity basis risk

Longevity basis risk arises because different populations, or subpopulations, will inevitably experience different longevity outcomes. This is a significant issue for those wishing to hedge longevity risk using a published mortality index – whether they be pension schemes, insurers, reinsurers, or banks. Put simply, actual longevity outcomes, and therefore cashflows, of the hedged portfolio will differ from those under the hedging instrument.

In addition, longevity basis risk can also present a wider issue for insurers using, in their reserving models, external data, such as population data, rather than their own policy data. The need to quantify any potential basis risk is receiving increasing focus, particularly under Solvency II.

Phase 1 of this research

Phase 1 of this research focused on producing a methodology that could be used to quantify longevity basis risk and assess the level of risk reduction that might be brought about through the use of an index-based longevity hedge. It resulted in a framework that recognised the fact that different practitioners will have different portfolios, different volumes and histories of experience data, and different constraints of the models that they can use in practice. It provided specific models and techniques for different situations, which delivered a great starting point for Phase 2 of this work.

Phase 2 of this research

The purpose of Phase 2 is to:
- determine the most relevant metrics for measuring longevity basis risk and hedge effectiveness;
- apply the approach in Phase 1 to realistic worked examples based on appropriate data;
- present a robust quantification of basis risk to third parties such as regulators; and
- investigate the potential limitations of the time series processes.

This report therefore builds on the Phase 1 findings and applies the models and methodologies of Phase 1 to a number of different datasets, portfolio sizes, and hedge structures in order to assess the level of longevity basis risk that would result under a wide range of scenarios and under a range of potential risk reduction metrics. The report therefore contains an extensive set of results within which we hope that all practitioners will find information that is useful and relevant to their individual circumstances.

We are delighted to be able to present the results of this research and hope it will prove of value to practitioners and enable an important step change in the ability to assess longevity basis risk. On behalf of the Longevity Basis Risk Working Group, I would like to thank the research team for all of their hard work in delivering this unique and valuable output.

Robert Bugg

Chair of the LLMA and IFoA Joint Longevity Basis Risk Working Group

Reliance and Limitations

This report has been produced by Macquarie University for the Longevity Basis Risk Working Group (LBRWG) of the Institute and Faculty of Actuaries (IFoA) and the Life and Longevity Markets Association (LLMA).

This report is addressed to the LBRWG. It may be shared with members of the IFoA and LLMA and other relevant third parties. This report does not constitute advice and should not be considered a substitute for specific advice in relation to individual circumstances. Whilst care has been taken to ensure that the report is accurate, current, and useful, neither Macquarie University, University of Waterloo, Mercer Australia, the IFoA, nor the LLMA (collectively, the Parties) makes any warranty or representation, express or implied, as to the report's accuracy, currency, and usefulness. The Parties disclaim all liability for any loss or damage suffered of whatever nature (direct, indirect, consequential, or other) as a result of or in relation to the use of this report and for actions taken by third parties as a consequence of the information contained in this report.

Principal Researchers

Associate Professor Jackie Li   PhD, PhD, FIAA

Department of Applied Finance and Actuarial Studies

Macquarie University, Australia

Associate Professor Johnny Siu-Hang Li   PhD, FSA

Department of Statistics and Actuarial Science

University of Waterloo, Canada

Professor Leonie Tickle   FIAA

Department of Applied Finance and Actuarial Studies

Macquarie University, Australia

Dr Chong It Tan   PhD, FSA, FIAA, CERA

Department of Applied Finance and Actuarial Studies

Macquarie University, Australia

Executive Summary

As per the call for research project proposals by the Longevity Basis Risk Working Group (LBRWG), the aim of the overall project is 'to develop a readily-applicable methodology for quantifying the basis risk arising from the use of population-based mortality indices for managing the longevity risk inherent in specific blocks of pension benefits or annuitant liabilities'. The project was funded by the Institute and Faculty of Actuaries (IFoA) and the Life and Longevity Markets Association (LLMA). Phase 1 of this project was completed by Cass Business School and Hymans Robertson LLP in December 2014, in which a decision tree framework was developed to provide a guide on how to choose a two-population mortality model for the reference and book populations. It includes the M7-M5 model, the CAE+Cohorts model, and the characterisation approach.

The current Phase 2 of the project has been undertaken by Macquarie University. It focuses on putting the earlier work conducted in Phase 1 into practice and measuring longevity basis risk under practical circumstances. This longevity basis risk arises from the potential mismatch between the longevity hedging instrument (reference population) and the pension or annuity portfolio (book population) being hedged. Generally speaking, there are three main sources of longevity basis risk, namely demographic basis risk (demographic or socioeconomic differences), sampling basis risk (random outcomes of individual lives), and structural basis risk (differences in payoff structures). While Phase 1 has proposed some two-population mortality models for measuring demographic basis risk, Phase 2 takes all the three risk components into account and considers more realistic hedging scenarios using UK and Australian population and industry data.

In accordance with the call for research project proposals, the major objectives of Phase 2 are to: (a) determine the most relevant metrics for measuring longevity basis risk and hedge effectiveness; (b) apply the approach in Phase 1 to realistic worked examples based on appropriate data; (c) present a robust quantification of basis risk to third parties such as regulators; and (d) investigate the potential limitations of the time series processes.

Risk Metrics

The effectiveness of an index-based longevity hedge can be described as how much longevity risk is transferred away via the hedge. The retained portion of the risk can then be seen as longevity basis risk. In line with the current literature, we define the level of longevity risk reduction for a certain longevity hedge on a pension or annuity portfolio as:

$$\text{longevity risk reduction} = \left(1 - \frac{\text{risk(hedged)}}{\text{risk(unhedged)}}\right) \times 100\%$$

in which risk(unhedged) and risk(hedged) are the chosen measure of the portfolio's longevity risk before and after taking the hedge. This metric represents the percentage of the portfolio's initial longevity risk that is reduced by the hedge. Regarding the risk measure, we consider the usual choices in the literature including the variance, standard deviation, 99.5% value-at-risk (VaR), and 99.5% expected shortfall (conditional VaR). The 99.5% VaR is of particular interest in practice, as this concept is adopted in the assessment of the Solvency Capital Requirement (SCR) under Solvency II. Different longevity hedging schemes could lead to different levels of risk reduction and potentials for capital savings. More details about the risk metrics and risk measures can be found in Section 3 and Appendix I.

Realistic Examples

Hypothetical cases of pension portfolios, based on UK and Australian industry datasets, are studied in Section 4, where standardised longevity swaps are used to construct the longevity hedge. In such a swap, two series of future cash flows are exchanged: one is linked to the percentage of the reference population who are alive on the payment dates, and the other is fixed at the start of the transaction. A number of hedging scenarios are examined, including a single cohort or multiple cohorts, an open or closed pension plan, varying portfolio sizes, and different levels of precision in the use of longevity swaps, for each UK and Australian subgroup. The major finding, from applying the two-population mortality models from Phase 1, is that the risk reduction levels (regarding the present value of the aggregate position after hedging) are often around 50% to 80% for a large portfolio, while the risk reduction estimates are usually smaller than 50% for a small portfolio. The exact risk reduction level depends on the particular hedging scenario being considered. The modelling procedure, technical information, and Excel VBA codes are stated in Section 3 and Appendices I and II.

Robust Quantification

An extensive sensitivity analysis on the hedging results by varying the initial model settings and assumptions is set out in Section 5. Based on the simulated results, the most important modelling assumptions and settings are the coherence property (i.e. constant book-to-reference ratio of mortality rates at each age in the long run) and behaviour of simulated future variability, portfolio size, data size and characteristics, type of hedging instrument, simulation method, and additional model features such as mortality structural changes. Comparatively, the other variations tested seem to have rather limited impact on the computed hedging results. In addition, some backtesting and scenario testing are performed on the longevity hedging strategy. The testing results suggest that the hedging strategy works reasonably well when there are sizable unanticipated mortality improvements. By contrast, if the major mortality trends are well captured by the modelling process and these trends endure over time, the longevity hedge would not cause much change to the aggregate position of the pension portfolio. Moreover, when the longevity shocks are more significant for the book population than for the reference population, due to longevity basis risk, the reduction in the portfolio loss from the hedge would still be considerable, given that the longevity shocks on the reference and book populations are broadly in the same direction.

Furthermore, a brief summary of all the simulated hedging results is presented in Section 7. A qualitative assessment table and a simple 'rule-of-thumb' formula are demonstrated as a possible quick guide for practitioners and regulators in assessing the effectiveness of an index-based longevity hedge. For instance, an assessment table can contain a number of major qualitative questions about the extent of longevity basis risk, in which a score of 0 to 10 (from mild to significant) is given to each question. The total score then represents a rough estimate of the level of longevity risk reduction. Some relevant questions include the size of the pension plan, how the book and reference populations are related, the assumed pace of reaching coherence between the two populations, and the presence of mortality structural changes. Moreover, a simple linear regression can also be applied to the risk reduction levels using a range of explanatory variables, which are related to the pension plan characteristics, hedging structure, model settings, and assumptions, as a simple way of quantifying longevity basis risk. Note that these quick guides are model-dependent and are highly specific to the datasets being modelled. If suitable resources and expertise are available, practitioners are encouraged to use the detailed technical information in this report or other relevant references to construct their own models and perform more accurate calculations. Finally, a number of practical issues in implementing index-based hedging solutions are also discussed.

Time Series

A sensitivity analysis on time series modelling is conducted in Section 6. Given the likely short data length of the book population, the feasible choices of time series processes are quite limited. Several modifications to the original time series processes in the M7-M5 and CAE+Cohorts models are tested, such as choosing a higher order for the autoregressive processes in the book component, using univariate time series processes in the reference component under the M7-M5 model, relaxing the independence assumption between the time series error terms of the reference and book components, and applying some other extensions like introducing non-coherence via an extra linear term. The testing results indicate that the time series modelling assumptions sorted in a decreasing order of importance are the behaviour of simulated future variability of the book component, the pace of reaching coherence, and then the other correlation assumptions. Proper judgement, reference materials, experts' opinions, and thorough testing are needed for making appropriate time series modelling assumptions in practice. Further research is also required when more data of longer periods and for different kinds of book populations can be collected in the future. More technical details and the Excel VBA codes for time series modelling are provided in Appendices I and II.

Table of Contents

Table of Contents

1   Introduction

Continual decline in mortality is a global phenomenon. This persistent trend has been driven mainly by substantial enhancements in nutrition, hygiene, medical technology, health care, and education since the last century. While it is certainly no mean feat for humans, it imposes a significant challenge on pension plan sponsors and annuity providers. The so-called longevity risk is the risk that pension plans or annuity portfolios pay more than anticipated due to unexpected mortality improvements. Broadly speaking, the risk consists of two components, namely systematic longevity risk and non-systematic longevity risk. The latter can usually be alleviated by increasing the portfolio size, but the former cannot be diversified by pooling in the same way. At present, governments and insurance companies are generally very cautious about assuming too much longevity risk.

There are a few approaches for financial institutions to manage longevity risk (Cairns et al. 2008, Li and Haberman 2015). The first one is insurance and reinsurance, in which the risk is transferred to an insurer or reinsurer by paying a premium. For instance, an insurer can enter into a reinsurance contract with a reinsurer to hedge the risk, or a pension plan can buy annuities from a life office to transfer the risk away. The second is natural hedging, which makes use of the opposite changes in the values of annuities and life insurances. This diversification strategy may be feasible for certain large institutions which have the financial structure and resources to sell both kinds of policies. It may also be facilitated by using, say, a mortality swap to construct an external hedge between two independent parties, such as an annuity provider and another separate life insurer.

The third approach, which has been in the limelight in recent years, is to adopt capital market solutions, including insurance securitisation, mortality- or longevity-linked securities, and derivatives. Insurance securitisation involves securitising a line of business as a complex bundle and selling the resulting highly structured securities to market investors. Some popular bespoke de-risking solutions currently like buy-ins, buy-outs, and longevity swaps (LCP 2012) are also tailored transactions for hedging particular portfolios. In contrast, standardised mortality- or longevity-linked securities and derivatives have their cash flows linked to a selected reference population (i.e. index-based), rather than the population underlying the portfolio to be hedged. Consequently, there would be a potential mismatch between the hedging instrument and the portfolio, in terms of demographic differences. There are also other concerns, e.g. a small portfolio would have high sampling variability, which makes it more likely to deviate from the experience of the reference population, and also the payoff structures (i.e. timing and amounts) would often be different between the hedging tool and the portfolio being hedged. These discrepancies lead to the concept of longevity basis risk, which is at the heart of this research project.

In 2010, the Life and Longevity Markets Association (LLMA) was established in the UK by several global insurers, reinsurers, and investment banks. Its mission is to promote the development of a liquid 'life market', which provides a platform for insurers, reinsurers, and market investors to trade various longevity- and mortality-linked assets and liabilities. Particularly, the capital market is huge and has much potential to absorb longevity risk from insurers and pension plans in exchange for appropriate risk-adjusted returns. Moreover, certain investors may want to diversify across a new market sector of longevity, which is arguably uncorrelated with traditional asset classes. Since then, there have been many interesting developments in the life market (Tan et al. 2015). For example, in 2011, the LifeMetrics indices, released by J.P. Morgan, the Pensions Institute, and Towers Watson back in 2007, were transferred to the LLMA for the purpose of building a global benchmark. J.P. Morgan issued a £70 million 10-year q-forward contract for the Pall (UK) pension fund, which was designed to hedge the value of the pension liabilities. In 2012, Deutsche Bank provided a €12 billion index-based longevity solution for Aegon in the Netherlands, in which the Dutch population was taken as the index. Then in 2013, Deutsche Bank released the Longevity Experience Option, which was a 10-year call option on 10-year forward survival rates, based on the England and Wales and Netherlands LLMA longevity indices. In 2014, the Mercer Global Pension Buy-out Index was introduced, showing the benchmark prices of eighteen independent third-party insurers in the UK, US, Canada, and Ireland. In the same year, Berkshire Hathaway made a £780 million quota-reinsurance transaction with Pension Insurance Corporation, which is one of the possible signs of an increasing demand for pension annuity books to offset in-house life insurance books. Some more recent longevity transactions are listed below (www.artemis.bm/library/longevity_swaps_risk_transfers.html).

Table 1.1   Recent longevity swaps transactions

Pension Fund / Sponsor	Provider(s)	Solution(s)	Amount	Date
Delta Lloyd	RGA Re	Index-based longevity swap	€12 billion	Jun 2015
Aegon	Canada Life Re	Longevity swap and reinsurance	€6 billion	Jul 2015
Manweb (ScottishPower)	Abbey Life	Longevity swap	£1 billion	Aug 2016
AXA France	RGA Re	Longevity swap and reinsurance	€1.3 billion	Nov 2016
Pension Insurance Corporation	SCOR	Longevity swap and reinsurance	£1 billion	Jul 2017
British Airways Pension Scheme	Partner Re, Canada Life Re	Longevity swap and reinsurance	£1.6 billion	Aug 2017

While most longevity deals to date have been bespoke transactions, index-based solutions and standardised products could draw more interest from financial entities both within and

outside the insurance and pension world. They have considerable potential to provide effective risk management at lower costs and offer significant capital savings. The major types of index-based derivatives proposed in the literature include longevity bond (Blake and Burrows 2001), longevity swap (Dowd 2003), q-forward (Coughlan et al. 2007), S-forward (LLMA 2010), K-forward (Chan et al. 2014), mortality option (Cairns et al. 2008), and survivor option (Dowd 2003). Some of them have already been issued and tested in practice, with different levels of success. As noted in the third approach above, however, the potential mismatch between the longevity hedging instrument (reference population) and the pension or annuity portfolio (book population) being hedged gives rise to longevity basis risk, which comprises demographic basis risk (demographic or socioeconomic differences), sampling basis risk (random outcomes of individual lives), and structural basis risk (differences in payoff structures). In order to address this critical issue, Phase 1 of the research project have earlier developed a decision tree framework as a guide for selecting a suitable two-population mortality model for the reference and book populations. The major choices are the M7-M5 model, the CAE+Cohorts model, and the characterisation approach, with a focus on demographic basis risk. Accordingly, the current Phase 2 aims to put the work conducted in Phase 1 into practice and assess longevity basis risk under practical circumstances. Specifically, it takes all the three risk components into account and examines more realistic hedging scenarios using various UK and Australian population and industry datasets.

The structure of this Phase 2 report is as follows. Section 2 describes the historical mortality levels and improvements of different subgroups in three industry datasets. Section 3 explains the modelling procedure of longevity basis risk adopted in this project and the way to compute the level of longevity risk reduction from an index-based longevity hedge. Section 4 examines hypothetical hedging scenarios of pension portfolios using standardised longevity swaps, and estimates the corresponding risk reduction levels. Section 5 conducts an extensive sensitivity analysis on the hedging results from Subsection 4.1 via making a series of changes to the initial model settings and assumptions, and performs backtesting and scenario testing on the hedging strategy. Section 6 carries out a sensitivity analysis on the choice of time series modelling, under the data constraint of the book population. Section 7 provides a brief summary of the hedging results in this report, proposes a qualitative assessment table and a 'rule-of-thumb' formula as a rough guide for measuring hedge effectiveness, and discusses a number of practical issues in employing index-based hedging solutions. Finally, Appendix I sets forth the technical details of the models, and Appendix II lists the Excel VBA codes that have been used for the various computations in Phase 2.

2   Mortality Patterns of Different Groups

In Phase 2, there are four sources of mortality data, in particular the Continuous Mortality Investigation (CMI), Office for National Statistics (ONS), Mercer Australia, and Human Mortality Database (HMD). The first three datasets are used to represent the underlying experience of the pension or annuity portfolio to be hedged, and the last is used as the reference population of the longevity hedging instrument. The following provides a description of the historical mortality levels and improvements of different groups in each dataset. As in Phase 1, the focus is on the age range of 60 to 89. In general, there has been a continual decline in mortality over time at all age groups, and the smaller the data size, the more volatile the experience.

2.1   CMI dataset

The dataset provided by the CMI comes from the self-administered pension scheme (SAPS) mortality investigation on UK pensioners. It includes the number of deaths and central exposed-to-risk, and covers ages 20 to 100+ and years 2000 to 2014. It is split by sex, industry class, pension amount, and retirement type.

The plots of the logit mortality rates (i.e. $\ln\left(q/(1-q)\right)$; see Subsections 3.1 and 3.2) for different groups are given in Figure 2.1 below. For demonstration purposes in this section, some pension groups smaller in size are combined such that their mortality levels decrease generally over time with tractable variability and that there are distinct differences in mortality levels between the groups. For male pensioners under normal retirement, the solid line refers to the pension range of £1 to £8,500 p.a., the dashed line refers to £8,500+ p.a. (for technology, the split is by £4,500 p.a. instead), and the dotted lines represent their potential underlying linear trends. For female pensioners under normal retirement, because of the smaller data sizes, all the pension groups are aggregated (except that there is a split by £1,500 p.a. for local authority). For ill-health retirement, some industries are further grouped together due to insufficient exposures.

There are a few major observations for pensioners under normal retirement. First, the higher pension groups generally have lower mortality and usually have more volatile experience. But the differences in (logit) mortality levels between the two pension groups have a tendency to reduce over age. Second, while the plots are constructed for each industry based on the data given, rather small differences in mortality levels between different industries can be seen, especially at older ages. Moreover, the mortality levels of female pensioners in aggregate are fairly comparable to those of the higher pension male groups, whereas the overall differences in mortality levels between females and males usually become smaller at older ages. Similar patterns are observed for ill-health pensioners, and as expected, ill-health mortality is generally higher than normal mortality.

6

Figure 2.1   Logit mortality rates of UK pensioners from 2000 to 2014

Basic materials (normal retirement; females (top row) and males (bottom row))
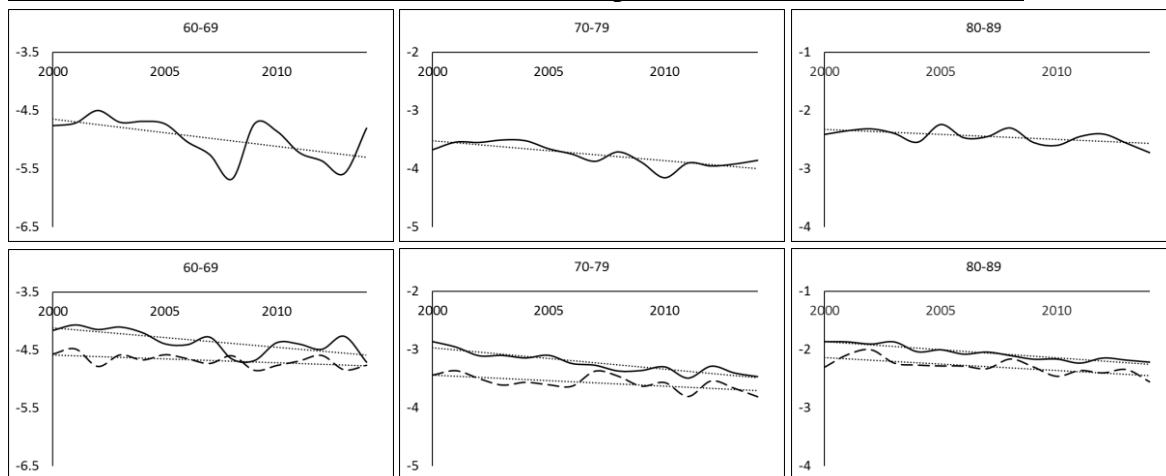
Industrials (normal retirement; females (top row) and males (bottom row))
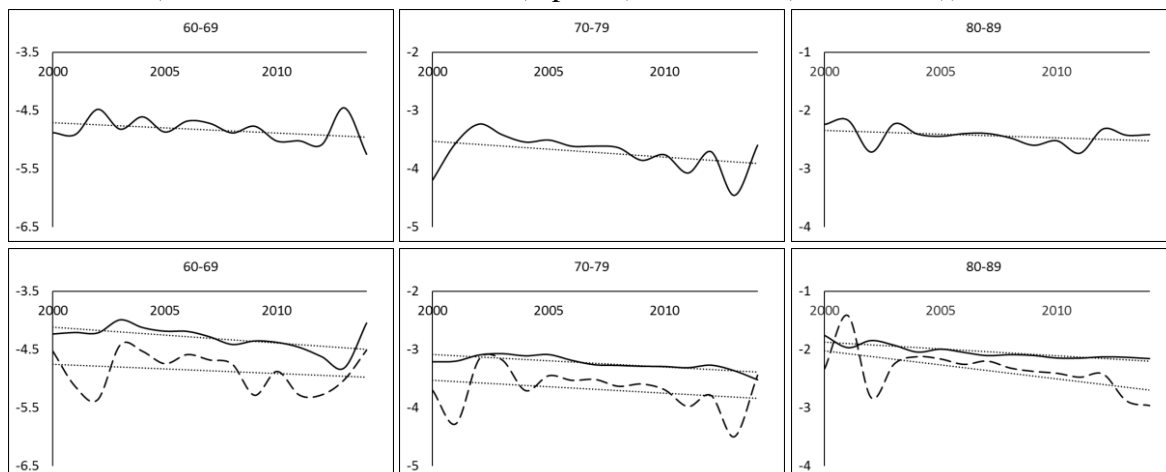
Consumer goods (normal retirement; females (top row) and males (bottom row))

Figure 2.1   Continued

Consumer services (normal retirement; females (top row) and males (bottom row))

Utilities (normal retirement; females (top row) and males (bottom row))
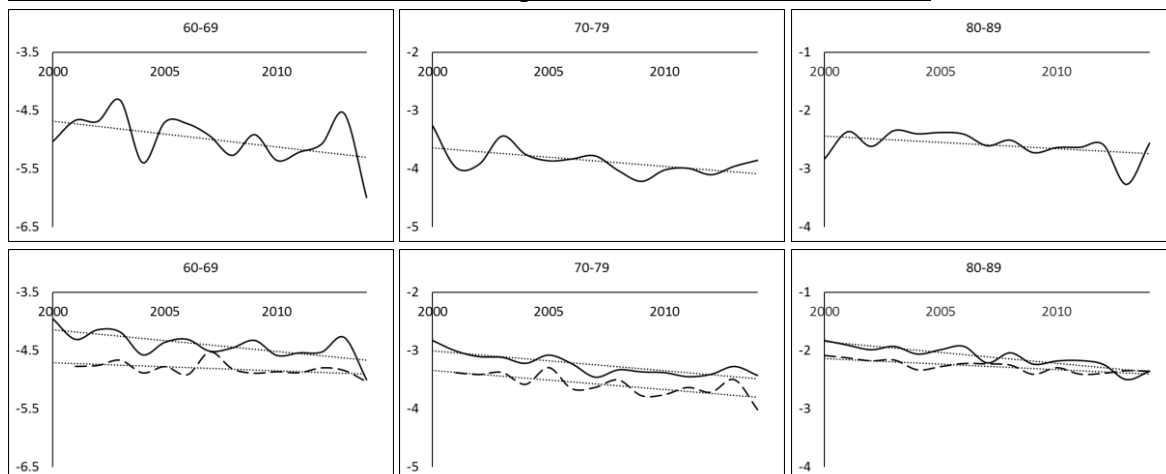
Financials (normal retirement; females (top row) and males (bottom row))
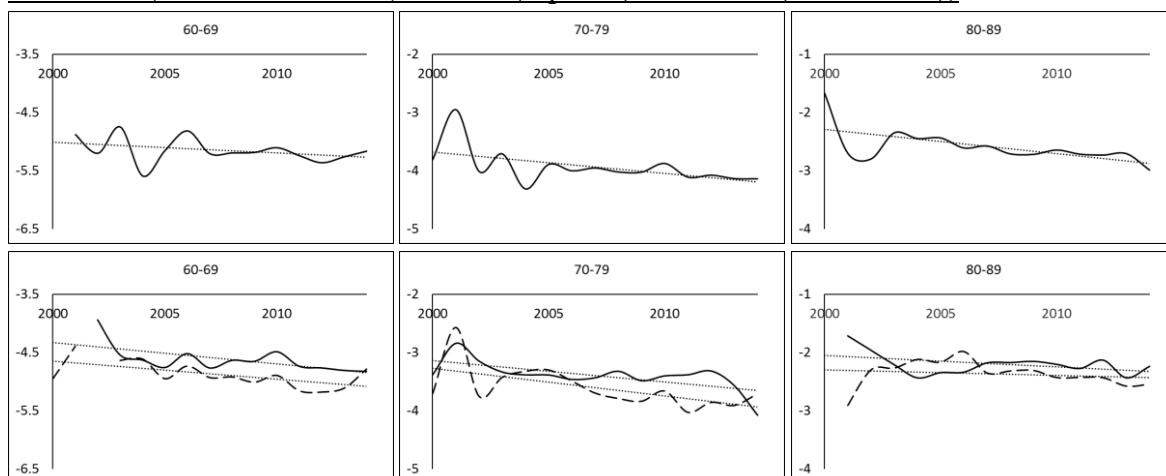
Figure 2.1 Continued

Technology (normal retirement; females (top row) and males (bottom row))
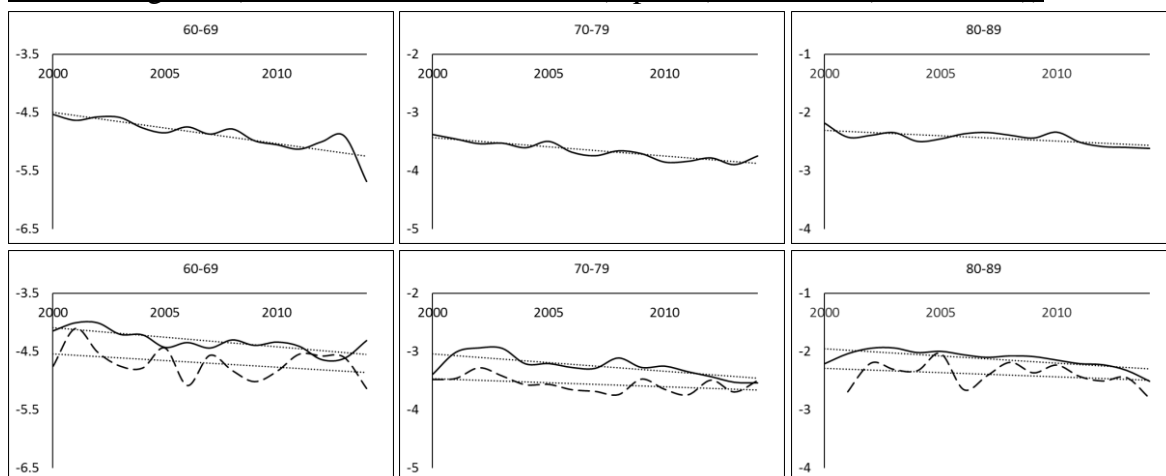

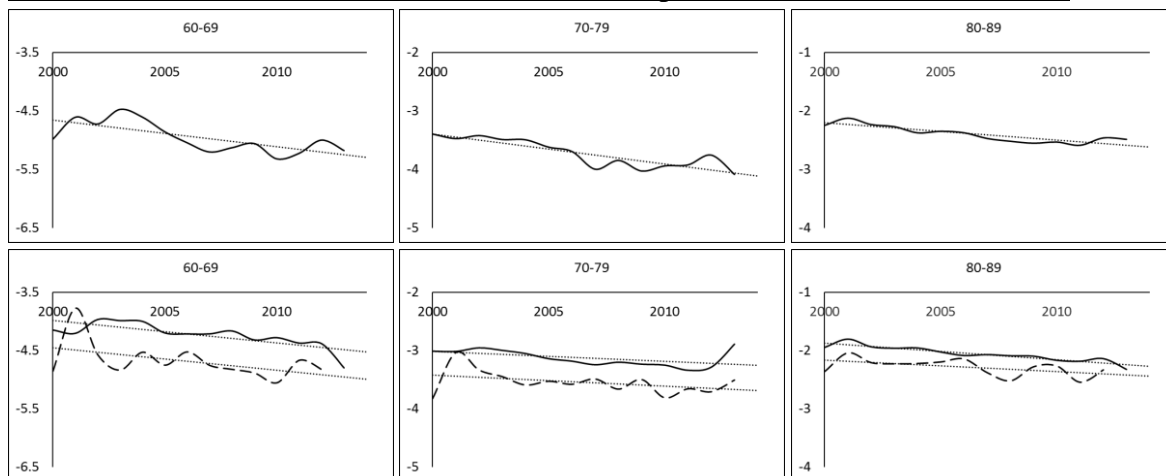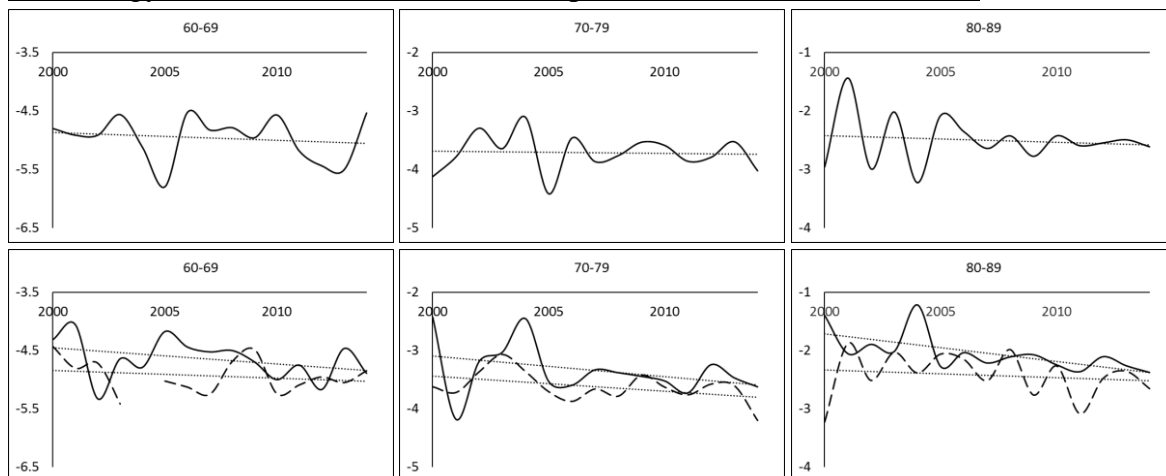Local authority (normal retirement; females (top row) and males (bottom row))
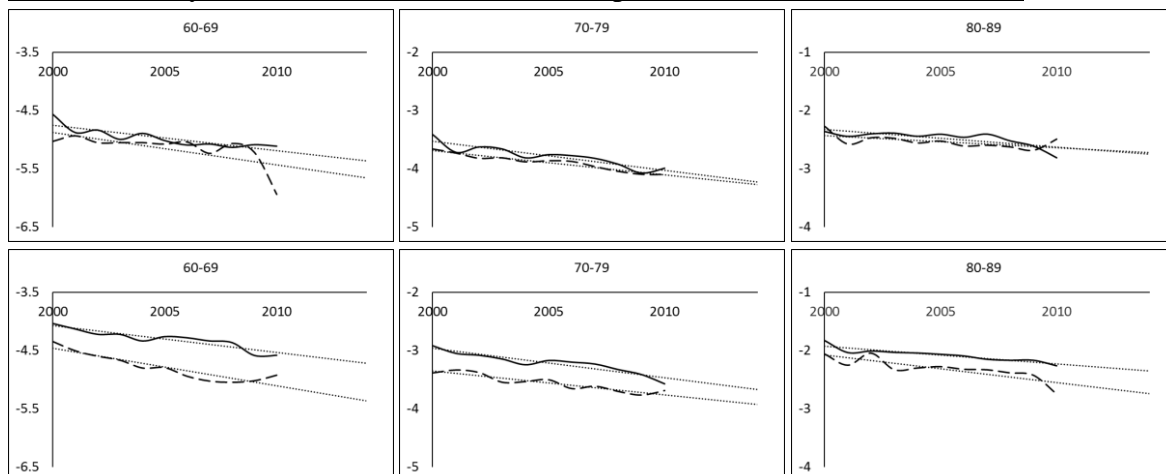
Basic materials and industrials (ill-health retirement; females (top row) and males (bottom row))

Figure 2.1   Continued

Consumer goods and consumer services (ill-health retirement; females (top row) and males (bottom row))

Utilities, financials, and technology (ill-health retirement; females (top row) and males (bottom row))

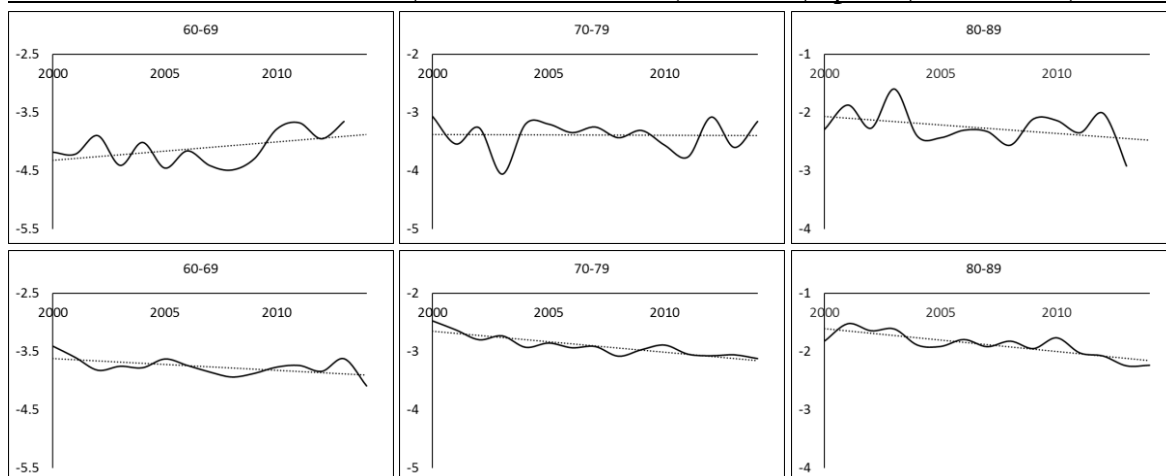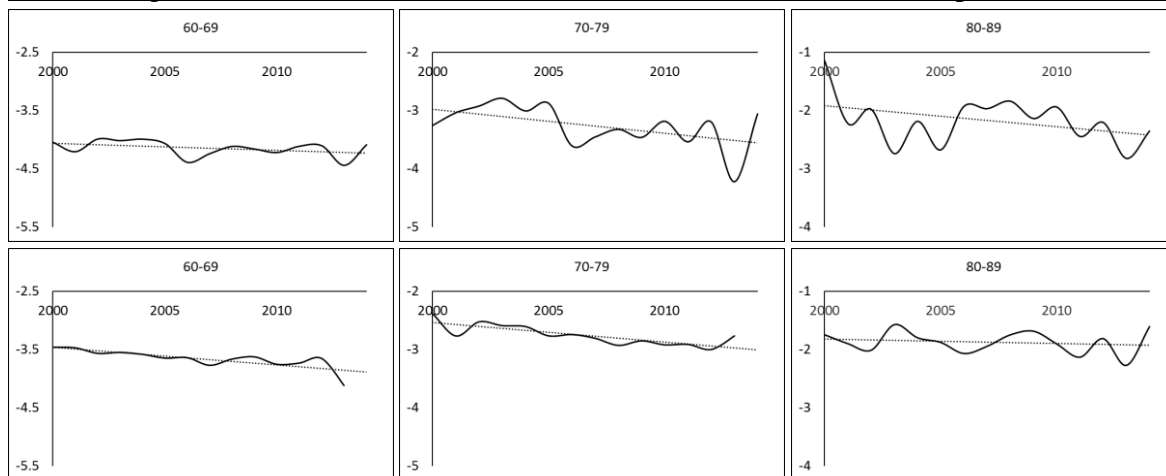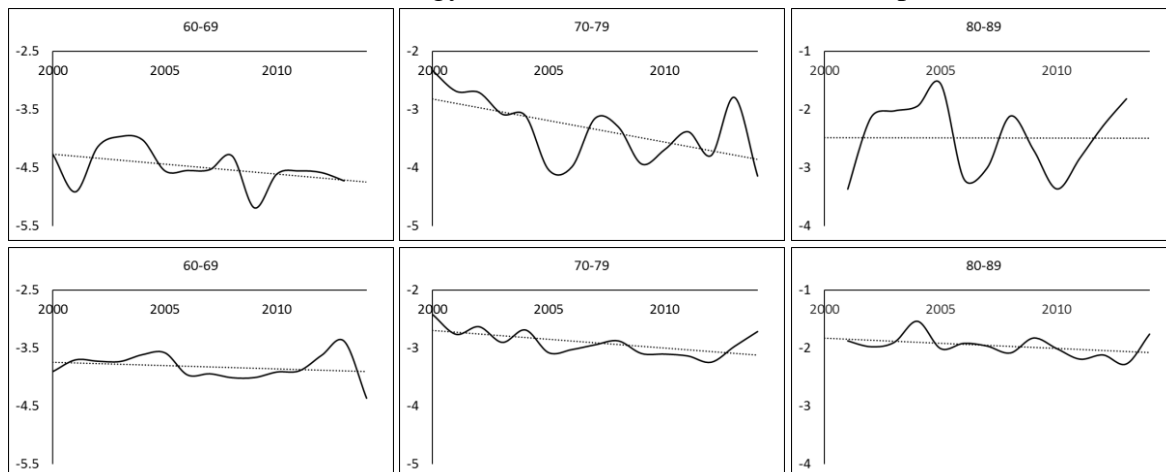Local authority (ill-health retirement; females (top row) and males (bottom row))
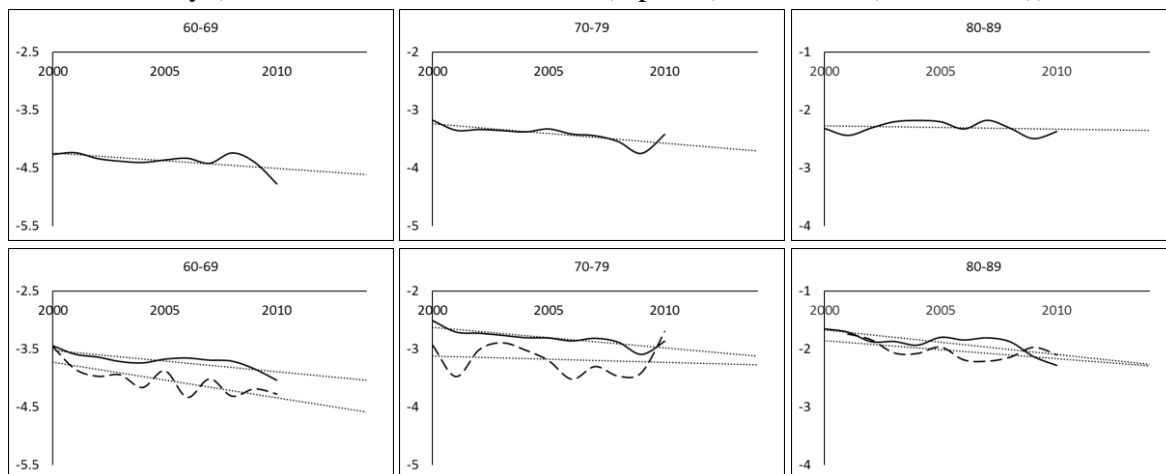

10

Table 2.1 provides the average mortality rate and the average annual rate of improvement in mortality rate for each group over the period. As above, some figures are given separately for the lower (left figure in a cell) and higher (right figure in a cell) pension groups. Again, for pensioners under normal retirement, the higher pension groups have lower mortality levels in general, and the differences in average mortality rates between different industries are small, though the pensioners in financials seem to experience slightly lower mortality levels compared to the others. The average mortality rates of female pensioners in aggregate are a little lower than those of the higher pension male groups. Regarding the improvement rates, there are a few things to note. First, the improvements, mostly ranging from about 1% to 5% p.a., tend to reduce over age, though the patterns are less clear for those groups smaller in size. Second, the lower pension groups, with higher mortality rates, look more likely to experience greater improvements in mortality, though not in all cases. In addition, the differences in improvement rates between different industries or between both sexes appear to be largely randomly scattered, though female pensioners in industrials and technology seem to experience smaller improvements relative to those in the other groups, while pensioners in local authority enjoy greater improvements. For ill-health pensioners, mortality levels are higher and some similar patterns as above can be seen.

Table 2.1   **Average mortality levels and improvements of UK pensioners from 2000 to 2014**

Group \ Age	Females		Males	
	Average Level	Improvement (p.a.)	Average Level	Improvement (p.a.)
Basic materials	(normal retirement)			
60-69	0.0072	4.6%	0.0129 / 0.0093	3.3% / 1.3%
70-79	0.0231	3.3%	0.0385 / 0.0275	3.5% / 1.8%
80-89	0.0801	1.6%	0.1136 / 0.0921	2.5% / 2.0%
Industrials	(normal retirement)			
60-69	0.0081	1.8%	0.0136 / 0.0080	2.7% / 1.6%
70-79	0.0246	2.6%	0.0380 / 0.0259	2.1% / 2.1%
80-89	0.0815	1.1%	0.1155 / 0.0905	2.1% / 4.3%
Consumer goods	(normal retirement)			
60-69	0.0078	5.4%	0.0133 / 0.0093	3.2% / 2.2%
70-79	0.0254	3.1%	0.0378 / 0.0278	2.8% / 1.4%
80-89	0.0809	1.7%	0.1076 / 0.0843	2.2% / 1.3%

Table 2.1 Continued

Group \ Age	Females		Males	
	Average Level	Improvement (p.a.)	Average Level	Improvement (p.a.)
Consumer services	(normal retirement)			
60-69	0.0072	4.5%	0.0145 / 0.0096	3.8% / 3.8%
70-79	0.0241	5.0%	0.0422 / 0.0285	1.6% / 1.8%
80-89	0.0840	2.7%	0.1139 / 0.0929	2.5% / 1.7%
Utilities	(normal retirement)			
60-69	0.0072	4.4%	0.0124 / 0.0081	3.7% / 1.4%
70-79	0.0211	3.1%	0.0380 / 0.0273	3.3% / 3.3%
80-89	0.0712	2.0%	0.1096 / 0.0936	3.3% / 1.7%
Financials	(normal retirement)			
60-69	0.0059	1.8%	0.0102 / 0.0078	3.6% / 3.1%
70-79	0.0202	3.6%	0.0332 / 0.0283	3.6% / 4.6%
80-89	0.0728	3.8%	0.1010 / 0.0871	1.7% / 0.8%
Technology	(normal retirement)			
60-69	0.0074	1.3%	0.0100 / 0.0074	2.8% / 1.3%
70-79	0.0248	0.4%	0.0374 / 0.0267	3.4% / 2.5%
80-89	0.0813	0.8%	0.1185 / 0.0854	4.1% / 1.2%
Local authority	(normal retirement)			
60-69	0.0070 / 0.0059	4.4% / 5.5%	0.0135 / 0.0085	4.6% / 6.4%
70-79	0.0227 / 0.0201	4.9% / 4.1%	0.0391 / 0.0280	4.8% / 4.0%
80-89	0.0778 / 0.0738	2.8% / 1.9%	0.1117 / 0.0906	2.7% / 4.3%
Basic materials, industrials	(ill-health retirement)			
60-69	0.0167	-3.1%	0.0229	2.0%
70-79	0.0336	0.1%	0.0528	3.4%
80-89	0.0977	2.6%	0.1339	3.4%
Consumer goods, consumer services	(ill-health retirement)			
60-69	0.0157	1.2%	0.0254	3.0%
70-79	0.0386	4.0%	0.0604	3.1%
80-89	0.1085	3.1%	0.1344	0.7%

Table 2.1   Continued

Group \ Age	Females		Males	
	Average Level	**Improvement (p.a.)**	**Average Level**	**Improvement (p.a.)**
Utilities, financials, technology	(ill-health retirement)			
60-69	0.0116	3.4%	0.0218	1.1%
70-79	0.0389	7.1%	0.0527	2.9%
80-89	0.0871	0.0%	0.1246	1.6%
Local authority	(ill-health retirement)			
60-69	0.0125	2.6%	0.0242 / 0.0179	3.6% / 6.0%
70-79	0.0324	3.2%	0.0576 / 0.0415	3.4% / 1.1%
80-89	0.0912	0.5%	0.1329 / 0.1168	3.7% / 2.7%

2.2   ONS dataset

The dataset offered by the ONS contains the number of deaths and mid-year population size, sorted by age, calendar year, sex, and index of multiple deprivation (IMD). The data cover ages 0 to 115 and years 2001 to 2015. The IMD combines information from seven domains to produce an overall relative measure of deprivation, including income, employment, education, health, crime, barriers to housing and services, and living environment.

The logit mortality rates of England IMD quintile groups are shown in Figure 2.2. The five different lines with progressively lighter shades in each graph, from top (black) to bottom (light grey), represent the most deprived areas to the least deprived areas consecutively. Clearly, for both sexes and all age groups, mortality increases with the level of deprivation. Furthermore, the differences in (logit) mortality levels between the quintile deprivation groups and between females and males tend to reduce over age. As the data exposures are large, the declining trends over time are quite steady at all age groups. Similar patterns can be observed when the data are further divided into decile groups instead, though the differences are then less clear-cut between the groups and so this level of division is not adopted here.

Figure 2.2   Logit mortality rates of England IMD quintile groups from 2001 to 2015

Most deprived to least deprived (top to bottom) areas (females (top row) and males (bottom row))
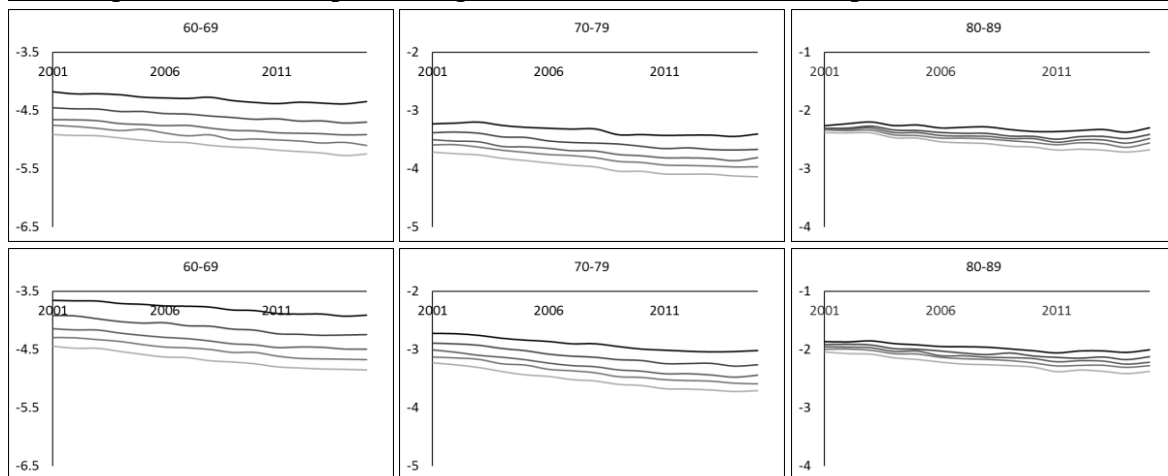
The average mortality rate and the average annual rate of improvement in mortality rate for each group during the period are given in Table 2.2. It is very clear that less deprived areas have not only lower mortality but also higher improvement rates than more deprived areas. Male improvement rates, ranging from 1.3% to 3.6% p.a., are higher than female improvement rates, which range from 0.9% to 3.2% p.a., for all areas and age groups. It is notable that while ages 80 to 89 have the lowest improvement rates, those aged 70 to 79 have the highest, for both sexes and all areas. All these patterns are very consistent across different groups, in line with the steady trends observed above.

One interesting contrast to the previous results is that while less deprived (wealthier) areas clearly have higher improvement rates as shown here, lower pension (less wealthy) groups for a number of industries seem more likely to have greater improvements as in Section 2.1. There may be a few causes behind this difference. First, the IMD in this section is based on a combination of seven domains compared to just using the pension amount in the previous section, so the two sets of results may refer to individuals with very different demographics. Second, the previous results are industry-specific, in which some cases have actually shown an opposite situation. Moreover, the data exposures of each industry are rather small and the corresponding results are unavoidably more volatile than those presented here. Another interesting difference is that male improvement rates are higher than female improvement rates based on the ONS data but no parallel results are produced using the previous CMI data. More investigation is possible only if more detailed individual information is made available by relevant organisations.

Table 2.2   Average mortality levels and improvements of England IMD quintile groups from 2001 to 2015

Group \ Age	Females		Males	
	Average Level	Improvement (p.a.)	Average Level	Improvement (p.a.)
Most deprived areas				
60-69	0.0134	1.4%	0.0221	2.1%
70-79	0.0343	1.7%	0.0522	2.4%
80-89	0.0914	0.9%	0.1237	1.3%
Second most deprived areas				
60-69	0.0101	1.9%	0.0162	2.6%
70-79	0.0282	2.3%	0.0428	2.9%
80-89	0.0839	1.2%	0.1139	1.7%
Third most deprived areas				
60-69	0.0082	2.1%	0.0129	2.8%
70-79	0.0243	2.6%	0.0367	3.3%
80-89	0.0805	1.6%	0.1082	1.9%
Fourth most deprived areas				
60-69	0.0072	2.4%	0.0111	3.0%
70-79	0.0219	3.0%	0.0331	3.4%
80-89	0.0773	1.9%	0.1032	2.2%
Least deprived areas				
60-69	0.0061	2.7%	0.0093	3.2%
70-79	0.0190	3.2%	0.0291	3.6%
80-89	0.0719	2.4%	0.0963	2.4%

2.3   Mercer dataset

Mercer Australia maintains a database that collects information about pensioners from public sector occupational superannuation funds to facilitate research on post-retirement mortality experience of working Australians. Since almost all superannuation pensions provided in Australia arise in the public sector, this dataset is uniquely placed to study funded retirees' mortality. The data have been collected from the major Australian public sector schemes covering the period from 2002 onwards. There are over three and a half million years of pensioners exposures to 2012. Comprehensive checks have been conducted on the data, including reconciliation with older data and correction for errors and omissions. In particular, late reported deaths and exits are allocated to the correct period, leading to an inherently reliable set of data, with a direct link between exposed-to-risk and recorded deaths.

The logit mortality rates of different types of pensioners for a number of states or groups in Australia and New Zealand (including the two most populous states in Australia, New South Wales and Victoria, and those operating Australia-wide noted as Commonwealth) are plotted in Figure 2.3 for both sexes. The solid line represents the observed logit mortality rates and the dotted line refers to their potential underlying linear trend. For pensioners under retirement, the differences in (logit) mortality levels between different states in Australia are quite small, while New Zealand tends to have higher mortality than Australia, especially at ages 60 to 69. Moreover, the differences in mortality levels between females and males become smaller at older ages in New Zealand, but this pattern over age is less obvious in Australia. In general, female pensioners have more volatile experience than male pensioners, due at least in part to the smaller portfolio sizes. Compared with female pensioners under retirement, (female) spouse pensioners generally have higher mortality levels; and as expected, mortality levels are significantly higher for invalidity pensioners.

Figure 2.3   Logit mortality rates of Australian and New Zealand pensioners from 2002 to 2012

New South Wales, Australia (retirement; females (top row) and males (bottom row))
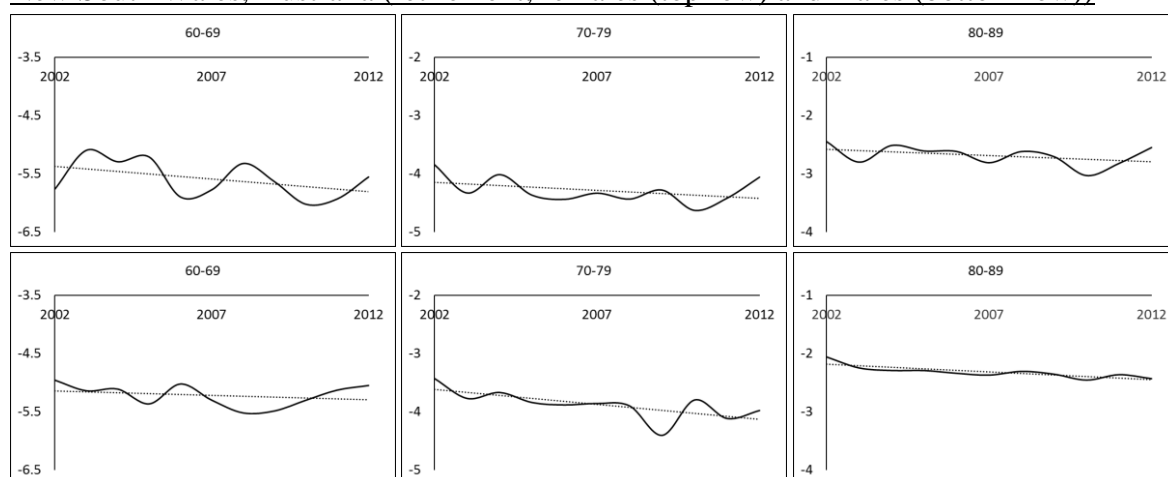
Figure 2.3   Continued

Commonwealth, Australia (retirement; females (top row) and males (bottom row))
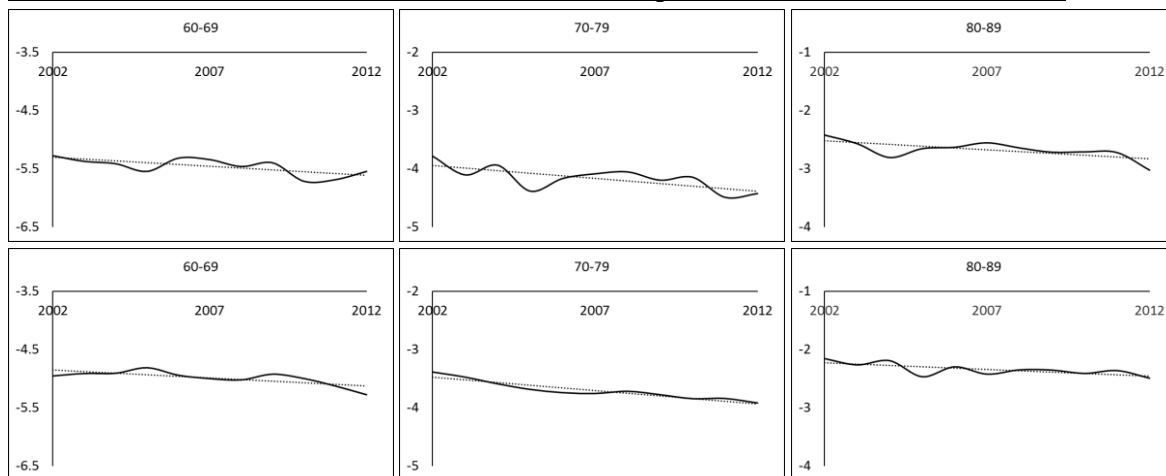
Victoria, Australia (retirement; females (top row) and males (bottom row))
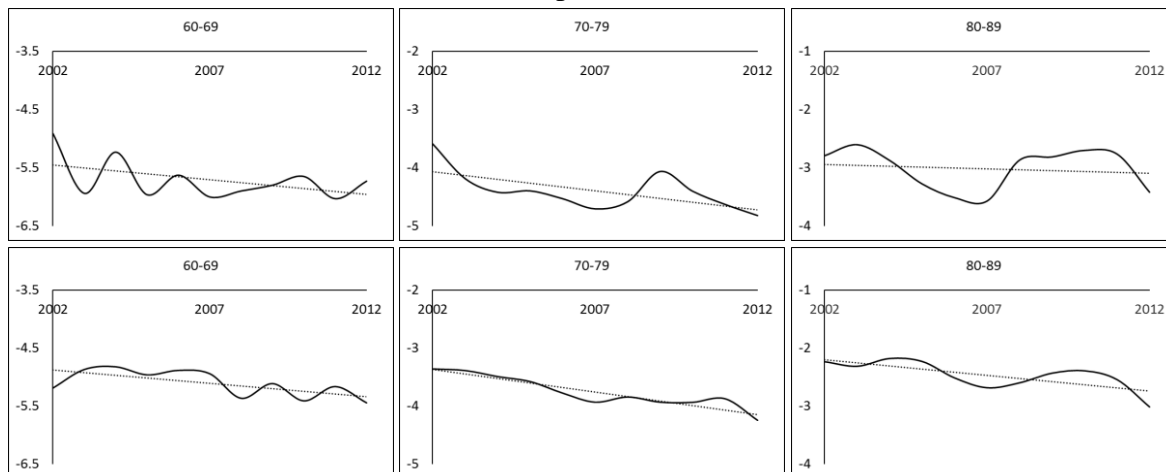
Other states in Australia (retirement; females (top row) and males (bottom row))
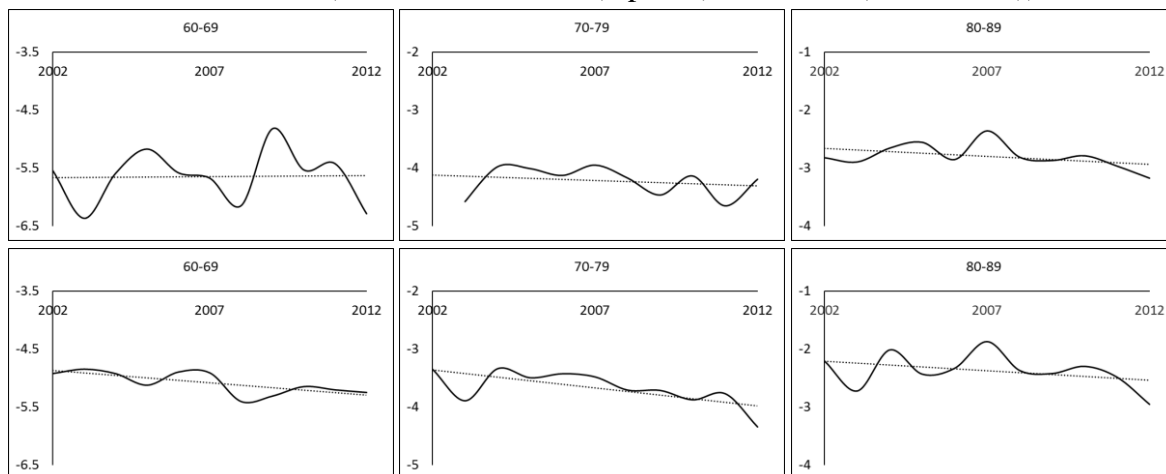
Figure 2.3   Continued

New Zealand (retirement; females (top row) and males (bottom row))
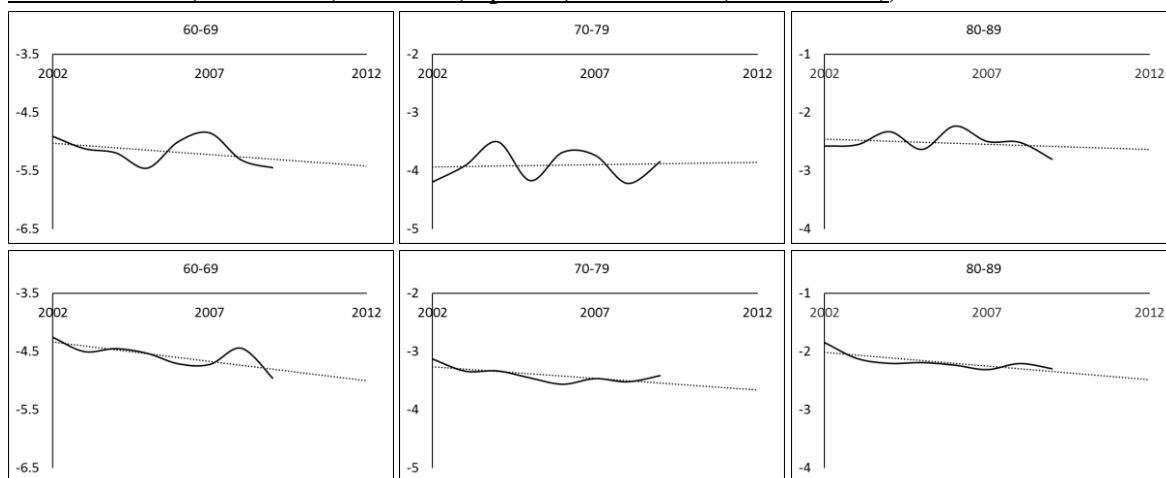
Australia and New Zealand (invalidity; females (top row) and males (bottom row))
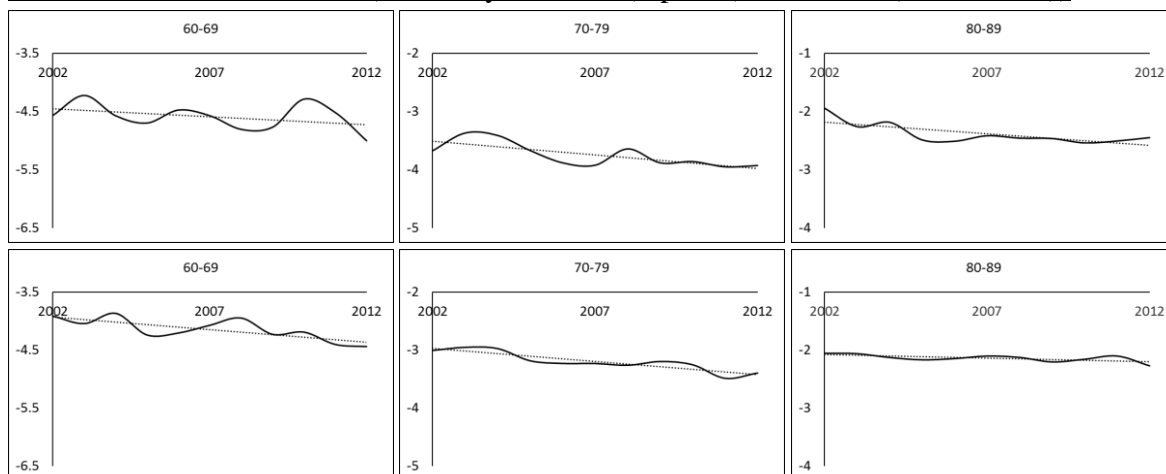
Australia and New Zealand (spouse; females)
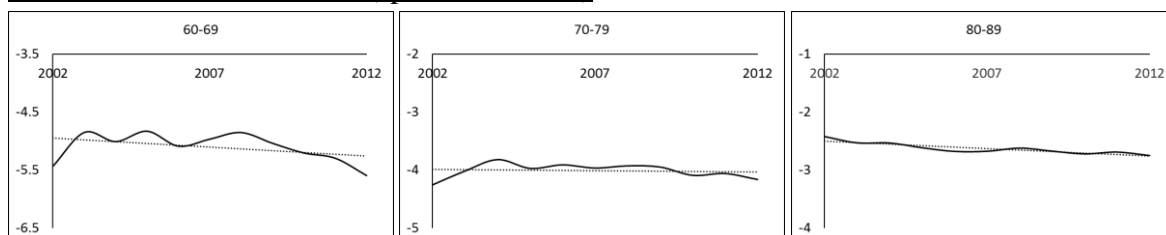
Table 2.3 records the average mortality rate and annual improvement rate for each group during the period. For pensioners under retirement, the differences in average mortality rates are quite small between different states in Australia, while New Zealand pensioners have higher average mortality rates for both sexes at all age groups. Spouse pensioners largely have higher mortality levels compared with female pensioners under retirement. Invalidity pensioners clearly have the highest mortality levels amongst all groups. Furthermore, the

18

improvement rates are mostly in the range of around 1% to 6% p.a. and have a tendency to decrease over age. Many groups show that those aged 70 to 79 have the highest improvement rates. The differences in improvement rates between different groups and between both sexes are not consistent and look quite random, though Victoria in Australia seems to experience greater improvements than the others.

Table 2.3   Average mortality levels and improvements of Australian and New Zealand pensioners from 2002 to 2012

Group \ Age	Females		Males	
	Average Level	Improvement (p.a.)	Average Level	Improvement (p.a.)
New South Wales, Australia	(retirement)			
60-69	0.0039	4.3%	0.0055	1.5%
70-79	0.0139	2.8%	0.0208	5.1%
80-89	0.0644	2.0%	0.0899	2.4%
Commonwealth, Australia	(retirement)			
60-69	0.0043	3.1%	0.0068	2.8%
70-79	0.0156	4.3%	0.0243	4.5%
80-89	0.0650	3.0%	0.0878	2.1%
Victoria, Australia	(retirement)			
60-69	0.0035	5.1%	0.0061	4.5%
70-79	0.0129	6.5%	0.0235	7.6%
80-89	0.0487	1.5%	0.0795	5.0%
Other states in Australia	(retirement)			
60-69	0.0039	-0.3%	0.0063	4.2%
70-79	0.0138	1.9%	0.0257	6.0%
80-89	0.0584	2.6%	0.0879	3.1%
New Zealand	(retirement)			
60-69	0.0058	3.9%	0.0104	6.6%
70-79	0.0203	-0.8%	0.0324	3.9%
80-89	0.0751	1.6%	0.1024	4.2%

Table 2.3   Continued

Group \ Age	Females		Males	
	Average Level	Improvement (p.a.)	Average Level	Improvement (p.a.)
Australia, New Zealand	(invalidity)			
60-69	0.0102	2.7%	0.0158	4.3%
70-79	0.0235	4.5%	0.0396	4.3%
80-89	0.0855	3.6%	0.1054	1.1%
Australia, New Zealand	(spouse)			
60-69	0.0062	3.0%	-	-
70-79	0.0179	0.4%	-	-
80-89	0.0675	2.4%	-	-

2.4   HMD dataset

The HMD provides detailed mortality and population data for many countries. The data of England and Wales, Australia, and New Zealand are collected for the analysis in this report. These data cover ages 0 to 110+ and the period from 1980 onwards. While data for earlier periods are also available, the starting year of 1980 is chosen here because structural changes in mortality improvement have occurred in the past for these countries (e.g. Booth et al. 2002, Renshaw and Haberman 2003, Li 2010) and it would be more relevant and convenient to use recent data when projecting future mortality levels.

Figure 2.4 plots the logit mortality rates of the three populations from 2000 for both sexes. The declining mortality trends over time are steady at all age groups and do not show much variability, as the population exposures are very large. The differences in (logit) mortality levels between females and males also tend to decrease over age. The corresponding average mortality rate and annual rate of mortality improvement are provided in Table 2.4. Male improvement rates have a range of 2.0% to 3.5% p.a., and are higher than female improvement rates, ranging from 1.3% to 2.9% p.a., for all the three populations and age groups. The improvement rates are lowest at ages 80 to 89 in all cases, and for England and Wales and Australia, those aged 70 to 79 experience the highest improvement rates.

Compared to UK pensioners under normal retirement in Section 2.1, the English and Welsh female population appears to have slightly higher mortality levels and lower improvement rates than female pensioners in most cases. On the other hand, the male population has mortality levels and improvement rates roughly comparable to those of the lower pension male groups. When compared with Australian and New Zealand pensioners under retirement in Section 2.3, the Australian and New Zealand populations largely experience higher mortality levels and lower improvement rates for both sexes.

Figure 2.4   Logit mortality rates of England and Wales, Australia, and New Zealand from 2000 to 2014

England and Wales (females (top row) and males (bottom row))
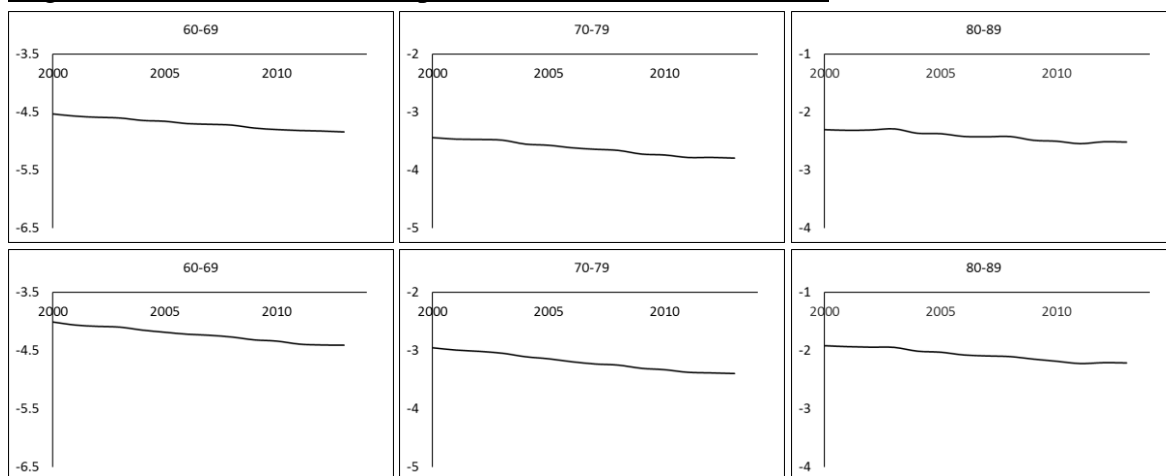
Australia (females (top row) and males (bottom row))
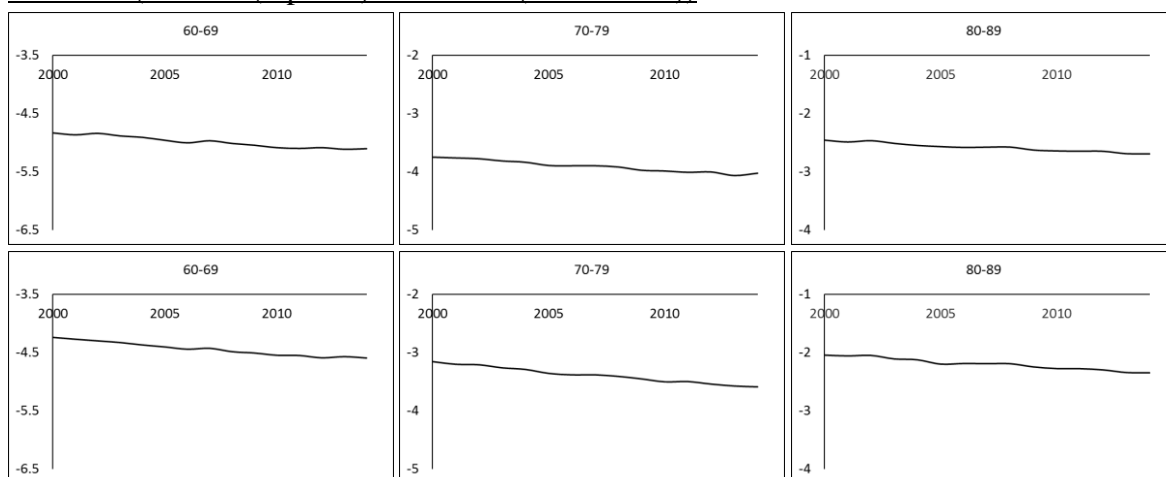
New Zealand (females (top row) and males (bottom row))
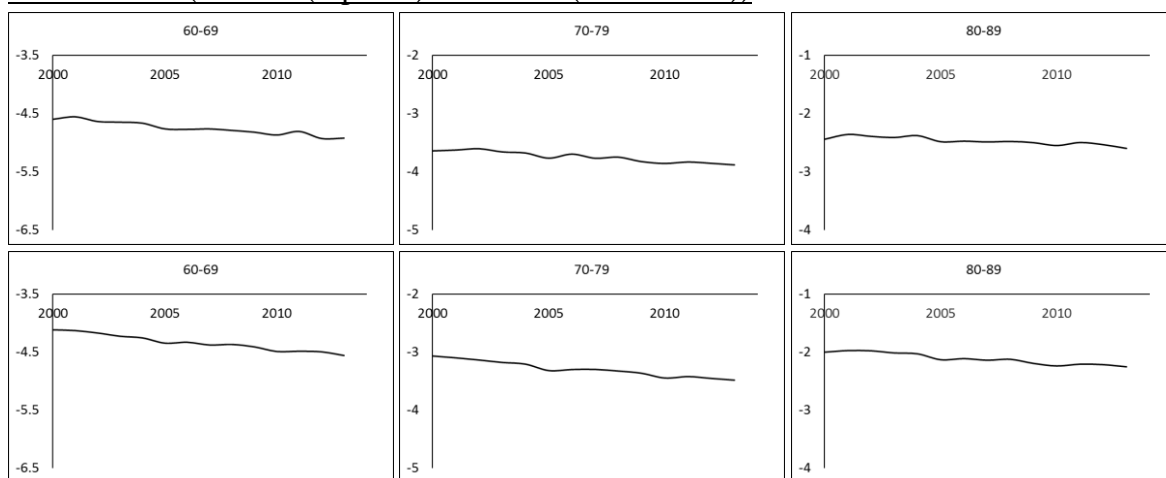
Table 2.4   Average mortality levels and improvements of England and Wales, Australia, and New Zealand from 2000 to 2014

Group \ Age	Females		Males	
	Average Level	Improvement (p.a.)	Average Level	Improvement (p.a.)
England & Wales				
60-69	0.0091	2.4%	0.0145	3.1%
70-79	0.0261	2.9%	0.0398	3.5%
80-89	0.0826	1.9%	0.1119	2.3%
Australia				
60-69	0.0068	2.2%	0.0117	2.6%
70-79	0.0198	2.2%	0.0329	3.0%
80-89	0.0707	1.6%	0.1000	2.0%
New Zealand				
60-69	0.0086	2.7%	0.0130	3.3%
70-79	0.0232	2.1%	0.0361	3.1%
80-89	0.0780	1.3%	0.1079	2.1%

2.5   Major patterns

All the observations so far can be briefly summarised as follows:
- Mortality rates within different age groups decline over time in general, but the movements are much more volatile for the pensioners than for the whole populations, mainly due to the smaller sizes of the former.
- Pensioners under normal retirement tend to have lower mortality levels and higher improvement rates than the whole populations.
- In the UK, financially wealthier groups generally have lower mortality levels. People living in less deprived areas have higher improvement rates.
- Improvement rates mostly range from about 1% to 5% p.a. They are usually lowest at ages 80 to 89. In the UK and Australia, for the population data and some pensioner data, those aged 70 to 79 have the highest improvement rates.
- Ill-health pensioners have higher mortality levels than normal pensioners and also the whole populations.

Evidently there are distinct differences between specific group experience and population experience, in terms of the level, rate of change, and variability of mortality rates. These observations point to the need of using an approach as in Phase 1, which caters for population experience in the first component of the model structure and then takes into account the differences between the pensioners and the overall population in the second component.

3 Modelling Longevity Basis Risk

Longevity basis risk arises from the potential mismatch between the longevity hedging instrument (based on a reference population) and the pension or annuity portfolio (with the underlying book population) being hedged. There are three main sources of longevity basis risk, including demographic basis risk (demographic or socioeconomic differences), sampling basis risk (random outcomes of individual lives), and structural basis risk (differences in payoff structures). Phase 1 has focused on measuring demographic basis risk, in which a decision tree framework (Figure 3.1) has been built for providing a practical guide on how to select an appropriate two-population mortality model for the reference and book populations. The major choices include the M7-M5 model, the CAE+Cohorts model, and the characterisation approach. The selection is based on the data size, portfolio compositions, inter-age relationships, and cohort effect.

Figure 3.1   Decision tree framework for selecting a two-population mortality model (extracted from Phase 1 report)

3.1 M7-M5 model

The M7-M5 model is a two-population extension of the CBD model (Cairns et al. 2006):

$$\text{logit } q_{x,t}^R = \ln \frac{q_{x,t}^R}{1-q_{x,t}^R} = \kappa_{t,1}^R + (x-\bar{x})\kappa_{t,2}^R + ((x-\bar{x})^2 - \sigma_x^2)\kappa_{t,3}^R + \gamma_{t-x}^R \qquad \text{(reference population)}$$

$$\text{logit } q_{x,t}^B - \text{logit } q_{x,t}^R = \kappa_{t,1}^B + (x-\bar{x})\kappa_{t,2}^B \qquad \text{(book vs reference population)}$$

For every year t, the first equation can be seen as a mortality curve over age x for the reference population. The notation $q_{x,t}^R$ is the mortality rate, $\kappa_{t,1}^R$, $\kappa_{t,2}^R$, and $\kappa_{t,3}^R$ refer to the level, slope, and curvature of the mortality curve respectively, and γ_{t-x}^R represents the cohort effect of those born in year $t-x$. Put simply, for each year, the logit mortality rates across age can be expressed in terms of the level, slope, curvature, and underlying cohorts. Then, in the second equation, the difference in the logit mortality rate between the book and reference populations ($\text{logit } q_{x,t}^B - \text{logit } q_{x,t}^R$) is modelled as a linear combination of $\kappa_{t,1}^B$ and $\kappa_{t,2}^B$, which are another two parameters for capturing the differences between the two populations. More technical details are provided in Appendix I.

3.2 CAE+Cohorts model

The CAE+Cohorts model is a two-population extension of the Lee and Carter (1992) model and the Kleinow (2015) model:

$$\text{logit } q_{x,t}^R = \ln \frac{q_{x,t}^R}{1-q_{x,t}^R} = \alpha_x^R + \beta_x^R \kappa_t^R + \gamma_{t-x}^R \qquad \text{(reference population)}$$

$$\text{logit } q_{x,t}^B - \text{logit } q_{x,t}^R = \alpha_x^B + \beta_x^R \kappa_t^B \qquad \text{(book vs reference population)}$$

In the first equation for the reference population, $q_{x,t}^R$ is the mortality rate, α_x^R depicts the mortality schedule over age x, κ_t^R is called the mortality index which captures the overall mortality improvement over time t, with β_x^R as the age-specific sensitivity measure, and γ_{t-x}^R represents the cohort effect of those born in year $t-x$. In the second equation, the difference in the logit mortality rate between the book and reference populations ($\text{logit } q_{x,t}^B - \text{logit } q_{x,t}^R$) is modelled as a linear combination of α_x^B and κ_t^B (again with β_x^R as the sensitivity measure), which are extra parameters catering for the differences between the two populations. Appendix I contains further technical details of this model.

3.3 Characterisation approach

When the past data size is too small or the period of reliable history is too short, the characterisation approach can be used to mimic the book population. Its rationale is that if the data size is too small, the resulting fitted model would induce higher variability in its

simulations, which may then exaggerate the level of uncertainty of the portfolio and so overestimate demographic basis risk. As shown in Figure 3.2, the major steps are to: (1) divide the book population into (say) three groups (B_1, B_2, B_3) based on distinct characteristics such as pension amounts and postcodes; (2) find an alternative data source with similar characteristics but a larger size and a longer and more reliable history for each of these characterising groups (C_1, C_2, C_3); (3) fit the M7-M5 model or CAE+Cohorts model to the alternative data and produce future simulations for each group; (4) use these simulations to generate 'proxy' future scenarios for each group of the actual book population; (5) combine the 'proxy' future scenarios of all groups to form a 'proxy' simulation of the entire book population. In effect, the book data are not modelled directly; instead, some alternative data, which mimic different segments of the book population, are modelled to generate future simulations as a proxy. Since the mapping is unlikely to be perfect and the characterising groups may not capture all the features of the book population, there would be residual basis risk being further introduced into the calculations.

Figure 3.2   Characterisation approach (extracted from Phase 1 report)
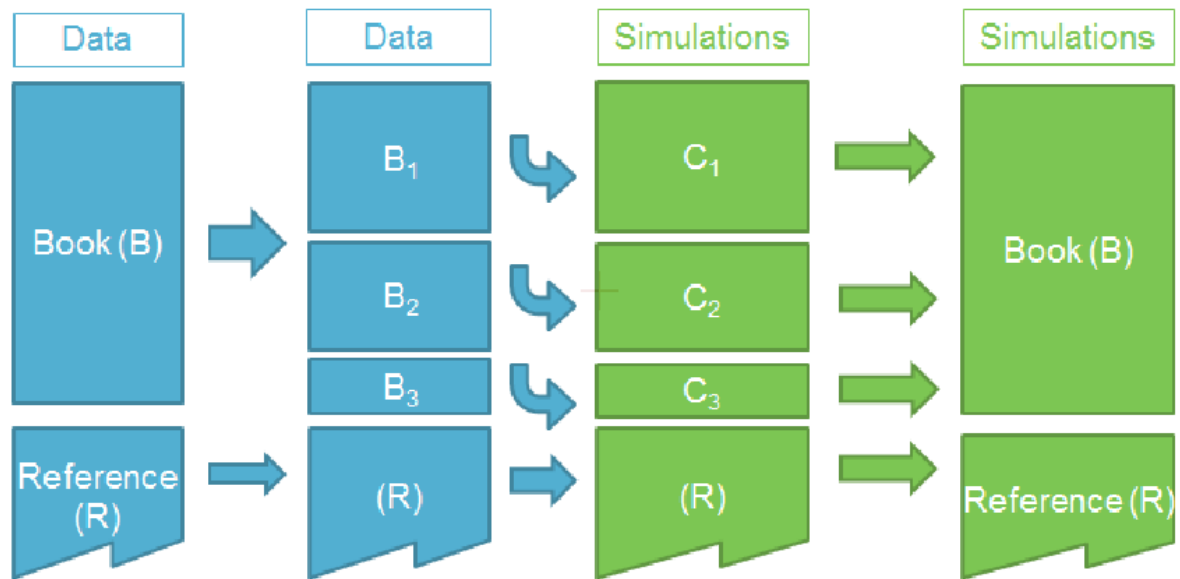
3.4   Future simulations

As per Phase 1, for the M7-M5 model, the time series of $\kappa_{t,1}^{R}$, $\kappa_{t,2}^{R}$, and $\kappa_{t,3}^{R}$ are modelled as a multivariate random walk with drift (MRWD). The time series of γ_c^{R} is modelled as an autoregressive integrated moving average process, ARIMA(1,1,0). The time series of $\kappa_{t,1}^{B}$ and $\kappa_{t,2}^{B}$ are modelled as a vector autoregressive process of order one, VAR(1). For the CAE+Cohorts model, the time series of κ_t^{R} is modelled as a random walk with drift (RWD).

The time series of γ_c^R is modelled as an ARIMA(1,1,0). The time series of κ_t^B is modelled as an autoregressive process of order one, AR(1). More details and discussions of time series modelling are set forth in Section 6.

According to the QIS5 Technical Specifications (CEIOPS 2010), assessing uncertainty of future outcomes includes process error (or process risk), parameter error (or parameter uncertainty), and model error (or model uncertainty). In order to incorporate both process error (variability in the time series) and parameter error (uncertainty in parameter estimation) in simulating future mortality rates, the method of residuals bootstrapping (Koissi et al. 2006; Li 2014) is adopted here (Figure 3.3). First, the residuals from fitting the M7-M5 model or CAE+Cohorts model to the actual data are resampled with replacement. Then an inverse formula is used to turn these resampled residuals into a pseudo sample of the number of deaths (i.e. pseudo data). The M7-M5 model or CAE+Cohorts model is fitted to this pseudo sample and the corresponding model parameters are computed. The time series processes are then applied to the temporal model parameters of the pseudo sample to simulate their future values. Finally, future mortality rates are generated using all the computed and simulated parameters of the pseudo sample. In effect, this pseudo sample gives rise to one future scenario. The whole process above is repeated to give (say) 5,000 scenarios. This bootstrapping procedure is further explained in Appendix I. Moreover, the extent of model error (uncertainty in the model choice) can be investigated by comparing the results from different models and assumptions.

Figure 3.3   Residuals bootstrapping
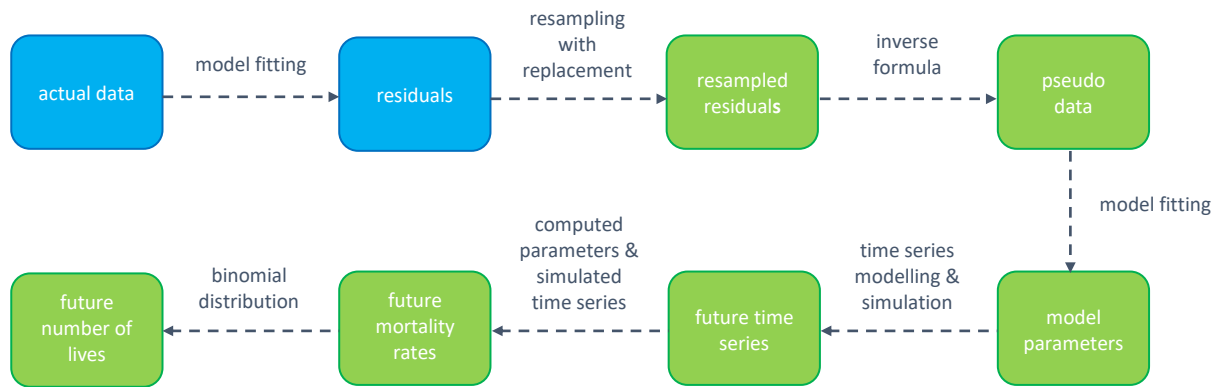
3.5   Sampling risk

Sampling basis risk exists because of the finite sizes of the reference and book populations and the randomness of outcomes of individual lives. Roughly speaking, by the law of large numbers, if the number of lives is infinite, the future outcomes will converge to the true underlying expected values. However, in reality, the number of lives is limited. Even though

one may argue that the number of lives is very large for bigger countries, the size of a pension or annuity portfolio is usually much smaller. So the outcomes of the book population and also the reference population (particularly if its size is small) will deviate randomly from their true underlying expected values and also from each other. To allow for the impact of the portfolio size, the number of survivors in the portfolio over time is simulated as:

$$l^B_{x+1,t+1} \sim \text{Binomial}\left(l^B_{x,t} , \, 1 - q^B_{x,t} \right)$$

The notation $l^B_{x,t}$ is the future number of lives aged x at time t in the book population. The future mortality rate $q^B_{x,t}$ at age x in year t (i.e. from time t to $t+1$) is simulated from residuals bootstrapping as noted in the previous subsection (Figure 3.3). Different starting values of $l^B_{x,t}$ at the valuation date may also be used in turn to study the effect of having different initial portfolio sizes. Note that while in principle the future number of lives in the reference population can be simulated in the same way, in practice this step of using a binomial distribution can be omitted for computation convenience if the reference population has a sizable exposure, such as English and Welsh population.

As discussed in Phase 1, for those portfolios with fewer than 25,000 lives in the past, however, the sampling variability in the observed data may flow through to the time series of the fitted model parameters. This effect may then lead to an overestimation of process error, parameter error, and so demographic basis risk. To alleviate this problem, one can use the characterisation approach in Subsection 3.3, instead of applying the M7-M5 and CAE+Cohorts models directly, in order to measure more accurately the true underlying level of process error and parameter error.

3.6 Structural risk

Structural basis risk arises from the differences in payoff structures between the hedging instrument and the portfolio being hedged. Both the timing and amounts of future cash flows of the hedging instrument will often be different to those of the portfolio. A potentially feasible approach to mitigate this problem is the use of numerical optimisation. In the process, specific objectives are first determined, and numerical algorithms are then adopted to find the most optimal strategy to achieve those objectives. In the context of hedging longevity risk, the objective can be set as minimising longevity risk and reducing this structural basis risk, and numerical optimisation can be used to find the optimal position of the hedging instrument, such that the objective is achieved under the simulated environment from above. Given the simulated future mortality rates and numbers of survivors, this calibration procedure is straightforward to apply with a spreadsheet or mathematical software (see Appendices I and II). Another calibration method called key q-duration (KQD) matching is also discussed in Section 5.

3.7   Hedge effectiveness

The effectiveness of an index-based longevity hedge can be expressed in terms of how much longevity risk is transferred away. The hedging positions are structured to achieve a desired level of effectiveness or even the maximum possible level, and the remaining portion can then be regarded as longevity basis risk. Accordingly, in line with Coughlan et al. (2011), we define the level of longevity risk reduction for a particular longevity hedge on a pension or annuity portfolio as:

$$\text{longevity risk reduction} = \left(1 - \frac{\text{risk(hedged)}}{\text{risk(unhedged)}}\right) \times 100\%$$

The notation risk(unhedged) and risk(hedged) are the portfolio's longevity risk before and after taking the hedge. That is, this metric gives the percentage of the portfolio's initial longevity risk that is being hedged away. We consider the common risk measures in the literature (e.g. Dowd and Blake 2006) including the variance, standard deviation, 99.5% value-at-risk (VaR), and 99.5% expected shortfall (conditional VaR). The 99.5% VaR is of particular concern in practice, as Solvency II (Directive 2009/138/EC) stipulates that the Solvency Capital Requirement (SCR) should be determined as the economic capital to ensure that 'ruin occurs no more often than once in every 200 cases' or that there is 'a probability of at least 99.5% to meet the obligations to policyholders and beneficiaries over the following 12 months'. This economic capital must be based on the 'true risk profile', allowing for 'the impact of possible risk-mitigation techniques as well as diversification effects'. Hence it would be of practical interest to measure the levels of risk reduction from different longevity hedges with a view to identify potential opportunities for capital savings.

3.8   Overall framework

Figure 3.4 summarises how all the three components of longevity basis risk can be modelled and addressed using the approaches outlined previously. Further technical details are given in Appendices I and II. Note that there are alternatives to certain parts of this modelling framework. For example, other two-population models discussed in Phase 1 may be used; Bayesian analysis and simulation can be adopted to allow for process error, parameter error, and even model error; more advanced financial hedging strategies may be adapted to the context of hedging longevity risk. However, many of these alternatives are much more technically and computationally demanding, and practitioners need to weigh the pros and cons when choosing a particular method.

Figure 3.4   Modelling longevity basis risk
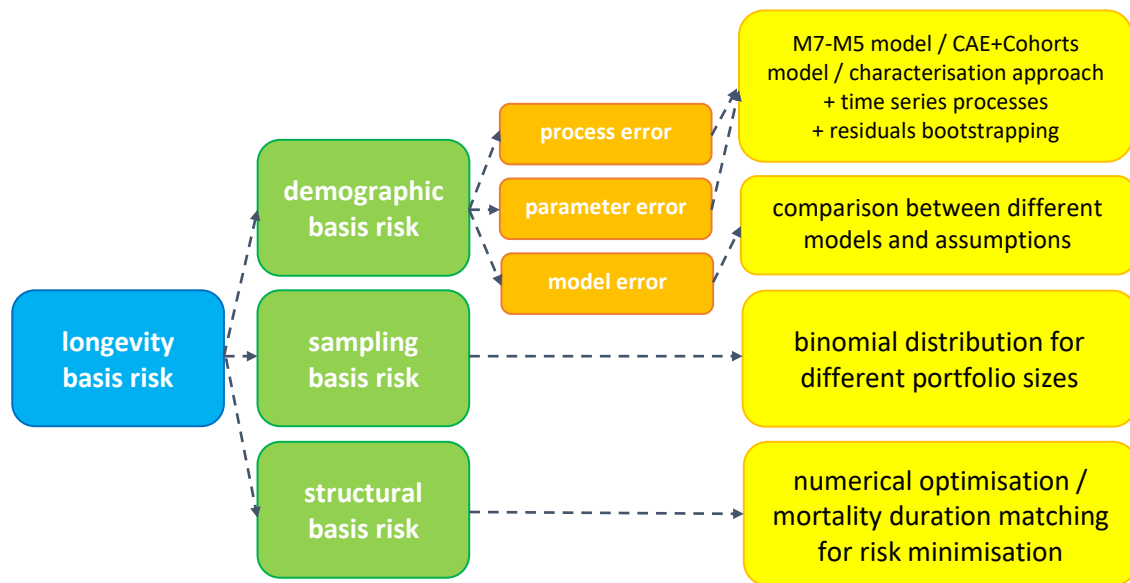
4   Longevity Hedging Scenarios

There are a number of mortality-linked and longevity-linked securities proposed in the literature. Some of these securities can serve as hedging instruments for a pension or annuity portfolio. The major types are listed briefly in Table 4.1, including longevity / survivor bond, longevity / survivor swap, q-forward, S-forward, K-forward, mortality option, and survivor option. The very first authors who suggested these ideas in the literature are also stated in the table. The focus here is on those that are standardised in nature, i.e. based on a particular reference population, such as English and Welsh males aged 65 in a certain calendar year.

Table 4.1   Mortality-linked and longevity-linked securities

Security Type	Key Features
Longevity bond (Blake and Burrows 2001, Dowd 2003)	Coupon sizes are linked to the percentage of the reference population who are still alive (i.e. survivor index) on the coupon payment dates, in which the survivor index at time t is calculated as $_t p_{65} = (1 - q_{65,0})(1 - q_{66,1})...(1 - q_{65+t-1,t-1})$ and $q_{x,t}$ is the mortality rate of the reference population observed in year t
Longevity swap (Dowd 2003, Dowd et al. 2006)	Two series of future cash flows are exchanged, one of which is linked to the percentage of the reference population who are still alive (i.e. survivor index) on the payment dates, and the other series is fixed at time 0
q-forward (Coughlan et al. 2007)	To a fixed rate receiver, a payoff of $\left(q_{x,T}^{\text{forward}} - q_{x,T}\right)$ is made after $T+1$ years (maturity), in which $q_{x,T}^{\text{forward}}$ is the forward mortality rate set at time 0 and $q_{x,T}$ is the actual mortality rate of the reference population observed in year T; for a floating rate receiver, the payoff is $\left(q_{x,T} - q_{x,T}^{\text{forward}}\right)$ instead
S-forward (LLMA 2010)	The payoffs are similar to those of the q-forward, with the mortality rate being replaced by the percentage of the reference population who are still alive (i.e. survivor index) on maturity
K-forward (Chan et al. 2014, Tan et al. 2014)	To a fixed rate receiver, a payoff of $\left(\kappa_{T,i}^{\text{forward}} - \kappa_{T,i}\right)$ is made after $T+1$ years (maturity), in which ($\kappa_{t,i}$ is the ith CBD model parameter as the ith type of mortality index in year t) $\kappa_{T,i}^{\text{forward}}$ is the forward mortality index set at time 0 and $\kappa_{T,i}$ is the mortality index calculated from the actual observations of the reference population in year T; for a floating rate receiver, the payoff is $\left(\kappa_{T,i} - \kappa_{T,i}^{\text{forward}}\right)$ instead
Mortality option (Cairns et al. 2008)	To a call holder, a payoff of $\max\left(q_{x,T} - q_{x,T}^{\text{strike}}, 0\right)$ is made after $T+1$ years (maturity), in which $q_{x,T}^{\text{strike}}$ is a fixed rate set at time 0 and $q_{x,T}$ is the actual mortality rate of the reference population observed in year T; for a put holder, the payoff is $\max\left(q_{x,T}^{\text{strike}} - q_{x,T}, 0\right)$ instead

Table 4.1   Continued

Security Type	Key Features
Survivor option (Dowd 2003)	The payoffs are similar to those of the mortality option, with the mortality rate being replaced by the percentage of the reference population who are still alive (i.e. survivor index) on maturity

We now consider the use of standardised longevity swaps, in which two series of future cash flows are exchanged. One series is linked to the percentage of the reference population who are still alive (a random quantity) on the payment dates, and the other is fixed since the start of the contract. (Bespoke longevity swaps are by far the most commonly used hedging instruments in practice.) Note that S-forwards can be seen as the building blocks for both the longevity swap and longevity bond. Some other types of securities are also tested in Section 5.

4.1   Preliminary analysis

To start with, we construct a simple hypothetical scenario of a pension portfolio in the UK. The current date ($t = 0$) is set as the beginning of year 2014. All the pensioners are currently aged 65 and each pension pays £1 per year on survival from ages 66 to 90. The pension plan is closed and there are no new entrants after the current date. The pensioners' (book population) mortality experience is the same as that in the CMI data or the ONS data. The pension plan sponsor wants to minimise its own longevity risk exposure by implementing a longevity hedge using standardised longevity swaps. Assume that a longevity swap for the same birth cohort as the pensioners is available in the life market, in which the underlying reference population is English and Welsh population and the maturity is 25 years from now. The corresponding fixed series of payments (i.e. forward survivor index) is calculated from the best (central) estimates of future mortality rates for convenience. (This simple assumption indicates a zero risk premium and would affect only the price but not the effectiveness of the longevity hedge.) Suppose the interest rate is 1% p.a. flat during the period (considering UK Gilt 10-year and 30-year yields of 1.05% p.a. and 1.67% p.a. as at 24 April 2017). When the Mercer data are used for the book population instead, the reference population is taken as Australian or New Zealand population accordingly, and the interest rate is changed to 3% p.a. (considering Australian Government Bond 10-year and 15-year yields of 2.59% p.a. and 2.98% p.a., and New Zealand Government Bond 10-year yield of 3.00% p.a., as at 24 April 2017). The current dates are then reset as the beginning of years 2012 and 2010 respectively.

The present value of all future cash outflows of the pension portfolio over the next 25 years is a random quantity and can be expressed as:

$$\text{PV}\left(\text{unhedged portfolio} \right) = \sum_{t=1}^{25} l_{65+t,t}^{B} \left(1 + i \right)^{-t}$$

The future number of lives $l_{x,t}^B$ aged x at time t is simulated as in Subsections 3.4 and 3.5 (Figure 3.3) based on the book population. The interest rate for discounting is denoted as i. Moreover, assuming payments are made annually, the random present value of all future cash inflows of the longevity swap (as a floating rate receiver; based on the reference population) for the same birth cohort as the pensioners is equal to:

$$\text{PV}\left(\text{hedging instrument}\right) = \sum_{t=1}^{25} \left({}_t p_{65}^R - {}_t p_{65}^{R;\text{forward}} \right) (1+i)^{-t}$$

The survivor index (or survivor rate) ${}_t p_{65}^R = \left(1 - q_{65,0}^R\right)\left(1 - q_{66,1}^R\right)\ldots\left(1 - q_{65+t-1,t-1}^R\right)$ is simulated as in Subsection 3.4 using the reference population, assuming the sampling variability in the future simulations is immaterial for such a large population. The corresponding forward survivor index ${}_t p_{65}^{R;\text{forward}}$ is calculated similarly from the best (central) estimates of future mortality rates, i.e. setting all the time series error terms to zero in the projection.

Then the random present value of the aggregate pension portfolio position after taking the longevity hedge is stated as:

$$\text{PV}\left(\text{hedged portfolio}\right) = \sum_{t=1}^{25} l_{65+t,t}^B (1+i)^{-t} - w\sum_{t=1}^{25} \left({}_t p_{65}^R - {}_t p_{65}^{R;\text{forward}} \right) (1+i)^{-t}$$

and the random cash outflow of the net position at each time $t = 1, 2, \ldots, 25$ is denoted as:

$$\text{CF}_t\left(\text{hedged portfolio}\right) = l_{65+t,t}^B - w\left({}_t p_{65}^R - {}_t p_{65}^{R;\text{forward}} \right)$$

The weight w is the required notional amount of the longevity swap. It can be estimated by numerical optimisation with an objective to minimise the risk or uncertainty of the random present value of the aggregate position. In line with the Solvency Capital Requirement (SCR) in Solvency II, the 99.5% value-at-risk (VaR) is used as the risk measure in the optimisation process here. In practice, other risk measures or a mix of different objectives may also be adopted, depending on the particular circumstances and the purpose of the analysis. In contrast, if the hedge is customised perfectly to the pension portfolio (which is more common in current industry practice; see LCP 2012) rather than being constructed with standardised instruments, and so there is no longevity basis risk, the weight can simply be set to be equal to the initial portfolio size and then effectively all longevity risk is transferred away.

In order to examine the actual effectiveness of this longevity hedge, we calculate the level of longevity risk reduction as defined in Subsection 3.7. For the terms risk(unhedged) and risk(hedged), the portfolio's longevity risk before and after taking the hedge respectively, we consider the variance, standard deviation, 99.5% VaR minus the mean, and 99.5% expected shortfall (conditional VaR) minus the mean of the present value of the aggregate position. The first two risk measures are estimated as the sample variance and sample standard deviation. The 99.5% VaR is computed as the sample 99.5[th] percentile. The 99.5% expected shortfall is calculated as the sample mean of all those simulated outcomes above the sample 99.5[th] percentile. The mean is taken as the best (central) estimate. Note that while the VaR is

the most common choice in financial practice and is the prescribed measure in Solvency II, the expected shortfall (conditional VaR) contains more information about the tail events and has better mathematical properties and more reasonable implications when compared to the VaR (Dowd and Blake 2006). There are also other decent alternatives such as spectral risk measures and distortion risk measures.

Table 4.2 (Figure 4.1) shows the level of longevity risk reduction estimates using the ONS dataset (with England IMD quintile groups). The percentage estimate represents the proportion of the portfolio's initial longevity risk that is being hedged away. In each cell of the table, the left figure is calculated from the M7-M5 model, and the right from the CAE+Cohorts model. In this dataset, the average annual number of lives is around one million for each quintile group and the data period covers more than 10 years, which suggest that either the M7-M5 model or CAE+Cohorts model (i.e. direct modelling) can be used (Figure 3.1). Four hypothetical cases are first considered, in which one has an infinitely large portfolio size (i.e. no sampling basis risk explicitly allowed for in the future simulations) and the others have an initial portfolio size of 100,000, 25,000, and 2,500 respectively. As noted in Phase 1, an exposure of 100,000 lives is the largest any scheme or insurer is likely to have in the UK.

There are a number of main observations from the table. Firstly, the longevity risk reduction estimates regarding the standard deviation, 99.5% VaR, and 99.5% expected shortfall (conditional VaR) range mostly from around 60% to above 90% for the three larger portfolio sizes. The estimates for these three risk measures are fairly close to one another in most cases. (On the other hand, the estimates for the variance, which has a different scale to the other risk measures, are generally larger at 80% or above. The following analysis will focus on the three risk measures above.) Secondly, as the initial portfolio size decreases, sampling basis risk in the future simulations increases, and so the level of longevity risk reduction declines. This drop is somewhat uneven – the longevity risk reduction estimates do not change much until the portfolio size decreases to around 25,000, below which the decline gets more and more significant. For a small portfolio size of only 2,500, for example, the longevity risk reduction estimates largely drop to around 50% or less. These results imply that sampling basis risk is rather immaterial when the portfolio size is more than about 25,000, but it can have a significant effect when the portfolio has only, say, a few thousand lives.

Moreover, for the larger portfolio sizes, the third most deprived areas have many estimates being close to or above 90% while the other areas have around 60% to 90% risk reduction. The former is the group with 'medium' or 'average' income, employment, health, education, crime rates, etc. and so it can potentially be matched more closely by the overall reference population, which would explain its higher risk reduction levels. Between both sexes, the differences in the hedging results are quite small and mostly less than 10% in magnitude,

except for the least deprived areas, in which the risk reduction estimates are clearly larger for males in all the cases shown.

Table 4.2 Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of England IMD quintile groups of a single cohort

Group \ Size	Females				Males			
	Variance	Standard Deviation	99.5% VaR	99.5% ES	Variance	Standard Deviation	99.5% VaR	99.5% ES
Most deprived areas								
Infinite	82 / 92	58 / 72	65 / 76	66 / 76	87 / 90	64 / 68	71 / 68	70 / 69
100,000	81 / 91	57 / 71	63 / 73	64 / 74	87 / 90	64 / 69	72 / 68	70 / 68
25,000	80 / 89	56 / 68	62 / 67	63 / 68	85 / 87	62 / 64	67 / 63	67 / 63
2,500	70 / 73	46 / 48	46 / 48	46 / 48	73 / 73	48 / 48	51 / 46	50 / 48
Second most deprived areas								
Infinite	94 / 97	76 / 83	77 / 83	77 / 83	96 / 98	81 / 85	83 / 87	83 / 86
100,000	94 / 96	75 / 81	75 / 78	76 / 79	96 / 97	79 / 83	81 / 83	80 / 83
25,000	93 / 94	73 / 76	72 / 72	72 / 72	94 / 95	76 / 78	77 / 76	77 / 77
2,500	80 / 75	55 / 50	52 / 47	52 / 45	80 / 77	55 / 52	54 / 52	53 / 52
Third most deprived areas								
Infinite	100 / 100	94 / 95	93 / 95	93 / 95	100 / 100	94 / 95	94 / 94	94 / 94
100,000	99 / 99	91 / 90	91 / 89	91 / 89	99 / 99	90 / 90	91 / 90	91 / 90
25,000	98 / 97	85 / 82	84 / 80	84 / 79	98 / 97	84 / 82	84 / 81	84 / 81
2,500	83 / 75	59 / 50	56 / 48	57 / 47	83 / 78	59 / 53	56 / 51	57 / 50
Fourth most deprived areas								
Infinite	99 / 100	91 / 94	89 / 94	85 / 93	98 / 99	88 / 88	81 / 84	76 / 82
100,000	99 / 99	89 / 89	88 / 89	85 / 89	98 / 99	86 / 88	81 / 87	75 / 85
25,000	97 / 96	83 / 81	83 / 81	82 / 80	96 / 96	80 / 81	76 / 79	71 / 79
2,500	83 / 75	59 / 50	58 / 47	58 / 46	82 / 77	57 / 52	57 / 50	55 / 50
Least deprived areas								
Infinite	92 / 95	72 / 77	60 / 65	49 / 61	97 / 98	83 / 87	75 / 83	70 / 82
100,000	92 / 93	71 / 73	59 / 64	49 / 61	96 / 98	80 / 85	76 / 81	71 / 79
25,000	90 / 90	69 / 69	55 / 63	45 / 61	95 / 96	78 / 79	73 / 78	69 / 77
2,500	78 / 67	53 / 42	50 / 38	44 / 34	80 / 75	55 / 50	54 / 49	52 / 48

Figure 4.1   Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of England IMD quintile groups of a single cohort (males)
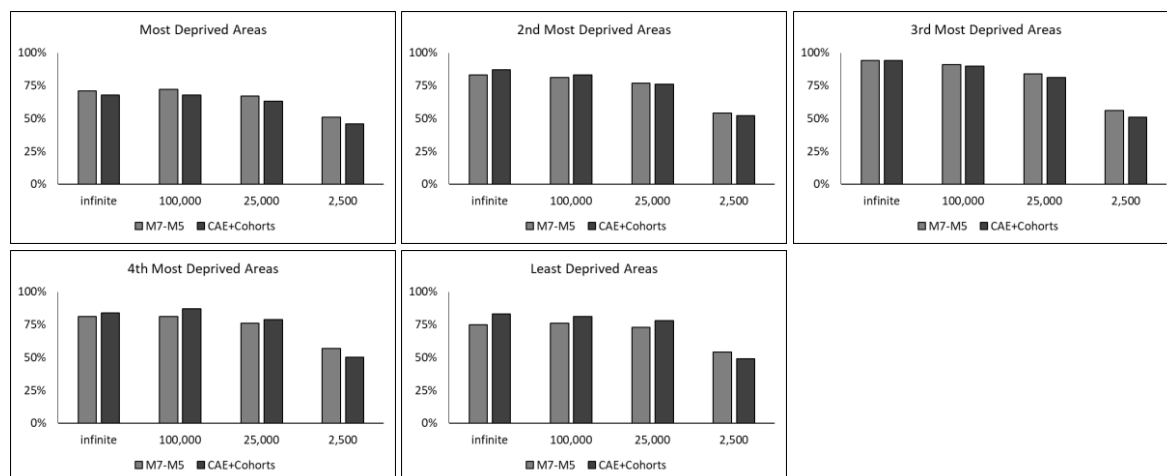
Finally, it is worth noting that the two models do not give rise to too much difference in the hedging results, which is less than 10% in magnitude for most of the estimated figures. Interestingly, while the CAE+Cohorts model tends to produce larger risk reduction estimates compared to the M7-M5 model for a large portfolio, the situation seems to be reversed for a small portfolio. It means that the hedging results calculated under the CAE+Cohorts model tend to cover a wider range than those under the M7-M5 model for different portfolio sizes.

Tables 4.3 and 4.4 (Figures 4.2 and 4.3) present the level of longevity risk reduction estimates for the CMI (with UK pensioners) and Mercer (with Australian and New Zealand pensioners) datasets. Again, in each cell of the tables, the left figure is computed using the M7-M5 model and the right from the CAE+Cohorts model. Regarding these two datasets, the groups chosen for demonstration here have around 20,000 lives or more per year on average and their data periods are at least 8 years, so the M7-M5 and CAE+Cohorts models would be suitable choices (Figure 3.1). For the hypothetical scenario, two initial portfolio sizes of 100,000 and 1,000 are examined in this part of the analysis.

Similar observations can be made on the two tables. First, the longevity risk reduction estimates with regard to the standard deviation, 99.5% VaR, and 99.5% expected shortfall (conditional VaR) range largely from around 50% to 80% for a portfolio size of 100,000. By contrast, for a small portfolio size of 1,000, the estimates decrease to only around 30% to 40%. Combined with the earlier observations from Table 4.2, all these results suggest that for a portfolio which has at least a few thousand lives, roughly speaking, more than 50% of the portfolio's longevity risk exposure can be hedged via the longevity swap. But for a small portfolio with only one or two thousand or even fewer lives, less than 50% of the risk can be transferred away. Note that sampling basis risk would unavoidably 'destroy' the relationship in future mortality movements between the book and reference populations, especially so for

a small portfolio. The implication is that index-based hedging looks more effective for reasonably large portfolios than for those that are too small. Second, for the larger portfolio sizes, the UK, Australian, and New Zealand pensioners generally have smaller risk reduction estimates compared to the England IMD quintile groups. The IMD groups represent significant segments of the entire reference population (i.e. smaller demographic basis risk), and so it is reasonable to expect that the risk reduction effect would be less obvious for the pensioners than for the IMD groups. For the smaller portfolio sizes, however, the differences in the hedging results between the pensioners and the IMD groups become less obvious, which again highlight the influence of sampling basis risk. For instance, both have only around 30% risk reduction in most of the cases simulated for a portfolio size of 1,000 (not shown in Table 4.2).

Furthermore, it can be seen again that the hedging results computed under the CAE+Cohorts model have a tendency to cover a wider range. When compared to the M7-M5 model, the CAE+Cohorts model tends to produce larger risk reduction estimates for a large portfolio but smaller estimates for a small portfolio. However, the two models produce more variable results between them in these two tables. While most differences are still less than 15% in magnitude, a few cases show more significant discrepancies between the two models. The smaller data sizes of the CMI and Mercer datasets (compared with the ONS dataset) may be the underlying reason, in which the more noisy patterns may have been captured rather differently under the unique features of each model. In particular, for those few cases which demonstrate obvious inconsistency, the M7-M5 model is usually the one that gives much smaller risk reduction estimates.

Apart from investigating the level of longevity risk reduction after hedging, it is also informative to examine the size of the portfolio's initial longevity risk before hedging, as a percentage of the portfolio's mean present value. For all the groups in Tables 4.3 and 4.4 with a portfolio size of 100,000, under the M7-M5 model, the standard deviation mostly ranges from 1.5% to 2.5%, the 99.5% VaR (minus the mean) from 3.5% to 5.5%, and the 99.5% expected shortfall (conditional VaR; minus the mean) from 4% to 6.5%. In contrast, under the CAE+Cohorts model, the ranges of the three risk measures are 1% to 1.5%, 2.5% to 4%, and 2.5% to 4.5% respectively. In general, the M7-M5 model induces higher variability in the simulated present values of the pension portfolio, probably because it involves modelling and simulating five kappa time series, compared to only two in the CAE+Cohorts model. When the portfolio size is only 1,000, the standard deviation (as a percentage of the mean) increases by less than 0.5% in magnitude while the tail measures increase by around 1%. Note that the expected shortfall (conditional VaR) contains more information about the extreme events, possesses better theoretical properties, and provides more conservative estimates than the VaR does for a given confidence level. In fact, banking regulations are gradually moving towards the use of this new risk measure in determining regulatory capital for market risk (Hull 2015).

Table 4.3 **Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of UK pensioners of a single cohort**

Group \ Size	Variance	Standard Deviation	99.5% VaR	99.5% ES
Males				
Basic materials	(normal retirement; lower pension group)			
100,000	87 / 95	65 / 77	63 / 74	63 / 73
1,000	60 / 57	37 / 34	36 / 37	38 / 36
Industrials	(normal retirement; lower pension group)			
100,000	77 / 81	52 / 56	51 / 50	39 / 49
1,000	54 / 52	32 / 30	35 / 25	27 / 28
Consumer goods	(normal retirement; lower pension group)			
100,000	81 / 92	57 / 72	39 / 66	29 / 62
1,000	57 / 55	35 / 33	31 / 27	25 / 31
Commercial services	(normal retirement; lower pension group)			
100,000	95 / 97	78 / 83	79 / 79	79 / 79
1,000	66 / 55	41 / 33	41 / 30	41 / 28
Utilities	(normal retirement; lower pension group)			
100,000	65 / 93	41 / 73	58 / 70	56 / 69
1,000	50 / 54	29 / 32	30 / 27	33 / 28
Local authority	(normal retirement; lower pension group)			
100,000	94 / 98	75 / 85	73 / 83	58 / 83
1,000	59 / 54	36 / 32	33 / 30	25 / 27
All industries	(ill-health retirement)			
100,000	90 / 96	68 / 79	78 / 78	75 / 77
1,000	61 / 57	37 / 35	41 / 35	41 / 34

Table 4.4   Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of Australian and New Zealand pensioners of a single cohort

| Group \ Size | Males | | | |
	Variance	Standard Deviation	99.5% VaR	99.5% ES
New South Wales, Australia	(retirement)			
100,000	92 / 93	72 / 74	69 / 70	63 / 68
1,000	73 / 60	48 / 37	42 / 33	38 / 33
Commonwealth, Australia	(retirement)			
100,000	94 / 96	76 / 80	65 / 77	52 / 76
1,000	72 / 63	47 / 39	41 / 36	39 / 34
Victoria, Australia	(retirement)			
100,000	72 / 70	47 / 45	41 / 44	33 / 40
1,000	60 / 52	37 / 31	30 / 30	25 / 27
Other states in Australia	(retirement)			
100,000	76 / 85	51 / 62	45 / 62	30 / 60
1,000	62 / 60	38 / 37	37 / 37	26 / 34
New Zealand	(retirement)			
100,000	95 / 97	77 / 82	74 / 79	67 / 79
1,000	81 / 74	57 / 49	55 / 44	53 / 44
Australia, New Zealand	(invalidity)			
100,000	81 / 92	56 / 72	40 / 72	28 / 72
1,000	63 / 63	39 / 39	34 / 38	25 / 35

In the CMI and Mercer datasets, there are a number of groups or states with much fewer than 25,000 lives per year in their data, especially for females. As noted in Subsection 3.5, the sampling variability in such a small set of past data may flow through to the time series of the fitted model parameters, leading to an overestimation of process error, parameter error, and hence demographic basis risk. The characterisation approach described in Subsection 3.3 may then be adopted in this situation, with an attempt to provide a more reasonable account of the extent of process error and parameter error.

Figure 4.2 **Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of UK pensioners of a single cohort (males; normal retirement; lower pension group)**
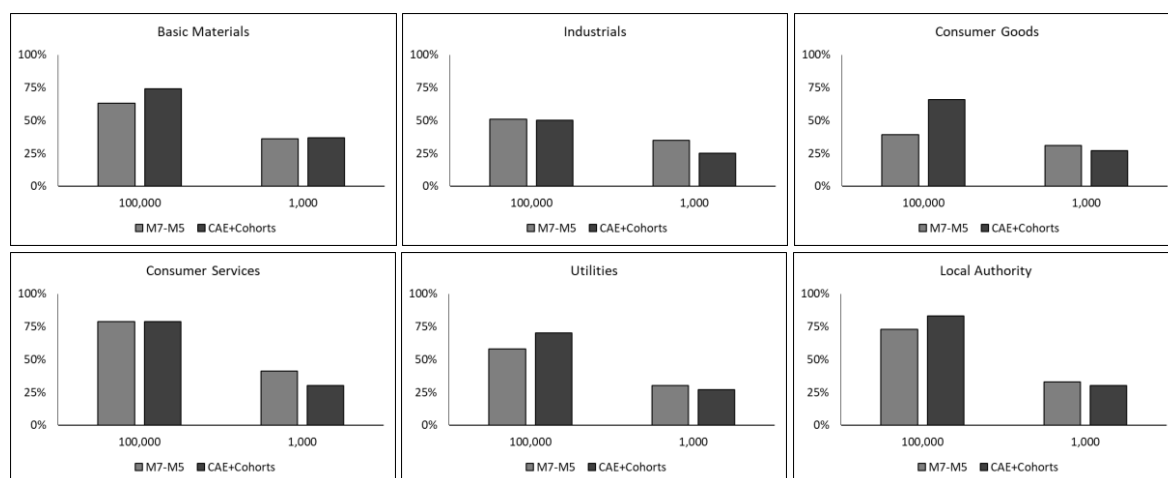
Figure 4.3 **Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of Australian and New Zealand pensioners of a single cohort (males)**
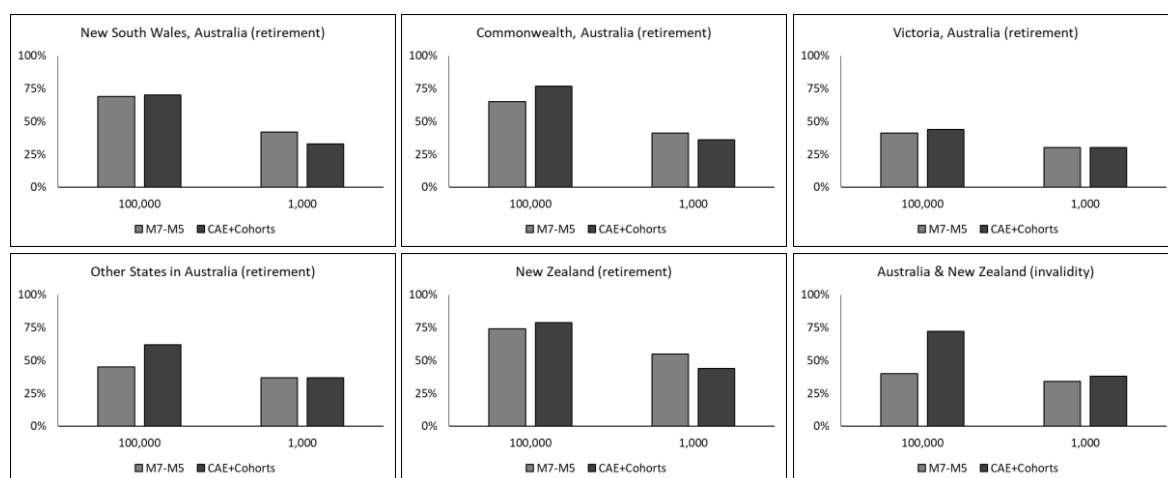
Beside the option of using the characterisation approach, as a further checking and a very rough approximation here, we still apply the M7-M5 and CAE+Cohorts models directly to those small data samples, regardless of their sizes. In effect, we implicitly allow the extra sampling variability in the past data to transfer in some (unknown) way to the future simulations, but at the same time omit the last step of using the binomial distribution to avoid 'double counting' the sampling variability, assuming the portfolio size remains more or less the same over time. Though this assessment is not precise at all, it turns out that many of the resulting longevity risk reduction estimates are well below 50% (not shown here), which are quite in line with the earlier discussion that less than half of the risk can be transferred away for a small portfolio.

As discussed in Subsection 3.6, structural basis risk exists because both the timing and amounts of future cash flows of the hedging instrument are often different to those of the portfolio being hedged. In this hypothetical scenario, the main problem is that the amounts of future cash flows of the longevity swap (based on the reference population) are obviously different to those of the pension portfolio (based on the book population). We use a simple numerical optimisation procedure (see Appendix I) to find the optimal position of the longevity swap to minimise longevity risk and reduce the effect of this structural basis risk. Accordingly, apart from studying the present value of all future cash flows as above, we also examine the individual cash flows (without discounting) and calculate their own risk reduction levels with respect to each future year. We observe that under the current model settings and assumptions, for those cases with a risk reduction well above 50% on the overall present value, the risk reduction estimates of the individual cash flows are indeed far below 50% in the early years from the current date ($t = 0$), but increase gradually every year to exceed 50% in the later years. From a modelling perspective, this effect appears to arise from the fact that the fitted time series processes for the reference population produce unbounded future variability while the fitted time series processes for the difference between the book and reference populations yield bounded future variability (see Section 6). As a result, the latter variability would reduce in significance relatively over time, which means that the simulated differences between the book and reference populations would become less and less important comparatively and so there is a decline in demographic basis risk being projected across time. From a practical perspective, it would also be rational to argue that the two related populations' mortality improvements may deviate in the short term but would be more in line (and so there is less demographic basis risk) over the long term.

4.2 Realistic scenarios

We now consider a more realistic hypothetical scenario of a pension portfolio with multiple cohorts. Suppose the total number of pensioners is currently 30,000. Figure 4.4 shows the split between different ages (60 to 89) in the portfolio, which is based roughly on the average proportions observed in the three datasets. We examine both an open pension plan and a closed one. For the open pension plan, we assume that there is exactly the same number of new members (1,400) joining the plan at age 60 every year, and that a standardised longevity swap for the cohort currently aged 60 (with a maturity of 30 years) is used to construct the hedge. For the closed pension plan, we assume that there are no more new entrants after the current date ($t = 0$), and that two standardised longevity swaps for the two cohorts now aged 60 and 70 (with maturities of 30 and 20 years respectively) are available for building the hedge. Each pensioner receives £1 per annum on survival from ages 61 to 90. The time horizon of the hedging exercise is 30 years from the current date. All the other settings are the same as previously.

Figure 4.4   Initial demographic structure of a hypothetical pension plan with multiple cohorts
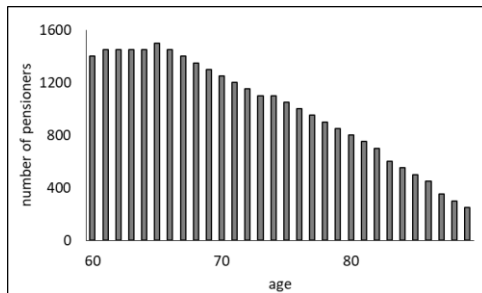
Tables 4.5 to 4.7 (Figures 4.5 to 4.7) set forth the level of longevity risk reduction estimates using the three datasets. First of all, the longevity risk reduction estimates regarding the standard deviation and the extreme risk measures vary mostly from around 30% to 60% for the open pension plan, and from about 50% to 80% for the closed pension plan, except for a small number of cases (e.g. industrials of UK pensioners and Victoria of Australian pensioners). These estimates are often smaller than those calculated in the previous subsection which involves only one cohort, as the two portfolios considered here have multiple cohorts being hedged by only one or two longevity swaps, and so the extent of demographic basis risk (in terms of age or cohort difference) is greater. Secondly, the risk reduction estimates are smaller for the open pension plan than for the closed pension plan. It appears to be caused by the fact that the open pension plan uses only one longevity swap to construct the hedge but for a larger number of cohorts (i.e. existing pensioners plus new entrants). Note that when the life market is still in its infancy, those longevity swap transactions available may be limited to a small number of cohorts, such that market liquidity can be focused and enhanced. For the open pension plan, those currently aged 31 to 59 will join the plan gradually in the future, while the existing pensioners are aged 60 to 89 now. So a longevity swap for the cohort currently aged 60 (in the middle of the whole age range considered) is chosen here for convenience. In contrast, for the closed pension plan, there are no new entrants in the future. Accordingly, two longevity swaps for the two cohorts now aged 60 and 70 are selected – the first one is included for covering the entire time horizon while the second one of age 70 refers approximately to the weighted (by the initial numbers of lives) average age of the current pensioners. If longevity swaps for more cohorts and with different maturities are also available in the market and being included, one may use more sophisticated numerical optimisation techniques or adapt certain financial hedging strategies to further improve the overall hedge effectiveness by reducing the extent of demographic basis risk. Nevertheless, the number of pensioners at each single age in the example is actually quite small, and using longevity swaps of too many cohorts would induce more sampling basis risk. A fine balance must be struck between these two offsetting effects. A further potentially feasible solution in practice is for those smaller pension plans to join bigger foundations (LCP 2012), which can then be hedged by more standardised instruments with less concern on sampling basis risk.

41

Table 4.5 **Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of England IMD quintile groups of multiple cohorts**

Group	Females				Males			
	Variance	Standard Deviation	99.5% VaR	99.5% ES	Variance	Standard Deviation	99.5% VaR	99.5% ES
Most deprived areas								
Open	69 / 62	44 / 38	48 / 39	47 / 42	75 / 59	50 / 36	53 / 38	51 / 37
Closed	79 / 87	54 / 64	58 / 63	59 / 63	81 / 81	56 / 57	61 / 55	59 / 55
Second most deprived areas								
Open	81 / 66	56 / 42	54 / 38	55 / 40	82 / 67	57 / 42	53 / 42	52 / 38
Closed	90 / 92	69 / 71	67 / 68	67 / 67	92 / 93	72 / 74	73 / 71	72 / 71
Third most deprived areas								
Open	85 / 70	62 / 45	62 / 44	61 / 43	84 / 69	60 / 44	59 / 43	60 / 41
Closed	96 / 95	80 / 78	79 / 77	80 / 77	96 / 96	80 / 79	79 / 78	79 / 78
Fourth most deprived areas								
Open	85 / 71	61 / 46	63 / 45	61 / 44	84 / 70	60 / 45	59 / 48	58 / 46
Closed	95 / 95	79 / 78	81 / 77	77 / 77	95 / 95	77 / 78	77 / 78	73 / 75
Least deprived areas								
Open	80 / 67	55 / 43	50 / 43	44 / 44	81 / 71	57 / 46	55 / 46	55 / 44
Closed	86 / 90	63 / 68	57 / 65	48 / 60	92 / 94	71 / 75	70 / 75	66 / 73

Some similar patterns to the previous subsection can be observed here. For instance, the third most deprived areas tend to have the largest risk reduction estimates, probably because this group has around average living standards and so it can be hedged more closely by the reference population. Moreover, the England IMD quintile groups often demonstrate higher risk reduction levels when compared to the UK and Australian pensioners. The IMD groups represent major portions of the reference population, which may involve smaller demographic basis risk and explain their larger risk reduction estimates. Lastly, most differences in the hedging results between the two models are less than 20% in magnitude. A few cases show larger differences, in which the CAE+Cohorts model tends to generate smaller risk reduction estimates. It is also noted that the differences in the hedging results between the two pension plans under the CAE+Cohorts model are often greater than those under the M7-M5 model. While the overall patterns observed in the two sets of results are

broadly comparable, the extent of model error is not negligible, particularly for a smaller data size (as mentioned in the previous subsection) and for a less precise hedging scheme (like using only one longevity swap for multiple cohorts here). The impact of different model choices is further explored in Sections 5 and 6.

Table 4.6   **Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of UK pensioners of multiple cohorts**

| Group | Males | | | |
	Variance	Standard Deviation	99.5% VaR	99.5% ES
Basic materials	(normal retirement; lower pension group)			
Open	76 / 66	51 / 42	50 / 43	49 / 44
Closed	84 / 88	60 / 66	60 / 64	58 / 62
Industrials	(normal retirement; lower pension group)			
Open	67 / 58	43 / 35	46 / 32	29 / 30
Closed	69 / 71	44 / 47	45 / 42	37 / 39
Consumer goods	(normal retirement; lower pension group)			
Open	70 / 64	45 / 40	38 / 35	29 / 36
Closed	79 / 86	54 / 62	45 / 56	37 / 49
Commercial services	(normal retirement; lower pension group)			
Open	84 / 65	60 / 41	58 / 40	57 / 40
Closed	93 / 93	73 / 74	74 / 70	72 / 70
Utilities	(normal retirement; lower pension group)			
Open	54 / 63	32 / 39	47 / 38	46 / 35
Closed	71 / 89	47 / 66	56 / 63	51 / 60
Local authority	(normal retirement; lower pension group)			
Open	81 / 66	56 / 42	54 / 44	47 / 44
Closed	91 / 94	69 / 75	67 / 73	55 / 72
All industries	(ill-health retirement)			
Open	77 / 56	52 / 34	58 / 36	57 / 35
Closed	85 / 92	61 / 71	73 / 69	73 / 69

Table 4.7 **Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of Australian and New Zealand pensioners of multiple cohorts**

| Group | Males | | | |
	Variance	Standard Deviation	99.5% VaR	99.5% ES
New South Wales, Australia	(retirement)			
Open	77 / 76	52 / 51	48 / 45	45 / 45
Closed	89 / 90	68 / 69	64 / 64	60 / 63
Commonwealth, Australia	(retirement)			
Open	79 / 77	54 / 52	51 / 45	43 / 46
Closed	90 / 92	68 / 72	64 / 66	52 / 65
Victoria, Australia	(retirement)			
Open	59 / 58	36 / 35	30 / 29	25 / 25
Closed	61 / 59	37 / 36	34 / 32	26 / 26
Other states in Australia	(retirement)			
Open	67 / 68	42 / 43	38 / 40	26 / 39
Closed	74 / 74	49 / 49	44 / 51	32 / 50
New Zealand	(retirement)			
Open	74 / 72	49 / 47	50 / 44	49 / 41
Closed	88 / 92	66 / 71	62 / 68	61 / 68
Australia, New Zealand	(invalidity)			
Open	74 / 74	49 / 49	40 / 53	25 / 50
Closed	83 / 91	58 / 69	51 / 70	45 / 69

Figure 4.5 **Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of England IMD quintile groups of multiple cohorts (males)**

Figure 4.5   Continued

Figure 4.6   Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of UK pensioners of multiple cohorts (males; normal retirement; lower pension group)
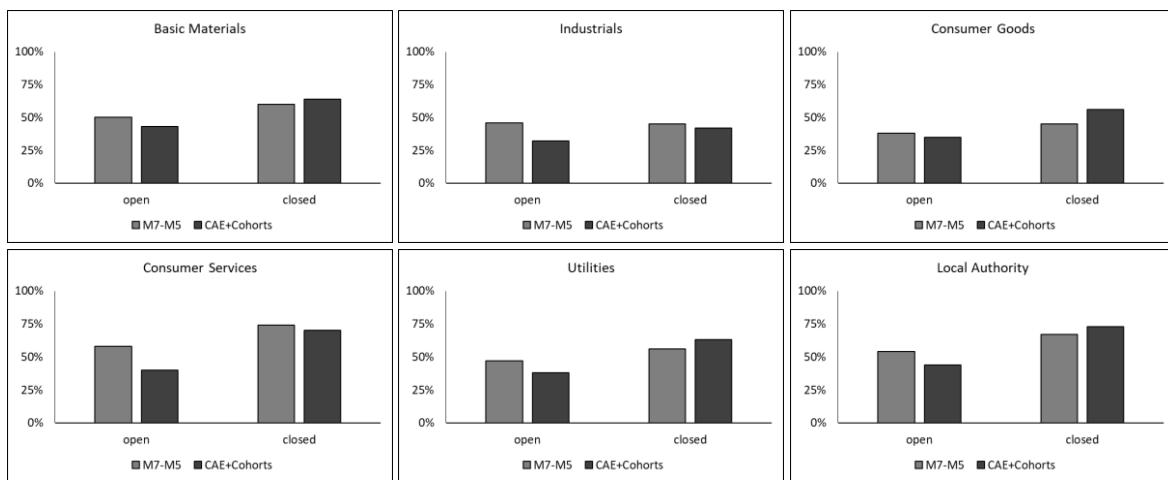
Figure 4.7   Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of Australian and New Zealand pensioners of multiple cohorts (males)
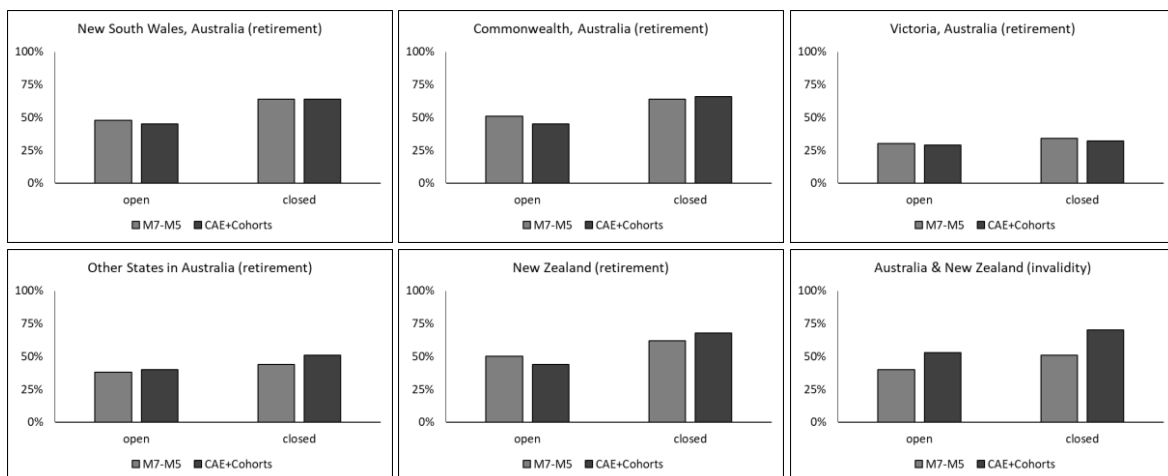
As discussed earlier, if more longevity swaps are included and calibrated separately for each cohort involved, there would be more sampling basis risk because the number of pensioners at each single cohort is quite small in the example. A possible alternative is to calibrate the swaps altogether as a group to the entire portfolio, rather than catering specifically for each cohort. In this way, the overall hedge effectiveness could further be improved, though there would be some implicit offsetting effects between different cohorts in the hedge. We now experiment with eleven standardised longevity swaps for the cohorts currently aged 35, 40, 45, … , 85 (with a (delayed) maturity of 5, 10, 15, 20, 25 years for those who have not entered the plan yet and maturities of 30, 25, 20, 15, 10 , 5 years for existing pensioners) for the open pension plan. We also try using six standardised longevity swaps for the cohorts now aged 60, 65, 70, … , 85 (with maturities 30, 25, 20, … , 5 years respectively) for the closed pension plan. As shown in Tables 4.8 to 4.10 (Figures 4.8 to 4.10), this new strategy of adopting 5-year age buckets and aggregate calibration clearly increases the risk reduction estimates in many of the cases. The effect is much more apparent for the open pension plan, which contains a larger number of cohorts (current pensioners plus new entrants) and has employed only one swap in the original setting. Although there would be complex subsidising effects operating between different cohorts in this revised hedge, the aggregate extent of demographic basis risk does appear to be reduced by the use of more swaps, and the overall hedge effectiveness is obviously enhanced further. On the other hand, the effect of incorporating more swaps is marginal for the closed pension plan, implying that the two swaps in the original setting (over 10-year age buckets) would probably be adequate to deliver a similar level of hedging performance and adding more swaps would not bring about much improvement.

Figure 4.8   **Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of England IMD quintile groups of multiple cohorts (males; with swaps over 5-year age buckets)**
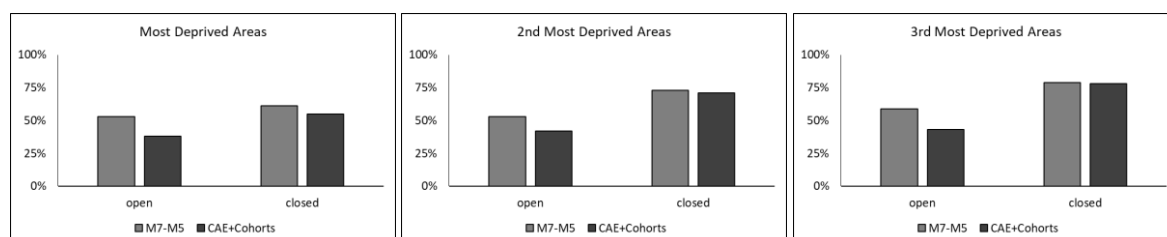
Table 4.8 **Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of England IMD quintile groups of multiple cohorts (with swaps over 5-year age buckets)**

Group	Females				Males			
	Variance	Standard Deviation	99.5% VaR	99.5% ES	Variance	Standard Deviation	99.5% VaR	99.5% ES
Most deprived areas								
Open	79 / 92	55 / 72	62 / 74	62 / 75	85 / 93	62 / 73	68 / 72	67 / 72
Closed	80 / 87	55 / 64	60 / 63	61 / 63	84 / 87	60 / 64	67 / 63	66 / 62
Second most deprived areas								
Open	92 / 96	72 / 79	71 / 75	71 / 76	95 / 96	77 / 81	78 / 81	78 / 80
Closed	90 / 93	69 / 74	67 / 71	68 / 70	93 / 93	73 / 74	67 / 74	74 / 73
Third most deprived areas								
Open	98 / 97	85 / 83	85 / 83	85 / 82	97 / 97	84 / 84	84 / 84	84 / 83
Closed	97 / 96	83 / 81	83 / 81	83 / 79	97 / 96	82 / 81	83 / 79	83 / 79
Fourth most deprived areas								
Open	96 / 97	79 / 83	81 / 82	78 / 81	96 / 97	79 / 83	77 / 83	70 / 83
Closed	96 / 96	81 / 80	82 / 79	79 / 78	96 / 96	79 / 79	77 / 78	74 / 75
Least deprived areas								
Open	86 / 92	63 / 72	54 / 67	45 / 65	95 / 96	77 / 79	75 / 80	69 / 78
Closed	88 / 89	65 / 67	59 / 65	49 / 61	92 / 94	72 / 76	71 / 75	66 / 74

In all the hedging scenarios analysed above, it has been assumed that each pension pays only up to age 90, and so the hedging scheme does not need to cover beyond that age. In reality, however, insurers and pension plans do have exposures to those lives aged over 90. As discussed in Section 2 and Subsection 6.3, the disparities between different subgroups tend to reduce at older ages. This merge in mortality experience over age means that if lives at advanced ages are also included in the portfolio and covered by the index-based hedging scheme, and if the portfolio size is large, the overall risk reduction level, in principle, would actually increase, because of lesser mismatching and so smaller longevity basis risk at the very old ages. Subsection 5.6 later also shows that the hedging results are fairly robust to an inclusion of advanced ages in the modelling process and the portfolio setting.
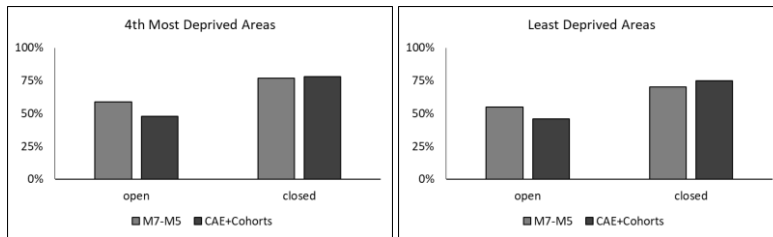
Table 4.9 **Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of UK pensioners of multiple cohorts (with swaps over 5-year age buckets)**

Group	Males			
	Variance	Standard Deviation	99.5% VaR	99.5% ES
Basic materials	(normal retirement; lower pension group)			
Open	88 / 94	65 / 75	65 / 74	63 / 74
Closed	85 / 90	62 / 69	62 / 69	59 / 67
Industrials	(normal retirement; lower pension group)			
Open	76 / 83	51 / 58	56 / 55	37 / 55
Closed	72 / 71	47 / 46	48 / 41	40 / 40
Consumer goods	(normal retirement; lower pension group)			
Open	79 / 91	54 / 70	46 / 66	30 / 64
Closed	79 / 85	54 / 61	45 / 55	38 / 49
Commercial services	(normal retirement; lower pension group)			
Open	95 / 96	77 / 80	79 / 76	78 / 76
Closed	93 / 94	74 / 75	76 / 72	75 / 71
Utilities	(normal retirement; lower pension group)			
Open	62 / 92	38 / 72	63 / 72	62 / 71
Closed	71 / 87	46 / 64	57 / 61	50 / 61
Local authority	(normal retirement; lower pension group)			
Open	91 / 96	70 / 80	68 / 80	54 / 78
Closed	91 / 94	70 / 76	68 / 75	57 / 74
All industries	(ill-health retirement)			
Open	88 / 95	66 / 78	77 / 77	75 / 77
Closed	86 / 92	62 / 72	75 / 71	75 / 70

Table 4.10   Level of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) in a hypothetical scenario of Australian and New Zealand pensioners of multiple cohorts (with swaps over 5-year age buckets)

Group	Males Variance	Standard Deviation	99.5% VaR	99.5% ES
New South Wales, Australia	(retirement)			
Open	89 / 90	67 / 69	64 / 65	60 / 64
Closed	91 / 91	70 / 70	65 / 65	60 / 64
Commonwealth, Australia	(retirement)			
Open	88 / 95	65 / 77	55 / 73	44 / 72
Closed	90 / 94	68 / 75	65 / 72	52 / 70
Victoria, Australia	(retirement)			
Open	67 / 68	42 / 44	37 / 41	27 / 34
Closed	59 / 64	36 / 40	34 / 34	25 / 28
Other states in Australia	(retirement)			
Open	75 / 80	50 / 56	44 / 56	26 / 55
Closed	74 / 76	49 / 51	46 / 53	32 / 52
New Zealand	(retirement)			
Open	92 / 95	72 / 78	69 / 77	67 / 76
Closed	92 / 93	71 / 73	67 / 71	65 / 70
Australia, New Zealand	(invalidity)			
Open	83 / 93	59 / 74	40 / 75	27 / 71
Closed	84 / 92	60 / 71	53 / 71	45 / 70

Figure 4.9   **Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of UK pensioners of multiple cohorts (males; normal retirement; lower pension group; with swaps over 5-year age buckets)**
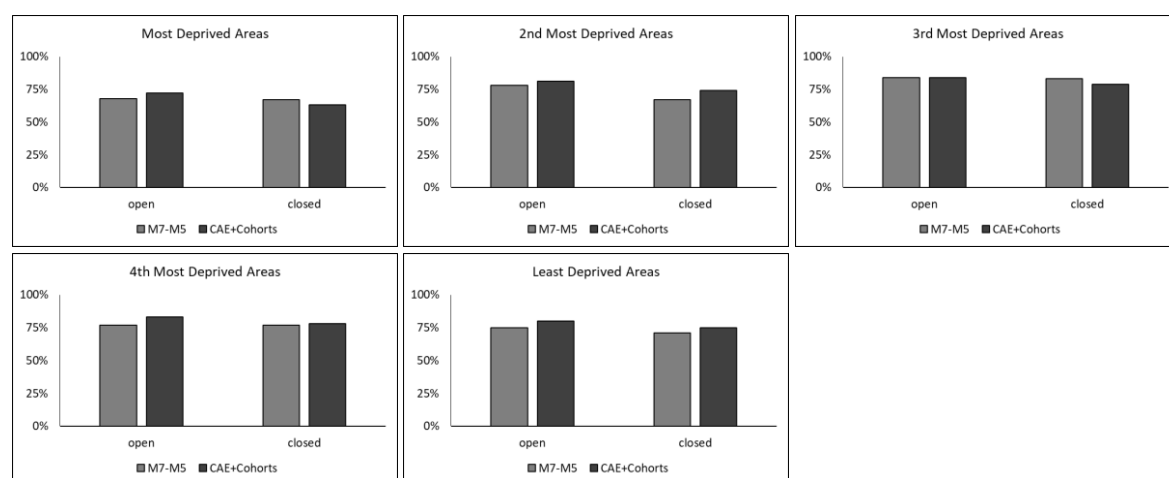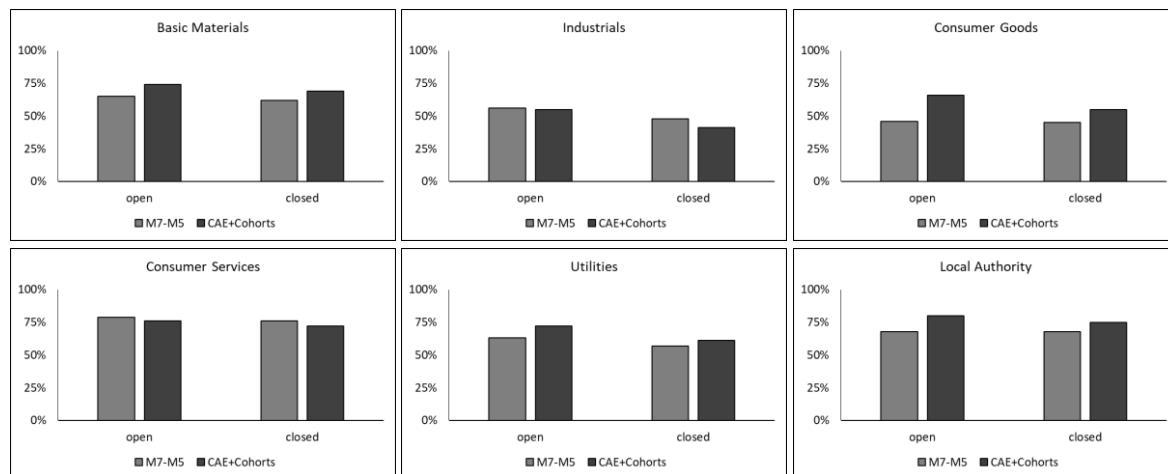

Figure 4.10   **Level of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) in a hypothetical scenario of Australian and New Zealand pensioners of multiple cohorts (males; with swaps over 5-year age buckets)**

50
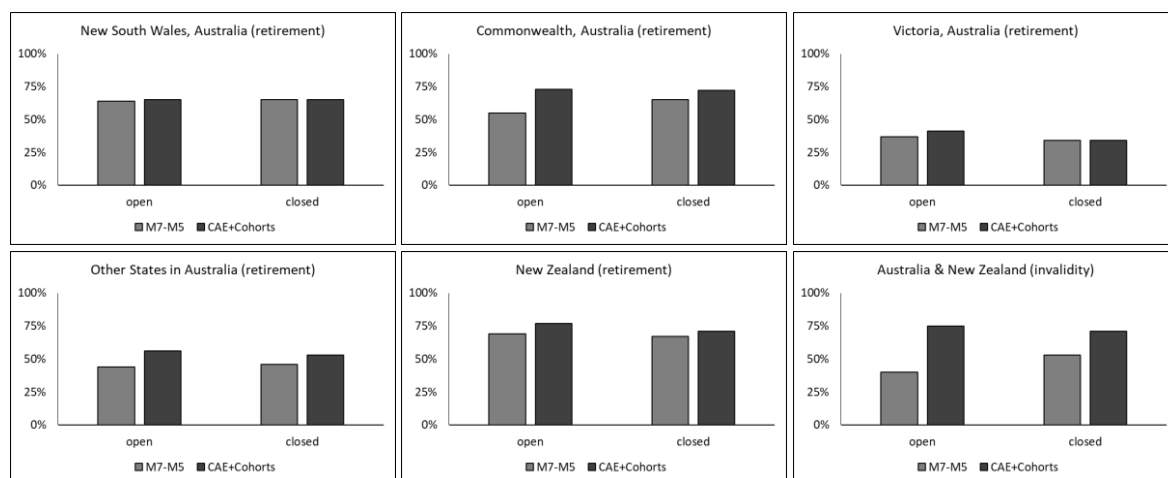
5 Sensitivity Analysis and Scenario Testing

In this section, we conduct an extensive sensitivity analysis on the hedging results from Subsection 4.1 via making a variety of changes to the initial model settings and assumptions. The items in the sensitivity testings include the portfolio size, interest rate, pension feature, type of hedging instrument, hedging strategy, age range for modelling, data fitting period, simulation method, model choice, and other additional features like mortality structural changes and mortality jumps. Whenever a change is made to an original setting, other things are kept equal unless otherwise specified. Through these testings, we can obtain a better picture of how robust (or volatile) the hedging results are under different circumstances. Lastly, we also perform some backtesting and scenario testing on the longevity hedging strategy. A sensitivity analysis on time series processes is given in Section 6.

5.1 Portfolio size

As noted earlier, in general, for a smaller initial portfolio size, sampling basis risk in the future simulations is higher and so the calculated level of longevity risk reduction is lower. Indeed, as shown in Table 5.1 (Figure 5.1), the longevity risk reduction estimates remain more or less the same (all above 60%) until the portfolio size drops to about 25,000, only under which the decline in the risk reduction estimates becomes more obvious. For a portfolio size of 2,500, the risk reduction estimates are around 50%. When there are just 1,000 lives in the portfolio, the extent of risk reduction is only about 30% to 40%. Based on all our simulated results so far, sampling basis risk can pose a significant, negative effect on longevity hedging when there are merely, say, a few thousand pensioners in the plan. These results suggest that index-based longevity hedging would likely be more feasible for sizable pension portfolios, or for larger foundations joined by small pension plans (LCP 2012).

Table 5.1   **Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) for different initial portfolio sizes**

Males Portfolio Size	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
Infinite	66 / 79	63 / 75	63 / 74	64 / 68	71 / 68	70 / 69	73 / 75	68 / 71	63 / 69
100,000	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
50,000	64 / 76	61 / 72	61 / 72	63 / 68	70 / 67	69 / 68	72 / 74	68 / 70	63 / 69
25,000	63 / 73	62 / 68	61 / 67	62 / 64	67 / 63	67 / 63	71 / 72	67 / 68	62 / 67
10,000	61 / 67	61 / 65	59 / 63	59 / 62	64 / 59	62 / 59	68 / 67	63 / 63	58 / 61
5,000	56 / 60	53 / 56	53 / 54	55 / 56	59 / 52	57 / 52	64 / 63	59 / 55	57 / 55
2,500	50 / 49	49 / 47	48 / 46	48 / 48	51 / 46	50 / 48	59 / 54	56 / 48	52 / 47
1,000	37 / 34	36 / 37	38 / 36	37 / 34	38 / 33	40 / 34	48 / 37	42 / 33	38 / 33

Figure 5.1   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) for different initial portfolio sizes (males)
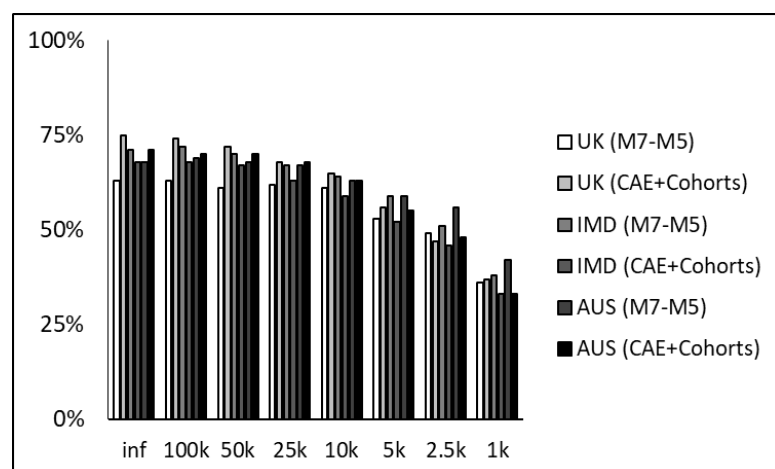
5.2   Interest rate

Table 5.2 (Figure 5.2) illustrates that generally, the higher the interest rate, the smaller the risk reduction estimates regarding the present value of the aggregate position. As the interest rate increases, the later individual cash flows, having greater risk reduction effect, are discounted more than the earlier individual cash flows, which have less risk reduction effect. Hence the extent of risk reduction for the overall present value decreases. Moreover, under the M7-M5 model, the differences in the risk reduction effect between the earlier and later cash flows are usually larger, which may explain the faster drop in the overall risk reduction as the interest rate rises. From an interest rate of 1% p.a. to 5% p.a., the changes in the risk reduction estimates given in the table are no more than 9% in magnitude and are not too significant.

In addition, the hedging results under a variable interest rate environment simulated from the discretised Cox-Ingersoll-Ross (CIR) model are also provided in the table. This interest rate model has the advantages of allowing for mean reversion and also avoiding negative interest rates (Cairns 2004). The most recent few years of historically low interest rates are used to calibrate the CIR model. Under variable interest rates, the risk reduction estimates drop further, especially under the CAE+Cohorts model, reflecting the presence of interest rate risk. While it is difficult to predict how long the current low interest rate levels would continue, it can be envisaged that higher interest rates with more fluctuations would reduce the hedge effectiveness to even a greater extent. Accordingly, interest rate swaps and government bonds may be added to the hedging scheme to mitigate the impact of interest rate risk (Tsai et al. 2011). For an insurer operating under Solvency II, however, the discount rate is prescribed and is based on the market swap rate, and this way of hedging interest rate would not produce the desired outcome in evaluating risk reduction.

Table 5.2   Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) under different interest rate assumptions

100,000 Males Interest Rate (p.a.)	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
1%	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	74 / 76	70 / 71	65 / 69
2%	63 / 77	61 / 73	60 / 73	63 / 69	69 / 67	68 / 67	73 / 75	69 / 71	64 / 69
3%	62 / 76	59 / 71	58 / 71	62 / 68	69 / 67	69 / 68	72 / 74	69 / 70	63 / 68
4%	60 / 75	58 / 72	57 / 71	62 / 68	68 / 67	67 / 68	71 / 74	67 / 69	62 / 68
5%	58 / 74	55 / 71	54 / 71	61 / 68	67 / 66	66 / 67	70 / 73	65 / 68	61 / 67
CIR	53 / 52	55 / 55	53 / 55	57 / 57	62 / 58	61 / 59	66 / 61	64 / 59	61 / 57

Figure 5.2   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) under different interest rate assumptions (100,000 males)
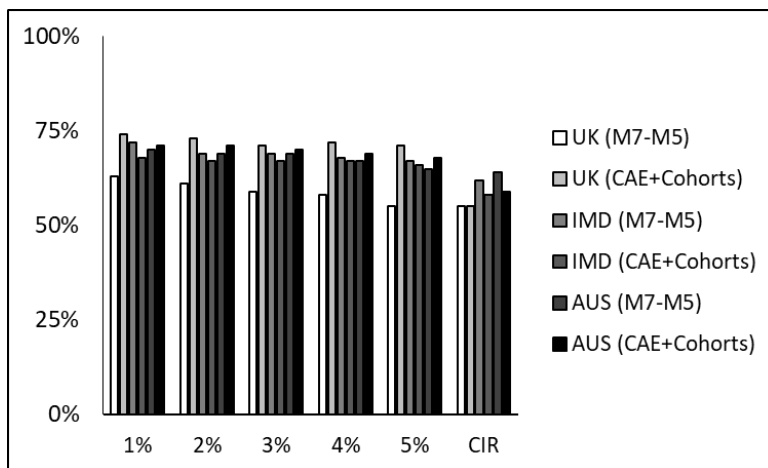
5.3   Pension feature

Table 5.3 (Figure 5.3) presents the hedging results for three different pensions. The first one pays £1 per year on survival from ages 66 to 90, the second one from ages 71 to 90, and the third one from ages 76 to 90. The pensioners are currently aged 65, 70, and 75 respectively. Assume that longevity swaps for the same birth cohorts as the pensioners are available, in which the maturities are 25 years, 20 years, and 15 years in these three cases. It can be seen that the risk reduction estimates are progressively smaller for those cases with shorter durations. This effect can be explained by the fact that under the current model settings and assumptions, the later the individual cash flows, the greater the risk reduction effect, and so those pensions with longer durations can achieve greater hedge effectiveness.

Table 5.3   Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) for different age ranges of payments

100,000 Males Pay Age Range	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
66-90	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
71-90	64 / 75	63 / 71	61 / 70	62 / 66	69 / 65	68 / 65	72 / 73	67 / 69	62 / 67
76-90	62 / 69	61 / 62	59 / 60	60 / 61	65 / 59	64 / 59	70 / 68	64 / 65	61 / 61

Figure 5.3   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) for different age ranges of payments (100,000 males)
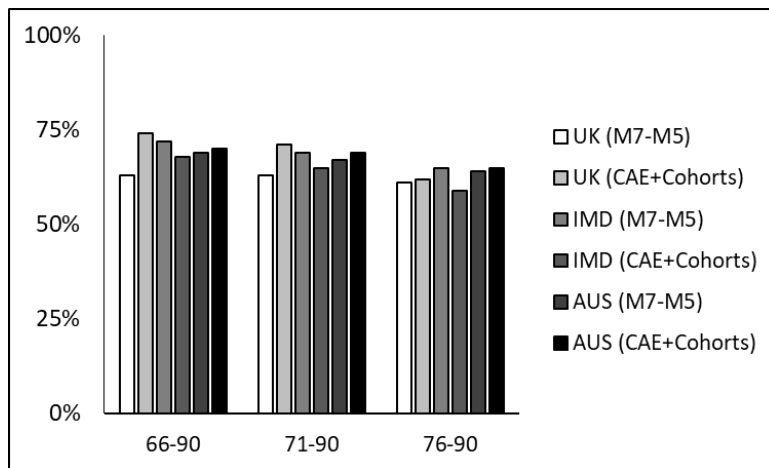
5.4   Hedging instrument

All the analysis above focuses on the use of standardised longevity swaps. There are other types of securities which can be tested (Table 4.1). One kind is the so-called *q*-forward, in which a payoff based on the difference between the forward mortality rate and the actual mortality rate of the reference population is made on maturity. On the other hand, for a S-forward, the payoff is fairly similar to that of the *q*-forward, but with the mortality rate being replaced by the survivor index, i.e. the percentage of the reference population who are alive. Note that S-forwards can be regarded as the fundamental building blocks for both the longevity swap and longevity bond. Again, numerical optimisation (Appendix I) can be adopted to determine the optimal positions of the hedging instruments. It is assumed that *q*-forwards and S-forwards with different maturities of 1 year to 25 years are available in the life market for the same birth cohort as the pensioners.

Table 5.4 (Figure 5.4) shows that although *q*-forwards with specific weights at different maturities are used in constructing the hedge, the resulting risk reduction is clearly less

effective (about 10% to 20% smaller) than that by using a longevity swap over the entire period. (The swap can simply be seen as a combination of a series of S-forwards with equal weights.) The more direct and natural relationship (in terms of survival probabilities) between the number of surviving pensioners and the survivor index (rather than the mortality rate) may explain the difference. In contrast, the risk reduction effect is slightly larger when the hedge is built on individual S-forwards with different weights at specific maturities, which allow more flexibility in dealing with the individual cash flows. When the life market is still immature, however, the availability of q-forwards and S-forwards may be limited to only certain 'key' ages and maturities, and the hedge would then be less effective than otherwise.

Table 5.4   Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) using different hedging instruments

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Hedging Instrument	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
Swap	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
q-forwards	56 / 62	54 / 58	52 / 59	47 / 56	47 / 52	48 / 54	60 / 62	56 / 56	53 / 54
S-forwards	65 / 77	63 / 74	63 / 73	64 / 69	71 / 68	70 / 69	74 / 76	70 / 71	66 / 69

Figure 5.4   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) using different hedging instruments (100,000 males)
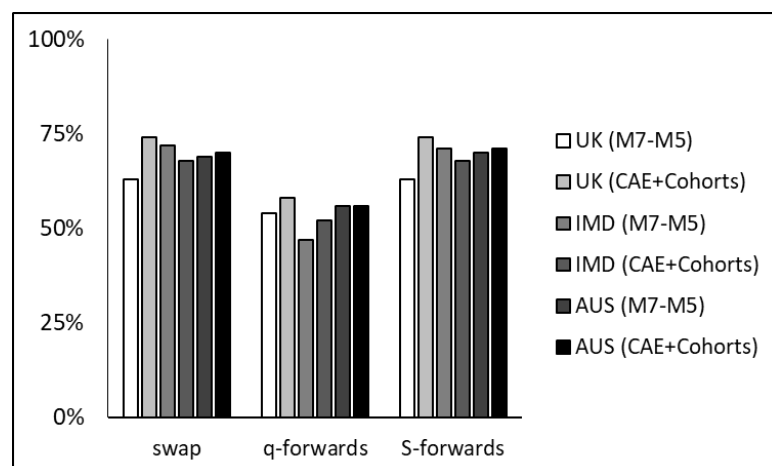
5.5   Hedging strategy

Thus far, the calculation of the optimal positions and hedging results has been based on simulated environments from either the fitted M7-M5 model or CAE+Cohorts model and numerical optimisation with respect to the 99.5% value-at-risk (VaR). It would be informative to examine some additional scenarios, for instance, the standard deviation or the

99.5% expected shortfall (conditional VaR) is taken as the risk measure instead in the optimisation process. Another one is that the swap size is simply equated to the initial portfolio size (i.e. a one-to-one hedge ratio), or that the swap weight is calculated from a 'wrong' model which is different to the model being used to generate the simulations.

As shown in Table 5.5 (Figure 5.5 – left graph), the differences are minimal between using different risk measures in the optimisation process. This observation may be a result of the distribution of the portfolio present value, which is fairly symmetric and does not demonstrate much skewness (Figure 5.6). Moreover, using a hedge ratio of one, the resulting differences in risk reduction are very small for the UK pensioners and the IMD quintile group, but are more obvious for those Australian pensioners. This result can be explained by the fact that the numerically optimised hedge ratio is close to one for the two UK groups but it is not the case for the Australian group. Overall, it seems that using simply a hedge ratio of one, or even the weight from a 'wrong' model, can still provide some 'basic' level of longevity risk hedging, while more precise modelling and hedging strategy could lead to better calculated hedge effectiveness, depending on the underlying data being modelled.

Besides numerically optimising the use of a swap, another calibration method called key q-duration (KQD) matching, proposed by Li and Luo (2012), may also be used when q-forwards are employed to build the hedge. An important and widely applied concept, bond duration, can be borrowed from the finance literature. Broadly speaking, the bond duration has two general meanings: first, it is the weighted average time of all future cash flows; at the same time, it also measures how the present value of future cash flows varies when the interest rate changes. In the context of hedging longevity risk, the KQD is defined as:

$$KQD = \frac{\text{resulting change in best estimate of present value}}{\text{small change in best estimate of a particular future mortality rate}}$$

The hedging positions can be determined by equating the KQDs of the hedging instruments and the corresponding KQD of the portfolio being hedged. This calibration procedure is similar to that in hedging a bond portfolio from interest rate fluctuations. The KQD has the advantages of being easy to understand and straightforward to apply. More technical details are provided in Appendix I.

The figures in Table 5.5 (Figure 5.5 – right graph) suggest that the KQD strategy can yield comparable results to those from numerically optimising the use of a longevity swap. Assume that q-forwards with maturities of 1 year to 25 years are available for the same birth cohort as the pensioners. Using these q-forwards, the risk reduction estimates regarding the present value of the aggregate position from applying the KQD are in fact mostly more than 10% larger in magnitude than those from optimising the individual cash flows, as the KQD is specifically designed for the present value. If q-forwards are then made available only at key ages of 70, 75, 80, and 85 (i.e. maturities of 6, 11, 16, and 21 years), the risk reduction estimates decrease only by less than 10%. However, we discover that although the KQD

strategy on *q*-forwards can produce decent hedge effectiveness with regard to the present value of the aggregate position, it does not necessarily lead to proper risk reduction effect on the individual cash flows, in which there are often very mixed results and offsetting effects between the cash flows in different years. Further research is required to examine the practical limitations of this method.

Table 5.5   **Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) using different hedging strategies**

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Longevity Swap	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
Optimisation (99.5% VaR)	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
Optimisation (99.5% ES)	63 / 77	63 / 74	63 / 73	64 / 68	71 / 67	70 / 69	72 / 75	67 / 70	63 / 68
Optimisation (standard dev)	65 / 77	63 / 74	63 / 72	64 / 69	71 / 67	70 / 68	72 / 75	69 / 70	63 / 68
One-to-one	65 / 77	62 / 74	62 / 74	63 / 66	70 / 65	69 / 66	67 / 64	62 / 52	57 / 52
Wrong model	65 / 77	62 / 72	61 / 71	63 / 67	70 / 67	69 / 68	70 / 74	66 / 68	62 / 66
100,000 Males	**UK Pensioners (basic materials; normal retirement; lower pension)**			**IMD Groups (most deprived areas)**			**Australian Pensioners (New South Wales; retirement)**		
q-forwards	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
Optimising CFs	56 / 62	54 / 58	52 / 59	47 / 56	47 / 52	48 / 54	60 / 62	56 / 56	53 / 54
KQD all ages	65 / 77	61 / 73	61 / 72	64 / 69	70 / 68	69 / 67	74 / 77	68 / 71	63 / 70
KQD key ages	62 / 72	60 / 67	58 / 65	59 / 64	63 / 59	62 / 59	70 / 72	62 / 63	58 / 62

Figure 5.5   **Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) using different hedging strategies (100,000 males)**
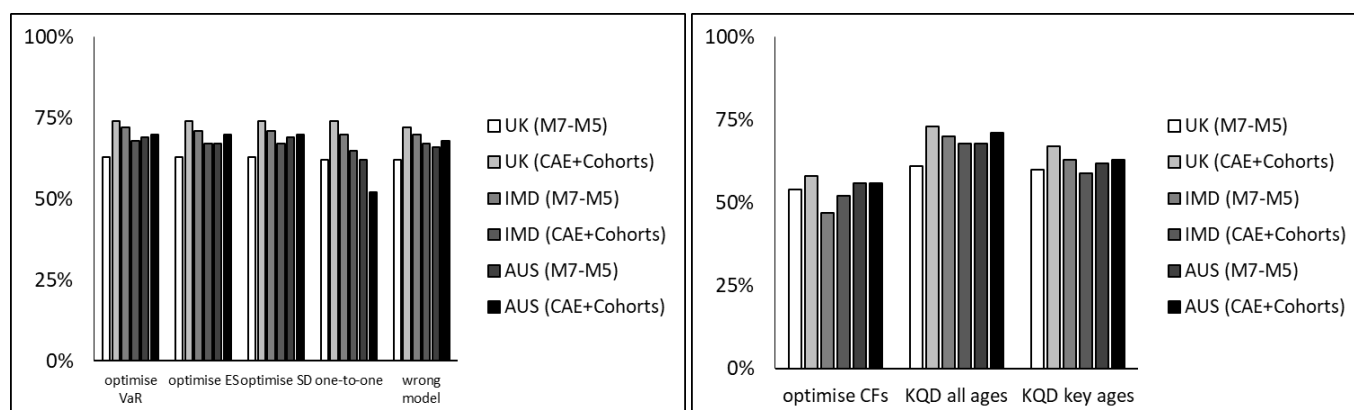
Figure 5.6   Simulated density of pension portfolio present value under M7-M5 model and original settings (males)
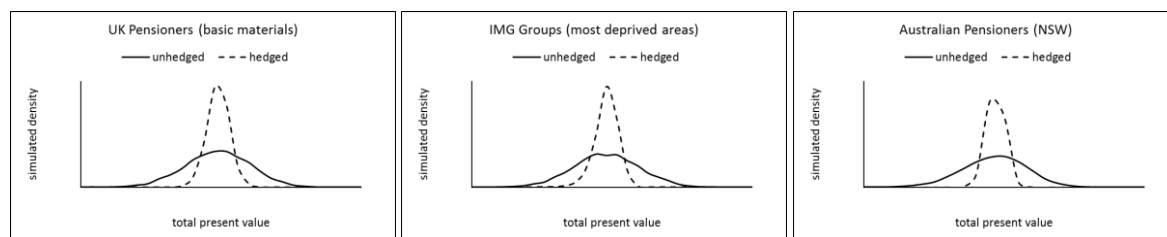
5.6   Age range

Both Phase 1 and Phase 2 focus on the age range of 60 to 89. The major problem of modelling older ages is that the data are scarce and the experience is volatile. For comparison purposes, we still extend the analysis to cover a wider age range of 60 to 99. The ONS population data are not split by single age for ages 90 and beyond and so the corresponding HMD proportions of those aged 90+ are used as an approximation here. First, the two models need to be re-fitted to this new age range. Then suppose that a few different pensions pay £1 per year on survival from ages 66 to 90 for those pensioners currently aged 65 (like previously), from ages 66 to 100 for those now aged 65, from ages 76 to 100 for those now aged 75, and also from ages 86 to 100 for those aged 85 at present. Assume that longevity swaps for the same birth cohorts as the pensioners are available with maturities of 25, 35, 25, and 15 years respectively. The results in the first two rows in Table 5.6 suggest that the models are fairly robust to an inclusion of advanced ages in the modelling process, except for some figures of the IMD quintile group, which may be due to the use of approximate data for ages 90+. Moreover, it can be seen again (Figure 5.7) that the risk reduction levels tend to be lower for those pensions with shorter durations. As discussed earlier, the later individual cash flows have greater risk reduction effect under the current model settings and assumptions.

Table 5.6   Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) using older ages in modelling

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Pay Age Range	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
66-90 (old)	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
66-90 (new)	64 / 78	62 / 76	59 / 74	52 / 59	60 / 62	58 / 63	73 / 75	63 / 73	62 / 72
66-100	76 / 86	76 / 85	74 / 84	56 / 61	69 / 65	69 / 66	81 / 84	76 / 82	74 / 81
76-100	74 / 80	73 / 78	72 / 76	56 / 58	65 / 62	64 / 63	80 / 81	75 / 78	74 / 78
86-100	65 / 64	63 / 63	63 / 59	55 / 53	60 / 55	61 / 57	74 / 71	69 / 70	69 / 70

Figure 5.7   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) using older ages in modelling (100,000 males)
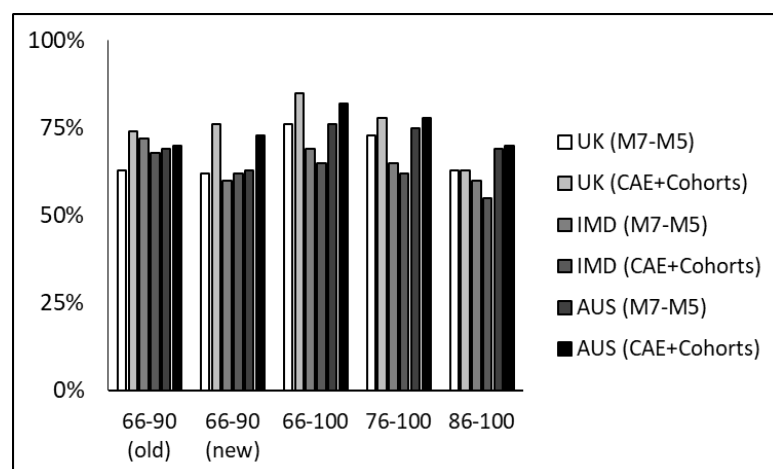
5.7   Data period

In order to test further the robustness of the modelling results, we re-fit the two models to a shorter data period for the reference population. It may have the advantage of allowing for more recent and relevant patterns, more in line with the book data which only start from year 2000 or later. But it also reduces the sample size and there may be more sampling variability. Table 5.7 (Figure 5.8) demonstrates that the risk reduction estimates do not differ much between the two fitting periods starting from years 1980 and 1990 (mostly within 5%). This observation may reflect that the reference population experience is relatively stable, the potential benefit and drawback of using a shorter data period largely offset each other, or the models are robust themselves. We also re-fit the models to a slightly shorter data period for the book population. The resulting risk reduction levels are again not very different (all by 6% or less). It would be insightful to consider more different data periods for the book population, but since its length is only a little over 10 years, more data need to be collected in order to conduct a more thorough testing.

Table 5.7   Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) using different data fitting periods

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Data Period	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
From 1980 (reference)	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
From 1990 (reference)	63 / 76	62 / 70	58 / 68	62 / 66	74 / 65	71 / 65	64 / 71	60 / 67	58 / 65

Table 5.7   Continued

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Data Period	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
From 1 year later (book)	64 / 77	64 / 76	65 / 75	64 / 70	72 / 70	73 / 70	71 / 72	68 / 70	66 / 65
From 2 years later (book)	63 / 76	63 / 73	61 / 72	69 / 73	77 / 74	76 / 74	66 / 73	68 / 71	61 / 68

Figure 5.8   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) using different data fitting periods (100,000 males)
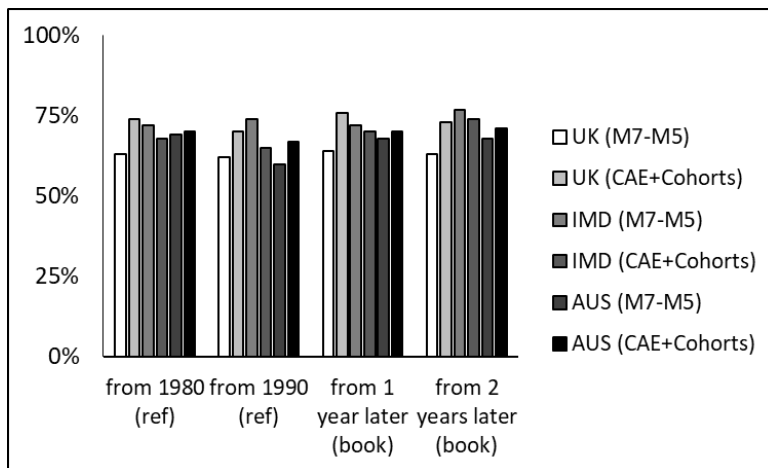
5.8   Simulation method

As noted in Subsection 3.4, residuals bootstrapping (Figure 3.3) can be adopted to include both process error (variability in the time series) and parameter error (uncertainty in parameter estimation). A simpler, parametric method is to perform Monte Carlo simulation on the error terms in the fitted time series processes to generate random future mortality rates. This way, however, considers only process error but not parameter error, and may underestimate longevity basis risk. The risk reduction estimates in Table 5.8 (Figure 5.9) suggest that the parametric method does understate longevity basis risk and so overestimates the extent of risk reduction by around 8% to 20% in magnitude. Note that with Excel VBA, the computation time of residuals bootstrapping for each single case of simulation can be up to a day or so, while the parametric method generally takes only a few hours. An alternative is to use the software R, but the computation time is usually about threefold in our simulations. Some potentially faster computation algorithms based on approximation methods are also provided in Appendix I.

Table 5.8   **Levels of longevity risk reduction (in % of initial longevity risk; M7-M5 vs CAE+Cohorts) using different simulation methods**

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Simulation Method	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
Bootstrapping	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
Parametric	73 / 87	73 / 86	72 / 84	84 / 83	84 / 82	84 / 81	84 / 88	84 / 86	83 / 86

Figure 5.9   **Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) using different simulation methods (100,000 males)**
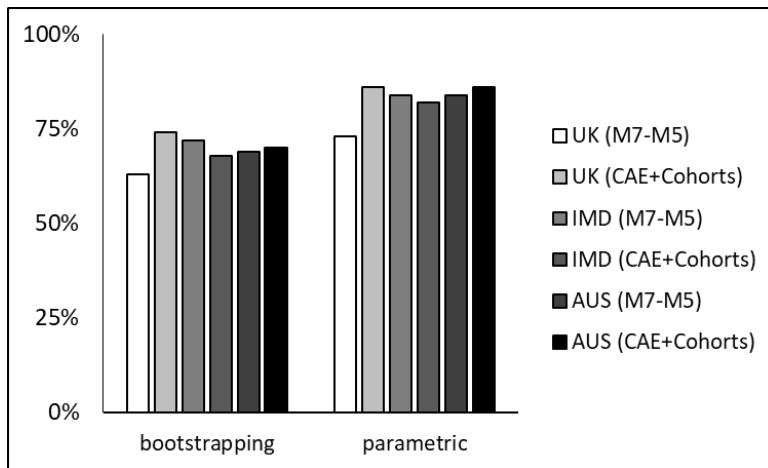
5.9   Model choice

Phase 1 has proposed a decision tree framework (Figure 3.1) on how to choose a two-population model for the reference and book populations, which includes the M7-M5 model, the CAE+Cohorts model, and the characterisation approach. These models are selected from a long list of potential candidates after an extensive investigation using several UK datasets. Besides process error and parameter error, the Solvency II Delegated Regulations (CEIOPS 2010) also requires an incorporation of model error (uncertainty in the model choice) in assessing uncertainty of future outcomes. The extent of model error can be inspected by comparing the hedging results under different models and assumptions. It can be noted from all the simulations above that while the M7-M5 and CAE+Cohorts models do not generate too different risk reduction estimates for larger data sizes and under more precise hedging, the extent of model error is not negligible otherwise. Accordingly, it would be helpful to test some other models for further comparison.

Three additional models are thus being considered here. The first one is modified from the first approach stated in Carter and Lee (1992), which fits the Lee and Carter (1992) model to each population and then captures the relationship between the two mortality indices:

$$\text{logit}\, q_{x,t}^R = \alpha_x^R + \beta_x^R \,\kappa_t^R + \gamma_{t-x}^R \qquad \text{(reference population)}$$

$$\text{logit}\, q_{x,t}^B = \alpha_x^B + \beta_x^B \,\kappa_t^B + \gamma_{t-x}^R \qquad \text{(book population)}$$

$$K_t = \Theta + K_{t-1} + \Delta_t \qquad \text{(bivariate random walk with drift, BRWD)}$$

It is assumed that $K_t = \left(\kappa_t^R, \kappa_t^B\right)'$, Θ is the vector drift term, Δ_t is the bivariate normal error term, and the other parameters are defined similarly as in Subsection 3.2. The second one is extended from the third approach in Carter and Lee (1992), which models the two mortality indices as a co-integrated process (Li and Hardy 2011):

$$\text{logit}\, q_{x,t}^R = \alpha_x^R + \beta_x^R \,\kappa_t^R + \gamma_{t-x}^R \qquad \text{(reference population)}$$

$$\text{logit}\, q_{x,t}^B = \alpha_x^B + \beta_x^B \,\kappa_t^B + \gamma_{t-x}^R \qquad \text{(book population)}$$

$$\kappa_t^R = \theta + \kappa_{t-1}^R + \delta_t \qquad \text{(random walk with drift, RWD)}$$

$$\kappa_t^B = a_0 + a_1 \,\kappa_t^R + \omega_t \qquad \text{(co-integrated process)}$$

The parameters θ, a_0, and a_1 govern the co-integrated process, and δ_t and ω_t are independent normal error terms. The final model is an extension of the idea proposed by Zhou, Li, and Tan (2013), which assumes a common age-specific sensitivity measure for the reference and book populations and then a weakly stationary autoregressive AR(1) process for the difference between the two resulting mortality indices:

$$\text{logit}\, q_{x,t}^R = \alpha_x^R + \beta_x^R \,\kappa_t^R + \gamma_{t-x}^R \qquad \text{(reference population)}$$

$$\text{logit}\, q_{x,t}^B = \alpha_x^B + \beta_x^R \,\kappa_t^B + \gamma_{t-x}^R \qquad \text{(book population)}$$

$$\kappa_t^R = \theta + \kappa_{t-1}^R + \delta_t \qquad \text{(random walk with drift, RWD)}$$

$$\kappa_t^R - \kappa_t^B = b_0 + b_1\left(\kappa_{t-1}^R - \kappa_{t-1}^B\right) + \omega_t \qquad \text{(AR(1) process)}$$

The notation θ, b_0, and b_1 are the parameters of the time series processes, and δ_t and ω_t are independent normal error terms.

It is important to recognise that the first two models are 'non-coherent' but the last model is 'coherent' (Li and Lee 2005). That is, the projected (best or central estimate) ratio of future mortality rates between the book and reference populations at each age converges to a constant in the long term under the last model but not for the first two. As noted previously, one could argue that the two assumingly related populations' mortality levels may diverge between their trends in the short run but would then move more closely in line over the long run. Failure to allow for this long-term coherence may overestimate longevity basis risk. Besides the coherence of the best estimates, it is also important to consider carefully the simulated future variability. For the first two models, both kappas of the reference and book populations have unbounded future variability. But for the last model, while the kappa of the reference population also has unbounded future variability, the difference between the kappas

of the reference and book populations has bounded future variability. The implication is that the two populations' future mortality movements could deviate more significantly under the first two models (particularly in the long term), but in contrast they would be more consistent under the last model.

Table 5.9 (Figure 5.10) below sets forth the risk reduction estimates calculated from the M7-M5 and CAE+Cohorts models and also the three additional models described above. It appears that the non-coherence feature and unbounded future variability of the first two additional models lead to very different hedging results, in which there is a significant overestimation of longevity basis risk and so the risk reduction estimates are much smaller. Comparatively, the M7-M5 model, the CAE+Cohorts model, and the last additional model are all coherent and have bounded future variability between the book and reference populations, and their risk reduction estimates reflect a more proper assessment of longevity basis risk. In particular, the differences between the results from these three models look quite small (within around 10%). These observations suggest that despite the obvious differences in the structures between the various types of models, the coherence property and the behaviour of simulated future variability are the major factors underlying the calculation of longevity risk reduction. Further analysis on time series modelling and simulated future variability are provided in Section 6.

Table 5.9   Levels of longevity risk reduction (in % of initial longevity risk) under different two-population models

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Model	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
M7M5	65	63	63	64	72	70	72	69	63
CAE	77	74	73	69	68	68	74	70	68
Extra model 1	7	8	4	28	22	22	3	5	5
Extra model 2	45	52	51	53	57	55	0	7	4
Extra model 3	77	74	74	67	64	64	75	70	69

5.10   Additional feature

All the analysis so far is based on the data fitting period starting from 1980 or later. There are more than 30 years of reference data and over 10 years of book data, and the projection period is 25 years. However, Phase 1 has stated that their research has focused on only up to a 15-year forecasting horizon and evaluating hedge effectiveness over longer horizons requires further analysis and more caution. In fact, mortality levels have not always been progressing so smoothly over time. There could be one-off incidents such as wars, catastrophes, and

epidemics, causing temporary drastic changes in mortality levels (i.e. mortality jumps). There could also be long-term effects like radical medical advances and climate changes, leading to permanent changes in mortality improvement rates (i.e. structural changes). For instance, Figure 5.11 plots the mortality indices (Subsection 3.2) of the English and Welsh and Australian male populations. It can clearly be seen that there are a few 'mortality spikes' (in red) before 1950, and that there is a major shift in the improvement trend (in blue) during around 1960 to 1970.

Figure 5.10   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) under different two-population models (100,000 males)
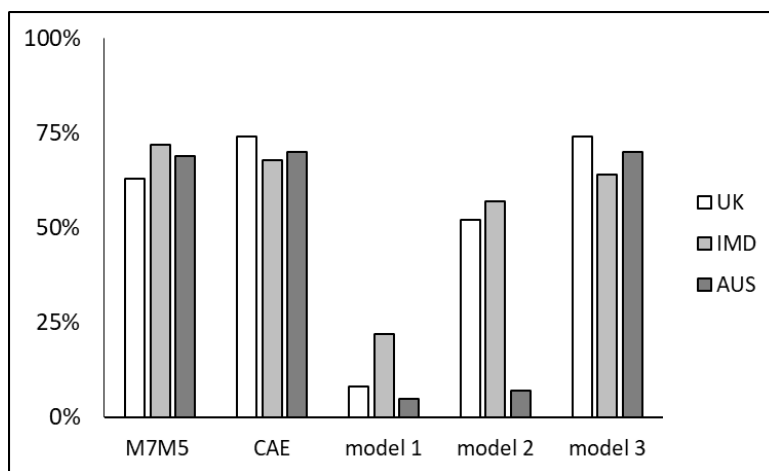

Figure 5.11   Mortality indices of England and Wales and Australia (males)
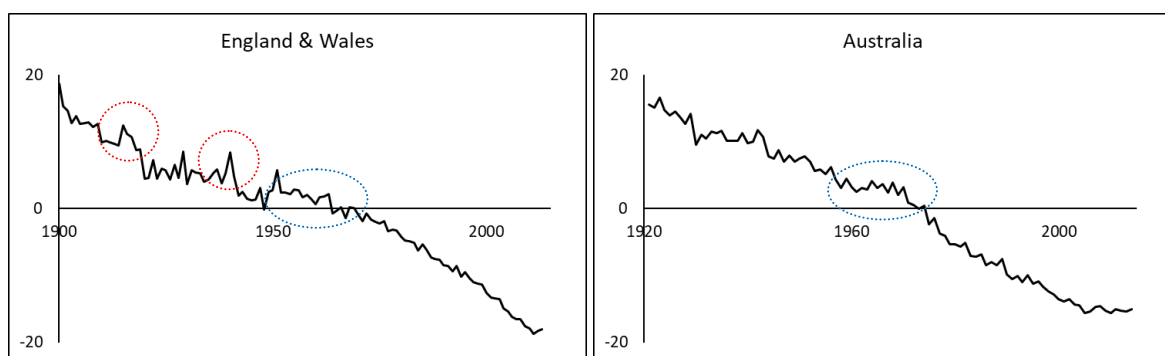
These events are either rare or hard to identify precisely, in which mortality data are usually available only annually for a limited period of time. So it is difficult to allow for them in the modelling process without taking references from other relevant data and related fields of studies, applying sound biological reasoning, and exercising appropriate judgement. Based on the data used in Figure 5.11 and also several other developed countries' population data, we consider three extra scenarios here and make some arbitrary assumptions – the first one with possible structural changes, the second one with possible mortality jumps, and the final one with both. We modify the last additional model in the previous subsection to allow for these

effects, which are implicitly assumed to have the same impact on both the reference and book populations:

$$\kappa_t^{R*} = \theta_t + \kappa_{t-1}^{R*} + \delta_t \qquad \text{(RWD modified with variable drift } \theta_t \text{, i.e. structural changes)}$$

$$\kappa_t^{R} = \kappa_{t-1}^{R*} + N_t\, Y_t \qquad \text{(RWD further modified with mortality jumps } N_t\, Y_t \text{)}$$

$$\theta_t = (2\theta, \theta, \theta/2)' \qquad \text{(structural changes' transition matrix } \begin{pmatrix} 0.99 & 0.01 & 0 \\ 0.01 & 0.98 & 0.01 \\ 0 & 0.01 & 0.99 \end{pmatrix} \text{)}$$

$$\Pr(N_t = 0) = 0.99 \quad \Pr(N_t = 1) = 0.01 \quad \text{(frequency of a mortality jump)}$$

$$Y_t \sim \mathrm{N}(30, 10^2) \qquad \text{(severity of a mortality jump)}$$

The probabilities and possible sizes of these events are roughly determined by inspecting the frequencies and severities of relevant historical incidents such as the World Wars, Spanish Flu, and past structural changes in mortality improvement. Users can seek other resources and experts' opinions in setting these assumptions. In addition, one may further assume that such events can have different extents of impact on the reference and book populations.

Table 5.10 shows that the incorporation of structural changes increases the risk reduction estimates regarding the tail risk measures by about 10% or more in magnitude, while the inclusion of mortality jumps does not lead to any obvious changes in risk reduction. As structural changes are assumed to have a long-lasting impact on the reference and book populations simultaneously in the model, the two populations' mortality levels would then move more consistently with each other over a long period of time, resulting in some further reduction in longevity basis risk (Figure 5.12). On the other hand, the effects of mortality jumps are assumed to be temporary and one-off, and accordingly the influence on the risk reduction levels is minimal.

Table 5.10   Levels of longevity risk reduction (in % of initial longevity risk) using models with additional features deduced from historical incidents

100,000 Males	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
Additional Feature	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
M7M5 / CAE	65 / 77	63 / 74	63 / 73	64 / 69	72 / 68	70 / 68	72 / 74	69 / 70	63 / 68
Extra model 3	77	74	74	67	64	64	75	70	69
Structural	83	85	84	74	80	79	78	81	80
Jump	80	72	71	70	64	65	71	68	66
Both	91	91	91	77	78	78	74	74	73

Figure 5.12   Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) using models with additional features deduced from historical incidents (100,000 males)
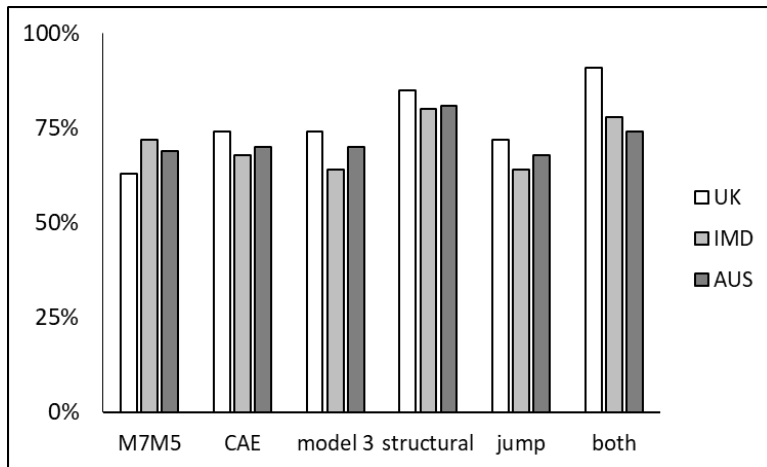
Based on all the results thus far, the most significant modelling assumptions and settings include the coherence property and behaviour of simulated future variability, portfolio size, data size and characteristics, type of hedging instrument, simulation method, and certain additional model features. Comparatively, the other conditions tested seem to have rather limited influence on the calculated hedging results. In constructing an index-based longevity hedge for a pension or annuity portfolio, it is very important to consider these factors carefully and conduct sufficient testing on their potential impact upon the hedge effectiveness that can be achieved.

5.11   Scenario testing

Besides the various sensitivity testings above, backtesting can also be performed on the longevity hedging strategy using relevant historical data to examine how well the strategy would have worked in the past. While the HMD dataset is dated back to a long time ago, the other three datasets cover a period of only around 10 years. In order to conduct a more meaningful backtesting exercise under this data constraint, a proxy is adopted here, in which the book population is assumed to follow the aggregate CMI mortality experience from 1983 to 2006. Though this aggregate CMI dataset is not sorted by the industry and pension amount and does not cover the most recent period, there are 24 years of data in total, which allows certain backtesting on some simple cases. (Note that the data have been used earlier by the first author in a few published journal articles.) The data are divided into two periods, in which the first period is used for fitting and the second for testing. The interest rate is assumed to be 5% p.a. in this analysis, more in line with the historical interest rate levels. It is also assumed that the portfolio size is very large and so sampling basis risk is minimal.

The difference between the actual present value and the expected (best estimate) present value (at the start of the testing period; as a percentage of the expected present value) for each case is given in Table 5.11. A positive figure (under the column 'unhedged') means that the overall mortality improvement is greater than anticipated and so the actual present value of the pension portfolio turns out to be larger than the expected value, i.e. there is an unexpected portfolio loss. It is observed (from the rows noted as 'actual') that the losses, if the portfolio is unhedged, are mostly around 1% to 3% only. This result reflects that the mortality experience is rather stable during these few decades and the main mortality trends are quite reasonably captured by the two models, despite a small degree of underestimation. Consequently, the longevity hedge would not impose much effect on the pension portfolio.

Further scenario testing can be conducted to illustrate the potential effect of the longevity hedge by considering more extreme situations. As per the Solvency II Standard Formula (CEIOPS 2010), the capital requirement for longevity risk is defined as the change in the net asset value under a longevity scenario in which there is a permanent 20% decrease in mortality rates (i.e. a longevity shock) for each age and each policy. Accordingly, nine different pairs of 'longevity shocks' are applied to the actual data in the testing period, ranging from a permanent 10% decrease to a 30% decrease for the reference and book populations. Table 5.11 shows that there is a wide range of hedging outcomes, from around 1% to 14% reduction in the portfolio loss in magnitude after taking the hedge. A few major patterns can be identified here. Firstly, as expected, the greater the extent of the shocks, the larger the portfolio loss, and the more obvious the corresponding hedging effect. The longevity hedge clearly comes into effect when there are substantial unexpected mortality improvements. Secondly, the portfolio loss and the hedging effect are greater for the older age range. This observation may be due to the fact that the mortality levels at ages 60 to 69 have already declined substantially over the years, while those aged 80 and above have more room for further improvement, which has a larger impact on the present value of a pension. So the hedging outcome is more noticeable for older ages. Finally, when the shock levels are different between the reference and book populations, which is a reflection of longevity basis risk, the longevity hedge still appears to provide a decent protection for the pension portfolio, cutting down at least one third of the portfolio loss in most of the cases shown.

In summary, the analysis above suggests that the longevity hedging strategy works reasonably well when there are considerable unanticipated mortality improvements. On the other hand, if the major mortality trends are well captured in the modelling process and these trends continue to exist, the longevity hedge would not bring much impact to the pension portfolio. In addition, when the longevity shock is greater for the book population than for the reference population, due to the presence of longevity basis risk, the reduction in the portfolio loss from the hedge would still be sizable, as long as the shocks on the two populations are in the same direction.

Table 5.11   Backtesting and scenario testing on longevity hedging strategy – difference (in %) between actual and expected present values when the pension portfolio is unhedged or hedged (M7-M5 vs CAE+Cohorts)

Males		Difference (in %) between Actual and Expected Present Values	
Reference Scenario	Book Scenario	Unhedged	Hedged
E&W (fitting 1970-1992 / testing 1993-2006)			
CMI (fitting 1983-1992 / testing 1993-2006)			
Pay age range 66-79			
actual	actual	3.2 / 1.7	2.5 / 2.2
10% ↓	10% ↓	4.7 / 3.1	2.2 / 1.9
10% ↓	20% ↓	6.1 / 4.5	3.7 / 3.4
10% ↓	30% ↓	7.6 / 6.0	5.1 / 4.8
20% ↓	10% ↓	4.7 / 3.1	0.5 / 0.2
20% ↓	20% ↓	6.1 / 4.5	1.9 / 1.7
20% ↓	30% ↓	7.6 / 6.0	3.4 / 3.1
30% ↓	10% ↓	4.7 / 3.1	-1.3 / -1.5
30% ↓	20% ↓	6.1 / 4.5	0.1 / -0.1
30% ↓	30% ↓	7.6 / 6.0	1.6 / 1.4
E&W (fitting 1970-1996 / testing 1997-2006)			
CMI (fitting 1983-1996 / testing 1997-2006)			
Pay age range 66-75			
actual	actual	1.5 / 1.0	0.8 / 1.0
10% ↓	10% ↓	2.4 / 1.8	0.5 / 1.1
10% ↓	20% ↓	3.2 / 2.7	1.3 / 1.9
10% ↓	30% ↓	4.1 / 3.5	2.2 / 2.8
20% ↓	10% ↓	2.4 / 1.8	-0.8 / 0.3
20% ↓	20% ↓	3.2 / 2.7	0.1 / 1.1
20% ↓	30% ↓	4.1 / 3.5	1.0 / 2.0
30% ↓	10% ↓	2.4 / 1.8	-2.0 / -0.6
30% ↓	20% ↓	3.2 / 2.7	-1.1 / 0.3
30% ↓	30% ↓	4.1 / 3.5	-0.3 / 1.1

Table 5.11   Continued

Males		Difference (in %) between Actual and Expected Present Values	
Reference Scenario	Book Scenario	Unhedged	Hedged
E&W (fitting 1970-1992 / testing 1993-2006)			
CMI (fitting 1983-1992 / testing 1993-2006)			
Pay age range 76-89			
actual	actual	0.6 / -1.5	-0.2 / 0.0
10% ↓	10% ↓	4.6 / 2.5	-0.4 / -0.1
10% ↓	20% ↓	8.9 / 6.7	3.9 / 4.1
10% ↓	30% ↓	13.5 / 11.1	8.5 / 8.5
20% ↓	10% ↓	4.6 / 2.5	-4.9 / -4.5
20% ↓	20% ↓	8.9 / 6.7	-0.6 / -0.3
20% ↓	30% ↓	13.5 / 11.1	4.0 / 4.1
30% ↓	10% ↓	4.6 / 2.5	-9.7 / -9.2
30% ↓	20% ↓	8.9 / 6.7	-5.4 / -5.0
30% ↓	30% ↓	13.5 / 11.1	-0.9 / -0.6
E&W (fitting 1970-1996 / testing 1997-2006)			
CMI (fitting 1983-1996 / testing 1997-2006)			
Pay age range 76-85			
actual	actual	2.6 / 1.9	0.9 / 2.0
10% ↓	10% ↓	5.3 / 4.6	0.4 / 2.0
10% ↓	20% ↓	8.1 / 7.3	3.2 / 4.7
10% ↓	30% ↓	10.9 / 10.2	6.0 / 7.6
20% ↓	10% ↓	5.3 / 4.6	-2.9 / -0.8
20% ↓	20% ↓	8.1 / 7.3	-0.2 / 1.9
20% ↓	30% ↓	10.9 / 10.2	2.7 / 4.8
30% ↓	10% ↓	5.3 / 4.6	-6.4 / -3.7
30% ↓	20% ↓	8.1 / 7.3	-3.6 / -1.0
30% ↓	30% ↓	10.9 / 10.2	-0.8 / 1.8

6 Choice of Time Series Processes

A careful selection in time series modelling is very important as it determines how future outcomes are simulated from the modelling process. Table 6.1 summarises the time series processes chosen in Phase 1 (Subsection 3.4) and the key implications of these models. Some basic descriptions are also given below. More technical details can be found in standard econometrics references such as Tsay (2002) and also in Appendix I.

Table 6.1   Selected time series processes in Phase 1

Parameter(s)	Time Series Process	Condition(s)	Best Estimate(s)	Variability
M7-M5				
$\kappa_{t,1}^R$, $\kappa_{t,2}^R$, $\kappa_{t,3}^R$	MRWD	-	linear trends	unbounded
γ_c^R	ARIMA(1,1,0)	$\left\vert \phi_1 \right\vert < 1$	long-term linear trend	unbounded
$\kappa_{t,1}^B$, $\kappa_{t,2}^B$	VAR(1)	eigenvalues of $\begin{bmatrix} \varphi_{1,1} & \varphi_{1,2} \\ \varphi_{2,1} & \varphi_{2,2} \end{bmatrix} < 1$ in magnitude	convergence	bounded
CAE+Cohorts				
κ_t^R	RWD	-	linear trend	unbounded
γ_c^R	ARIMA(1,1,0)	$\left\vert \phi_1 \right\vert < 1$	long-term linear trend	unbounded
κ_t^B	AR(1)	$\left\vert \varphi_1 \right\vert < 1$	convergence	bounded

6.1   M7-M5 model

For the M7-M5 model (Subsection 3.1) regarding the reference population, the time series of $\kappa_{t,1}^R$, $\kappa_{t,2}^R$, and $\kappa_{t,3}^R$ are modelled as a multivariate random walk with drift (MRWD):

$$\begin{pmatrix} \kappa_{t,1}^R \\ \kappa_{t,2}^R \\ \kappa_{t,3}^R \end{pmatrix} = \begin{pmatrix} d_1 \\ d_2 \\ d_3 \end{pmatrix} + \begin{pmatrix} \kappa_{t-1,1}^R \\ \kappa_{t-1,2}^R \\ \kappa_{t-1,3}^R \end{pmatrix} + \begin{pmatrix} \varepsilon_{t,1} \\ \varepsilon_{t,2} \\ \varepsilon_{t,3} \end{pmatrix} \qquad \begin{pmatrix} \varepsilon_{t,1} \\ \varepsilon_{t,2} \\ \varepsilon_{t,3} \end{pmatrix} \sim N(0, \Sigma)$$

The parameters d_1, d_2, and d_3 are the drift terms (linear trends' slopes), and $\varepsilon_{t,1}$, $\varepsilon_{t,2}$, and $\varepsilon_{t,3}$ are the multivariate normal error terms with mean zero and covariance matrix Σ. Under the MRWD, the future variability of $\kappa_{t,1}^R$, $\kappa_{t,2}^R$, and $\kappa_{t,3}^R$ increase over time. Moreover, the

time series of γ_c^R is modelled as an autoregressive integrated moving average process, ARIMA(1,1,0):

$$\gamma_c^R - \gamma_{c-1}^R = \phi_0 + \phi_1\left(\gamma_{c-1}^R - \gamma_{c-2}^R\right) + \omega_c \qquad\qquad \omega_c \sim N\left(0, \sigma_\omega^2\right)$$

The slope of the long-term linear trend is given by $\phi_0/(1-\phi_1)$ if $|\phi_1| < 1$, the extent of autocorrelations is determined by the size of ϕ_1, and ω_c is the normal error term with mean zero and variance σ_ω^2. Under the ARIMA(1,1,0), if $|\phi_1| < 1$, the future variability of γ_c^R is finite and increases over time. It is assumed that $\left(\varepsilon_{t,1}, \varepsilon_{t,2}, \varepsilon_{t,3}\right)$ and ω_c are independent.

For the M7-M5 model regarding the difference between the book and reference populations, the time series of $\kappa_{t,1}^B$ and $\kappa_{t,2}^B$ are modelled as a vector autoregressive process of order one, VAR(1):

$$\begin{pmatrix} \kappa_{t,1}^B \\ \kappa_{t,2}^B \end{pmatrix} = \begin{pmatrix} \varphi_{1,0} \\ \varphi_{2,0} \end{pmatrix} + \begin{bmatrix} \varphi_{1,1} & \varphi_{1,2} \\ \varphi_{2,1} & \varphi_{2,2} \end{bmatrix} \begin{pmatrix} \kappa_{t-1,1}^B \\ \kappa_{t-1,2}^B \end{pmatrix} + \begin{pmatrix} \xi_{t,1} \\ \xi_{t,2} \end{pmatrix} \qquad\qquad \begin{pmatrix} \xi_{t,1} \\ \xi_{t,2} \end{pmatrix} \sim N(0, \Psi)$$

The extent of autocorrelations depends on the values of $\varphi_{1,1}$, $\varphi_{1,2}$, $\varphi_{2,1}$, and $\varphi_{2,2}$, and $\xi_{t,1}$ and $\xi_{t,2}$ are the bivariate normal error terms with mean zero and covariance matrix Ψ. If there is a further requirement that the projected (best estimate) book-to-reference ratio of mortality rates at each age should not diverge over the long term, all the eigenvalues of the matrix $\begin{bmatrix} \varphi_{1,1} & \varphi_{1,2} \\ \varphi_{2,1} & \varphi_{2,2} \end{bmatrix}$ have to be smaller than one in absolute value and so both $\kappa_{t,1}^B$ and $\kappa_{t,2}^B$ converge to a constant over time. Under such conditions, the future variability of $\kappa_{t,1}^B$ and $\kappa_{t,2}^B$ are bounded across time, in contrast to using the MRWD. It is also assumed that the time series error terms of the reference population and those of the difference between the book and reference populations are independent.

Figure 6.1 plots the various time series computed (solid lines), their best estimate projections (dashed lines), and their simulated 95% prediction intervals (dotted lines) for a particular UK pension group. As shown, $\kappa_{t,1}^R$, $\kappa_{t,2}^R$, $\kappa_{t,3}^R$, and γ_c^R have linear trends with unbounded prediction intervals, while $\kappa_{t,1}^B$ and $\kappa_{t,2}^B$ converge to a constant with bounded variability.

Note that under the characterisation approach (Subsection 3.3), the time series of $\kappa_{t,1}^B$ and $\kappa_{t,2}^B$ of different characterising groups are also co-modelled as a VAR(1).

Figure 6.1   Time series projections and 95% prediction intervals for a UK pension group (basic materials; normal retirement; lower pension group) under M7-M5 model
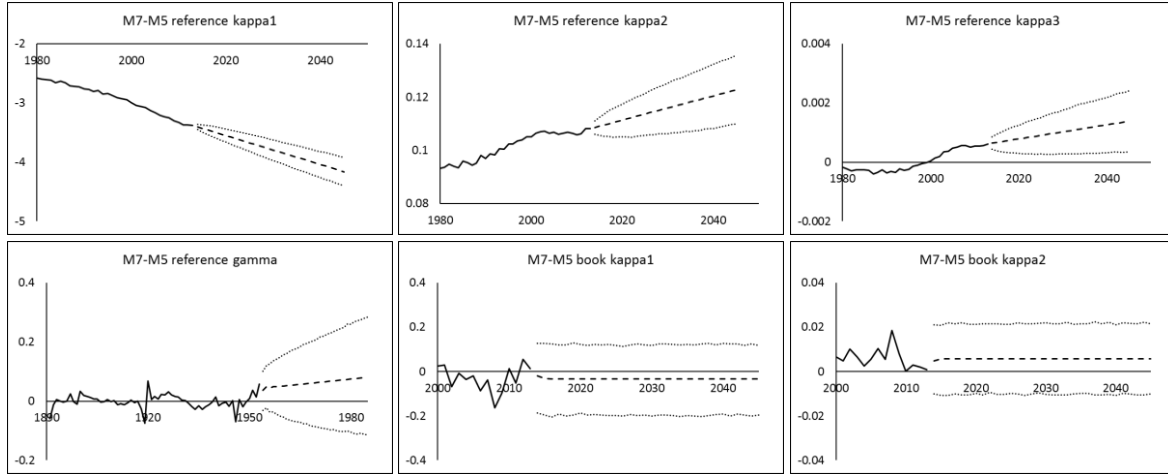
6.2   CAE+Cohorts model

For the CAE+Cohorts model (Subsection 3.2) regarding the reference population, the time series of κ_t^R is modelled as a random walk with drift (RWD):

$$\kappa_t^R = d + \kappa_{t-1}^R + \varepsilon_t \qquad\qquad \varepsilon_t \sim N\!\left(0, \sigma_\varepsilon^2\right)$$

The parameter d is the drift term (linear trend's slope) and ε_t is the normal error term with mean zero and variance σ_ε^2. Under the RWD, the future variability of κ_t^R increases over time. Furthermore, the time series of γ_c^R is modelled as an ARIMA(1,1,0), like in the M7-M5 model. It is assumed that ε_t and ω_c are independent.

For the CAE+Cohorts model regarding the difference between the book and reference populations, the time series of κ_t^B is modelled as an autoregressive process of order one, AR(1):

$$\kappa_t^B = \varphi_0 + \varphi_1 \, \kappa_{t-1}^B + \xi_t \qquad\qquad \xi_t \sim N\!\left(0, \sigma_\xi^2\right)$$

The long-term mean of κ_t^B is given by $\varphi_0/(1-\varphi_1)$ if $\left|\,\varphi_1\,\right|<1$, the extent of autocorrelations depends on the size of φ_1, and ξ_t is the normal error term with mean zero and variance σ_ξ^2. If it is required that the projected (best estimate) book-to-reference ratio of mortality rates at each age should not diverge in the long run, $\left|\,\varphi_1\,\right|$ has to be smaller than one and so κ_t^B converges to $\varphi_0/(1-\varphi_1)$ over time. Under this condition, the future variability of κ_t^B is bounded across time. It is further assumed that the time series error terms of the reference population and that of the difference between the book and reference populations are independent.

Figure 6.2 shows the different time series computed (solid lines), the best estimate projections (dashed lines), and their simulated 95% prediction intervals (dotted lines) for the same UK pension group as above. It can be seen that κ_t^R and γ_c^R have linear trends over time with unbounded prediction intervals, whereas κ_t^B tends to a constant value with bounded intervals.

Figure 6.2   Time series projections and 95% prediction intervals for a UK pension group (basic materials; normal retirement; lower pension group) under CAE+Cohorts model
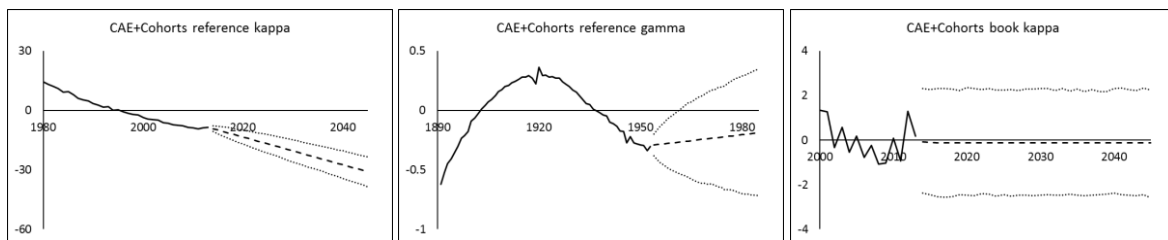
If the characterisation approach (Subsection 3.3) is adopted, the time series of κ_t^B of different characterising groups are modelled jointly as a VAR(1) process.

6.3   Mortality convergence

Figure 6.3 plots the book-to-reference ratios of mortality rates at each age range over time for the three datasets. There are a number of general patterns that can be observed. First, considering the UK and Australian pensioners, the ratios of females tend to be more volatile across time than those of males, probably due to the smaller data sizes of the former. For the IMD quintile groups, the data sizes are more comparable and there is no such difference between the two sexes. Second, the differences between the industries or groups or states become smaller at older ages. In particular, the ratios are quite close to one at ages 80 to 89. Moreover, for the UK and Australian pensioners, the ratios are usually below one and fluctuate fairly randomly over time, with no distinct level differences between the industries or states. In contrast, for the IMD quintile groups, the ratios clearly decrease from above one to below one for the most deprived areas to the least deprived areas, and since these groups represent significant segments of the reference population, their ratios are much more stable across time. Overall, despite the fluctuations, there do not seem to be any particular temporal trends in the ratios for the UK and Australian pensioners. Comparatively, there are some slowly increasing trends in the ratios for the most deprived areas, i.e. the mortality improvement of the most deprived group is not as fast as that of the general population. Since the data length is only a little more than 10 years, it is difficult to deduce whether these trends will continue in the long term. One may make an arbitrary adjustment (e.g. adding a term which changes linearly with time; see Subsection 6.4) to the time series modelling of the

difference between the book and reference populations to allow for this effect. Another possible alternative is the use of a new concept called semicoherence (Li et al. 2015), in which the projected mortality trajectories of two related populations are permitted to diverge over certain time periods, only within a specific tolerance corridor. But as these past trends are rather mild, and if the corresponding adjustment needed is small, the resulting impact on the calculated hedging results is unlikely to be material.

Note that both the M7-M5 and CAE+Cohorts models are 'coherent' (Li and Lee 2005). This feature means that the projected (best estimate) ratio of future mortality rates between the book and reference populations at each age tends towards a constant over the long term. This assumption appears to be roughly in line with the observations in Figure 6.3 and also follows the usual expectation that the two related populations' mortality trends would not diverge indefinitely. A lack of this long-term coherence in the modelling process may lead to an overestimation of longevity basis risk. In the short term, however, the two projected trends could diverge under the two models, the extent to which depends on how fast the fitted (weakly stationary) time series processes of the book component converge in the projections. For example, under the CAE+Cohorts model, the projected values of κ_t^B would approach the constant $\varphi_0/(1-\varphi_1)$ faster if φ_1 is smaller in magnitude and closer to zero.

Apart from the best estimates' coherence, it is also important to take full account of the simulated future variability. How the possible extent of co-movements (away from the projected trends) between the book and reference populations is being modelled is one major factor determining the calculated amount of longevity basis risk. Under the two models, the fitted time series processes for the reference population generate unbounded future variability. In contrast, those for the difference between the book and reference populations produce bounded future variability. The former variability would then dominate in the long run while the latter would diminish in its impact relatively. The consequence is that the extent of co-movements between the two populations (as a portion of the total variability or uncertainty which increases over time) would grow with time in the future simulations. This modelling effect is in agreement with the common view that it would be unlikely for the two related populations to continue to deviate significantly in their long-term future mortality improvements.

Figure 6.3   Book-to-reference ratios of past mortality rates

Basic materials, industrials, consumer goods, consumer services, and utilities (black to light grey) industries (normal retirement; females and males)
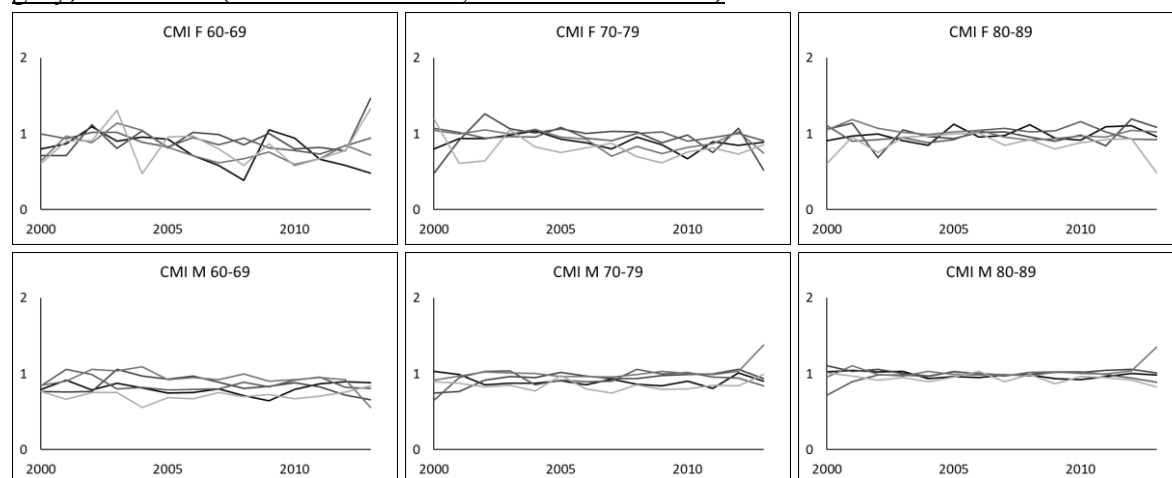
Most deprived to least deprived (black to light grey) areas (females and males)
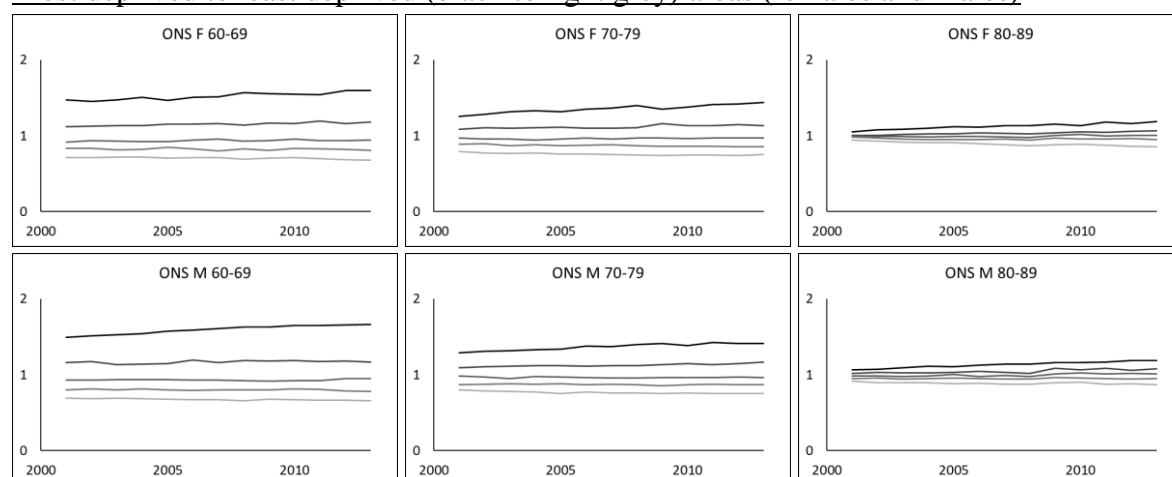
New South Wales, Commonwealth, Victoria, and other Australian states, and New Zealand (black to light grey) (retirement; females and males)
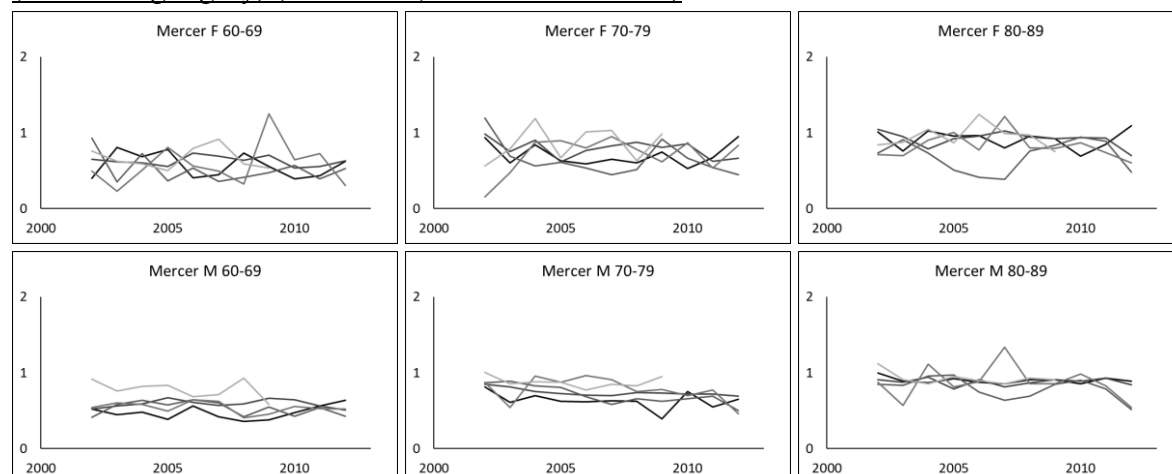
6.4   Sensitivity analysis

In fact, there exist a great variety of time series processes with different degrees of sophistication in the literature. Unlike econometric and financial data which usually have at least, say, hundreds of data points, however, the length of annual mortality data is generally much shorter and so limits the use of more advanced models. Despite this limitation, it would still be of practical interest to investigate how sensitive the calculated hedging results are to the use of different time series processes. To examine the potential significance of varying times series modelling assumptions, the following alternatives are considered: (a) choosing a higher order for the autoregressive processes in the book component; (b) using univariate (rather than multivariate) time series processes in the reference component under the M7-M5 model; (c) relaxing the independence assumption between the time series error terms of the reference and book components; and (d) applying some other combinations of time series processes given the data constraint. For the autoregressive processes, the selected order for further testing cannot be too high, as the data length of the book population is short. Moreover, the order should be chosen and the parameters estimated in such a way that the autoregressive processes are still weakly stationary. Accordingly, an order of two is tested in this analysis.

Table 6.2 (Figure 6.4) presents the level of longevity risk reduction for each combination of time series processes, and Table 6.3 provides the key characteristics of each case. When a higher order is used for the autoregressive processes in the book component (see (a)), the risk reduction levels tend to be lower. Although the VAR(2) and AR(2) processes are still weakly stationary, their pace of convergence would often be slower than that of the VAR(1) and AR(1) processes, because of the autoregressive parameters for the longer lag and so greater autoregressive effects. Hence the projected trends of the book and reference populations could diverge for a longer time before convergence, resulting in a smaller risk reduction. When independent RWD processes are used instead of the MRWD in the reference component under the M7-M5 model (see (b)), the risk reduction estimates are slightly smaller. The univariate RWD processes could produce a different level of total variability to that generated by the MRWD, which may explain the differences in risk reduction, though these differences seem to be immaterial. When the time series error terms of the book and reference components are treated as correlated rather than independent (see (c)), the changes in the risk reduction estimates are minimal. It seems that this independence or correlation assumption does not have a significant influence on the calculated hedging results.

If the MRWD in the reference component under the M7-M5 model is replaced by the VARIMA(1,1,0) process (see (d1)), the risk reduction estimates drop by about 10% in magnitude. If the RWD in the reference component under the CAE+Cohorts model is changed with the ARIMA(1,1,0) process, the risk reduction levels decline by around 5%. Though both the VARIMA(1,1,0) and ARIMA(1,1,0) processes, like the MRWD and RWD,

generate unbounded future variability, the actual levels of variability over time under these integrated autoregressive processes would be different to those under the random walk processes, which may lead to different risk reduction levels. Another important feature of the integrated autoregressive processes is that their projected paths converge to some linear trends and the pace of convergence depends on the autoregressive parameters estimated.

When the autoregressive processes in the book component are replaced by the random walk processes *without* drift (see (d2)), although the models are still coherent because of the flat trends (i.e. zero drift), the risk reduction estimates decrease significantly to very low values. The random walk processes produce unbounded future variability, which means that the two populations' future mortality movements could deviate significantly in the simulations, especially over the long term. One may argue that this implication is not biologically and socially reasonable, as the two populations are associated demographically, and the corresponding assumption should thus be avoided. Once again, it can be seen that the behaviour of simulated future variability of the difference between the book and reference populations is a decisive factor in determining the calculated level of longevity risk reduction. Extra caution must be exercised in choosing the time series process for the book component.

As noted in Subsections 2.5 and 6.3, the two populations may experience different mortality improvements in the short term. They then may or may not move more consistently in the long run, though the usual view in the literature is that it is unlikely for the two related populations to continue to diverge in their mortality movements for an extensive period. Despite this common view, even if an arbitrarily determined extra term, which changes linearly with time, is added within the autoregressive processes in the book component (see (d3)) to allow for an expected long-term discrepancy, the risk reduction estimates are actually more or less the same, though with some rather small differences for the IMG quintile group. This observation indicates that non-coherence alone (without unbounded future variability) would not have a significant impact on the calculated hedging results. The implication is that as long as the major mortality trends, diverging or not, are well incorporated in the model and calibration and do not deviate much from the model expectation in the future, there would still be a decent level of hedge effectiveness. It is mainly the potential random deviations (future variability) from the expect trends that would affect the hedging performance.

In brief, given the short data length of the book population, the practical, feasible choices of time series processes are rather limited. The analyses above and also in Subsection 5.9 suggest that the importance of the various time series modelling assumptions follows the order of the behaviour of simulated future variability of the book component, the pace of reaching coherence between the two populations, and then the other correlation assumptions. Appropriate judgement, reference materials, experts' opinions, and thorough testing are required in setting these modelling assumptions properly in practice. Further research is also needed when more data of longer periods and for different kinds of book populations are collected in the future.

Table 6.2 Levels of longevity risk reduction (in % of initial longevity risk) using different combinations of time series processes

100,000 Males Reference Component / Book Component	UK Pensioners (basic materials; normal retirement; lower pension)			IMD Groups (most deprived areas)			Australian Pensioners (New South Wales; retirement)		
	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES	Standard Deviation	99.5% VaR	99.5% ES
(Original) M7M5 MRWD / VAR(1)	65	63	63	64	72	70	72	69	63
CAE+Cohorts RWD / AR(1)	77	74	73	69	68	68	74	70	68
(a) M7M5 MRWD / VAR(2)	45	55	47	35	58	35	58	54	39
CAE+Cohorts RWD / AR(2)	75	73	72	69	72	71	36	65	57
(b) M7M5 RWD / VAR(1)	62	60	60	60	67	69	68	65	61
(c) M7M5 Correlated MRWD & VAR(1)	66	63	61	60	66	64	71	67	65
CAE+Cohorts Correlated RWD & AR(1)	79	76	75	66	65	66	75	72	68
(d1) M7M5 VARIMA(1,1,0) / VAR(1)	56	53	52	52	62	64	61	57	53
CAE+Cohorts ARIMA(1,1,0) / AR(1)	72	69	68	62	63	63	67	62	60
(d2) M7M5 MRWD / BRW	56	36	37	44	41	37	3	5	2
CAE+Cohorts RWD / RW	13	6	7	20	19	15	19	11	8
(d3) M7M5 MRWD / Modified VAR(1)	64	60	59	78	77	77	73	72	71
CAE+Cohorts RWD / Modified AR(1)	75	73	72	88	87	87	75	70	69

Note: The terms VAR(2), AR(2), VARIMA(1,1,0), ARIMA(1,1,0), BRW, and RW stand for the vector autoregressive process of order two, autoregressive process of order two, vector autoregressive integrated moving average process of order (1,1,0), autoregressive integrated moving average process of order (1,1,0), bivariate random walk *without* drift, and random walk *without* drift respectively. More details can be found in Appendix I. For (c), the error terms of the two time series processes are treated as correlated. For (d3), the modified VAR(1) and AR(1) have a (vector) term of δt, which

changes linearly with time t, in addition to the constant term(s) and autoregressive parameter(s), like how mortality jumps are added in a separate equation as in Subsection 5.10.

Table 6.3   Major characteristics of different combinations of time series processes

Model	Reference / Book Component	Time Series Process	Best estimate(s)	Variability
(Original)				
M7-M5	Reference	MRWD	linear trends	unbounded
M7-M5	Book	VAR(1)	convergence	bounded
CAE+Cohorts	Reference	RWD	linear trend	unbounded
CAE+Cohorts	Book	AR(1)	convergence	bounded
(a)				
M7-M5	Reference	MRWD	linear trends	unbounded
M7-M5	Book	VAR(2)	convergence	bounded
CAE+Cohorts	Reference	RWD	linear trend	unbounded
CAE+Cohorts	Book	AR(2)	convergence	bounded
(b)				
M7-M5	Reference	RWD	linear trends	unbounded
M7-M5	Book	VAR(1)	convergence	bounded
(c)				
M7-M5	Reference	Correlated MRWD	linear trends	unbounded
M7-M5	Book	& VAR(1)	convergence	bounded
CAE+Cohorts	Reference	Correlated RWD	linear trend	unbounded
CAE+Cohorts	Book	& AR(1)	convergence	bounded
(d1)				
M7-M5	Reference	VARIMA(1,1,0)	long-term linear trends	unbounded
M7-M5	Book	VAR(1)	convergence	bounded
CAE+Cohorts	Reference	ARIMA(1,1,0)	long-term linear trend	unbounded
CAE+Cohorts	Book	AR(1)	convergence	bounded
(d2)				
M7-M5	Reference	MRWD	linear trends	unbounded
M7-M5	Book	BRW	flat trends	unbounded
CAE+Cohorts	Reference	RWD	linear trend	unbounded
CAE+Cohorts	Book	RW	flat trend	unbounded
(d3)				
M7-M5	Reference	MRWD	linear trends	unbounded
M7-M5	Book	Modified VAR(1)	non-convergence	bounded
CAE+Cohorts	Reference	RWD	linear trend	unbounded
CAE+Cohorts	Book	Modified AR(1)	non-convergence	bounded

Figure 6.4 **Levels of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR) using different combinations of time series processes (100,000 males)**
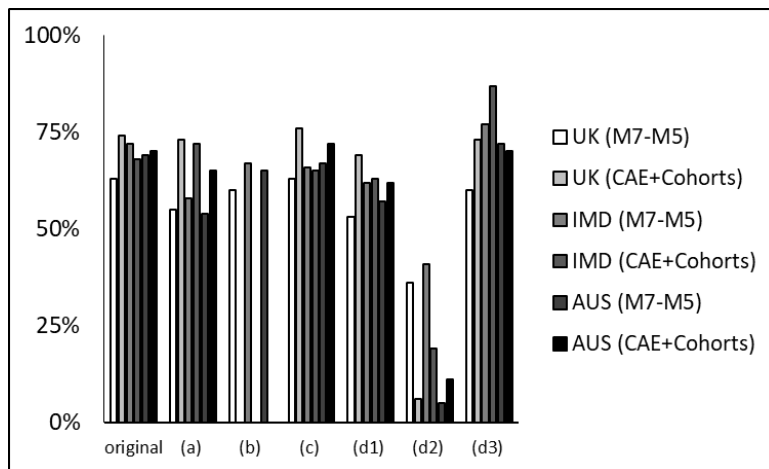
7 Practical Guidelines

An extensive investigation on modelling longevity basis risk and assessing longevity risk reduction has been conducted in the previous sections. Section 2 depicts the historical mortality levels and improvements of different subgroups in the three datasets. Section 3 summarises the modelling procedure of longevity basis risk and the means to calculate the level of longevity risk reduction. Section 4 considers hedging pension portfolios of a single cohort and multiple cohorts with standardised longevity swaps and estimates the resulting risk reduction levels. Section 5 carries out a detailed sensitivity analysis on the hedging results from Subsection 4.1 through making a range of changes to the initial model settings and assumptions, and performs backtesting and scenario testing on the hedging strategy. Section 6 conducts a sensitivity analysis on the choice of time series modelling, given the data constraint of the book population. In this final section, we set forth a brief summary of the hedging results so far, both qualitatively and quantitatively. Particularly, we present the results in a way that can possibly be used as a quick guide for practitioners and regulators, and also discuss a number of practical issues in adopting index-based hedging solutions.

7.1 Overall summary

Based on the hedging results in Sections 4 to 6, the level of longevity risk reduction in a particular hedging scenario may be determined very approximately in a qualitative way. Suppose that the book population under consideration has mortality experience similar to those of the three datasets, the mortality patterns are reasonably captured by the modelling procedure in this report, and the time series assumptions properly reflect the extent of future mortality co-movements between the book and reference populations. Further assume that the hedging strategy uses standardised longevity swaps of multiple cohorts which are calibrated closely to the demographic structure of the pension plan, and that interest rate risk is minimal due to a consistently low interest rate environment or an interest rate hedging scheme being in place. If the size of the pension plan is large, the risk reduction level would be in the range of about 50% to 80%. The more precise value would depend on a number of factors, including how related the book and reference populations are, the pace of reaching coherence between the two populations in the future, and whether additional features such as mortality structural changes are taken into account. Some extent of subjective judgement may be exercised in examining these factors, somewhat like how the level of correlation between two lines of general insurance business was determined qualitatively (as high, medium, or low) by senior actuaries as stated in Collings and White (2001).

For demonstration purposes, a qualitative assessment may be structured as in Table 7.1 below, in which a score of 0 to 10 (from mild to significant) is given to each question. For instance, consider the following hypothetical case study of a pension plan of 30,000 male lives with multiple cohorts. It is estimated that the plan's longevity risk (99.5% VaR minus the mean of

the present value of all future cash outflows) is around £250 million. After examining the underlying factors and conditions, the pension actuary has given a rank of 6 to the relationship between the book and reference populations, a rank of 5 to how fast the two populations' mortality trends would move back in line, and a rank of 4 to the possible occurrence of mortality structural changes and their potential impact on the two populations. As a result, the overall risk reduction level is estimated as 65% (= 50 + 6 + 5 + 4). Then the plan's remaining longevity risk after hedging, due to the existence of longevity basis risk, is about £87.5 million (= 250 × 35%).

Table 7.1   Qualitative assessment of longevity risk reduction (in % of initial longevity risk, in terms of 99.5% VaR)

Assessment Question	Longevity Risk Reduction
The pension plan has around 25,000 lives or more?	Yes (50) or No (20)
The book and reference populations are related demographically, e.g. the former is a general subset of the latter?	+ a rank of 0 to 10
The pace of reaching coherence between the book and reference populations is fast in the future?	+ a rank of 0 to 10
Mortality structural changes can affect the two populations' future mortality significantly in the same direction?	+ a rank of 0 to 10

Besides a qualitative assessment as above, the level of longevity risk reduction may also be measured approximately in a simple, quantitative way. The idea is similar to the 'rule-of-thumb' formula for calculating the diversification benefit for general insurance liabilities in Bateup and Reed (2001). Using the hedging results for males in Sections 4 to 6, we apply a simple linear regression to the risk reduction estimates (regarding the 99.5.% VaR) with the ten explanatory variables in Table 7.2. The first five variables refer to the pension plan and the hedging environment, and the others describe the model settings and assumptions. This approach provides a brief summary of the previous numerical results and takes all major factors into account simultaneously in a single equation. It is easy to understand and may serve as a quick guide for assessing the amount of risk reduction in an index-based longevity hedge. It can also be used as a rough check for those who want to perform detailed calculations with the technical information provided in Appendices I and II. There are, however, a number of limitations behind this approach. First, the computed coefficients of the regression depend mainly on the data and the models adopted, and they may not be applicable when other data and models are used. Second, although there are already 329 data points (cases) with ten explanatory variables, the cases simulated and the variables selected may not fully reflect all the underlying factors. In particular, the spread of the simulated cases between different variables is not even, which may result in certain biases in the results. Moreover, the use of linear relationships, without considering much about data transformations and potential interactions between variables, is admittedly an over-simplification of the real situation.

Users must be aware of all these limitations when taking this simple approach. The regression equation, based solely on the numerical results in Sections 4 to 6, is computed as (with explanatory variables x_1 to x_{10}):

$$\text{longevity risk reduction} = -0.0303 + 0.0644 x_1 + 0.0553 x_2 - 0.0784 x_3 - 0.8120 x_4 - 0.1662 x_5$$
$$+ 0.0006 x_6 + 0.0219 x_7 + 0.1204 x_8 + 0.1217 x_9 - 0.0762 x_{10}$$

subject to a minimum level of 0% and a maximum level of 100%.

Table 7.2   Explanatory variables for measuring longevity risk reduction

Explanatory Variable	Interpretation
(1) log portfolio size	The larger the portfolio size, the smaller the sampling basis risk level, the larger the risk reduction estimate
(2) 1: pensioners / 2: IMD groups 1 & 5 / 3: IMD groups 2 & 4 / 4: IMD group 3	As the data size of the book population and its demographic relationship to the reference population increases, the risk reduction estimate increases
(3) 1: tailored hedging / 2: less tailored hedging / 3: approximate hedging	The less precise the hedging scheme, the smaller the risk reduction estimate
(4) interest rate	Assuming the interest rate is constant, the higher the interest rate, the smaller the risk reduction estimate
(5) 1: swaps / 2: q-forwards	Optimising both the present value and individual cash flows, using q-forwards leads to a smaller risk reduction estimate
(6) 0: another model / 1: M7-M5	The risk reduction estimate depends on the model adopted
(7) 0: another model / 1: CAE+Cohorts	The risk reduction estimate depends on the model adopted
(8) 1: residuals bootstrapping / 2: parametric simulation	Using parametric simulation, instead of residuals bootstrapping, ignores parameter error and understates longevity basis risk, resulting in a larger risk reduction estimate
(9) 1: original / 2: structural changes	Further incorporating mortality structural changes that affect the two populations' future mortality significantly in the same direction leads to a larger risk reduction estimate
(10) 1: low AR order / 2: higher AR order	Assuming the book component has bounded future variability, the higher the autoregressive order, the slower the convergence of the time series process, the smaller the risk reduction estimate

Now consider another hypothetical case study below of a pension plan of 30,000 male lives with multiple cohorts. The plan's longevity risk (99.5% VaR minus the mean of the present value of all future cash outflows) is estimated to be about £250 million. The book population is a small, unique subset of the reference population. Standardised longevity swaps are calibrated closely to the demographic structure of the pension plan. The interest rate is assumed to be 2% p.a. flat. The CAE+Cohorts modelling assumption and residuals bootstrapping are selected, in which the coherence property holds and the book component's time series process has a high autoregressive order, as it is expected that the two populations'

mortality trends would deviate for a short while but move more closely with each other in the long term. It is further assumed that possible structural changes because of medical advances could impact on the two populations' future mortality significantly in the same direction. Consequently, the overall risk reduction level is estimated to be 66%, and the plan's residual longevity risk after hedging, due to the presence of longevity basis risk, is around £85 million ($= 250 \times 34\%$):

$$\text{longevity risk reduction} = -0.0303 + 0.0644(\ln(30000)) + 0.0553(1) - 0.0784(1) - 0.8120(0.02)$$
$$-0.1662(1) + 0.0006(0) + 0.0219(1) + 0.1204(1) + 0.1217(2) - 0.0762(2) = 0.66$$

The qualitative and quantitative assessment methods above are straightforward to understand and apply in practice. But it is very important to note that these quick guides are fundamentally a short summary of the numerical results in this report, which is model-dependent and is specific to the datasets being modelled. Given suitable resources and required expertise, practitioners are encouraged to follow the precise technical details to build up their own models (M7-M5, CAE+Cohorts, or others) and perform more accurate calculations. It may also be practically feasible that a thorough longevity hedging valuation is conducted only once per every few years, during which a qualitative assessment table or a simple linear regression 'rule-of-thumb' formula can be developed from the previous valuation for temporary use. Similarly, regulators may collect data from a wide range of pension plans and annuity portfolios and form these assessment guidelines as a general reference for index-based hedging exercises in the industry.

7.2   Practical issues

Ever since the very first pension buy-in was transacted in 2007, there has been much development in the UK life market regarding de-risking solutions. Figure 7.1 shows that the market has grown in size significantly over the period. From 2011 to 2015, the average annual volume of pension buy-ins, buy-outs, and longevity swaps has reached more than £18 billion (LCP 2016). This size, however, is only about 1.5% of the total private sector defined benefit pension assets, which is estimated to be around £1.2 trillion. Despite the recent strong growth, it is clear that there is still plenty of room for the market to develop in the coming years. Besides the UK, a few other countries like Ireland, the Netherlands, Switzerland, Canada, and the US have also seen some growth in their de-risking markets. Although regulations and market practices vary from country to country, there exist many opportunities for companies in different countries to de-risk their pension plans.

As pension plans mature over time and manage to reduce investment, interest rate, and inflation risks, longevity risk becomes the key risk for pension plan sponsors. From 2006 to 2015, the asset proportion in equities has reduced by about half, while the proportion in gilts, bonds, and other matching assets has nearly doubled (PIC 2015). A shift from these fixed interest investments to buy-ins, buy-outs, and longevity swaps can reasonably be regarded as

a sensible next stage for many pension plans. Moreover, apart from reducing longevity risk, de-risking can also lead to stronger corporate governance and reduced operational costs. These additional benefits would particularly be appreciated by those sponsors of smaller pension plans who want to avoid spending the time and resources required in managing their plans.

Figure 7.1   Buy-ins, buy-outs, and longevity swaps volumes from 2007 to 2015
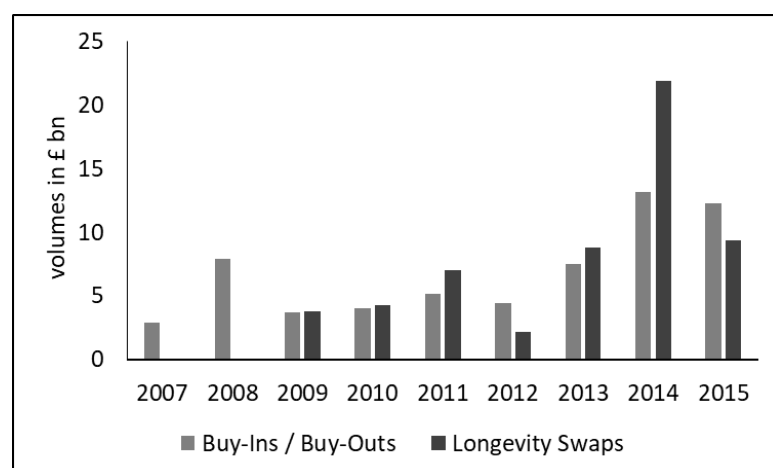
Currently, under a falling interest rate environment, most pension plans have substantial deficits on their books. Only around 10% of the FTSE 100 companies' pension plans are over 80% funded compared to their buy-out costs (LCP 2016). For those smaller pension plans that belong to large institutions with significant cash balances, a full buy-out would be a feasible option, leading to savings in operational costs and higher business flexibility. On the other hand, for larger pension plans, a buy-out would usually be too expensive; instead, a phased approach involving a series of well-timed buy-ins could be more flexible and cost-effective. Furthermore, for large pension plans which have the resources to maintain an internal investment strategy and manage their own investment, interest rate, and inflation risks, longevity swaps represent a decent alternative, in which only longevity risk is hedged. As noted above, there is a huge potential for the de-risking market to continue to grow. Looking forward, when the general funding levels improve, and also when de-risking tools become more widely affordable and more innovative solutions appear in the market, the buy-in, buy-out, and longevity swap market will have a great chance to flourish.

On the supply side, insurers and reinsurers had been quite conservative in accepting longevity risk. Recently, a number of events appear to have changed their appetite for this risk. First, there has been a sudden drop in the sale of individual annuities since the 2014 Budget was released. Consequently, insurers attempt to alleviate the problem by shifting the attention and capacity to pension plans. Second, reinsurers show stronger interest in recent years, in which those in the de-risking market generally have more mortality risk than longevity risk and so assuming longevity risk from pension plans can offer them diversification benefits. Moreover,

as Solvency II has finally settled down in 2016, after considerable effort in lobbying and creating innovative approaches in pricing and capital modelling, many insurers can now renew their focus and risk appetite and find access to new capital. Note that under Solvency II, insurers are required to match their asset and liability cash flows closely, and they are encouraged to implement more longevity risk hedging than previously. As a result, reinsurers' involvement in the transfer of longevity risk has enhanced prominently.

In spite of the increased insurer and reinsurer capacity, the demand for longevity risk transfer from pension plans is expected to continue to grow, and may exceed the supply and drive up the prices at some point. More capital is then needed in order to maintain a sufficient level of supply, and index-based solutions and standardised products could be the key to open up the gate to the wider capital market. Compared to most longevity transactions to date which have mainly been bespoke in nature, index-based longevity- or mortality-linked securities and derivatives are more transparent and standardised, and could draw more interest from both inside and outside the insurance industry. Market investors may take on longevity risk from insurers and pension plans in exchange for appropriate risk-adjusted returns. Some may use external longevity exposures to offset their own mortality exposures. Others may also diversify their investment portfolios over a new longevity asset class, which is arguably uncorrelated with traditional asset classes. Then with access to extra capital, index-based products can be priced competitively amongst insurers, resulting in lower costs for pension plans in using these products. A notable recent example (though there have been only a few indexed-based transactions so far) is the €12 billion longevity swap provided by Deutsche Bank to Dutch insurer Aegon, announced in 2012, in which the reference population was set as the Dutch population and the trade was targeted at capital market investors specifically.

Nevertheless, there are a number of concerns in adopting index-based hedging solutions. Firstly, longevity risk has a very long-term nature, and market investors would find it highly uncertain, especially when they are unfamiliar with the assessment of this risk. For instance, the Dutch population-based longevity swap transaction in 2012 was limited to 20 years, reflecting the market's perception on the potential open-ended losses. It would take some time for the market to build up a better understanding of longevity risk. Secondly, from the hedger's perspective, the presence of longevity basis risk means that an index-based hedge will be imperfect and there will be residual longevity risk. The main purpose of this report is to address this concern – according to the previous numerical results, using a range of model settings and industry datasets, the risk reduction level is computed to be around 50% to 80% for a large portfolio. This risk reduction effect has the potential to enable a reduction in capital for insurers, either through greater confidence in their own pricing or possibly favourable treatment from the regulator. But it is of critical importance to further test the index-based solutions extensively on more data and scenarios to identify how exactly they can help in capital assessment and pricing exercise, and also communicate the results properly with different stakeholders. Insurers may then be interested in working with

consultants to discuss with the regulator on what relief may be given, how to assess the impact on capital requirements, and what further steps are needed to gain the relief. If improvements in capital requirements turn out to be possible, insurers and consultants can work together to explain to the clients about the hedging mechanism and obtain the best pricing for the clients. Investment banks may also use the findings to assist in pricing hedging products or structured investment products and help them market the products to potential market investors.

Finally, it would be practically useful to have some standardisation on the key factors that determine how much longevity basis risk is present in an index-based hedge, and also some sensitivity analysis and scenario testing conducted to examine the impact of potential variations in the future. Sections 5 and 6 (sensitivity testing) and Subsection 7.1 (simple qualitative and quantitative assessments) have already made an initial attempt to propose some solutions for these problems. Further research is called for to test other models, data, and scenarios and to identify a more thorough, standardised list of major factors driving longevity basis risk.

References

Bateup, R. & Reed, I. (2001). *Research and Data Analysis Relevant to the Development of Standards and Guidelines on Liability Valuation for General Insurance*. The Institute of Actuaries of Australia, Tillinghast – Towers Perrin.

Blake, D. & Burrows, W. (2001). Survivor bonds: Helping to hedge mortality risk. *Journal of Risk and Insurance*, 68(2), 339-348.

Bell, W.R. (1997). Comparing and assessing time series models for forecasting age-specific fertility and mortality rates. *Journal of Official Statistics*, 13(3), 279-303.

Booth, H., Maindonald, J. & Smith, L. (2002). Applying Lee-Carter under conditions of variable mortality decline. *Population Studies*, 56, 325-336.

Brouhns, N., Denuit, M. & Van Keilegom, I. (2005). Bootstrapping the Poisson log-bilinear model for mortality forecasting. *Scandinavian Actuarial Journal*, 3, 212-224.

Brouhns, N., Denuit, M. & Vermunt J.K. (2002). A Poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics and Economics*, 31, 373-393.

Cairns, A.J.G. (2004). *Interest Rate Models*: *An Introduction*. Princeton University Press.

Cairns, A.J.G., Blake, D. & Dowd, K. (2006). A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk and Insurance*, 73(4), 687-718.

Cairns, A.J.G., Blake, D. & Dowd, K. (2008). Modelling and management of mortality risk: A review. *Scandinavian Actuarial Journal*, 2008(2-3), 79-113.

Carter, L.R. & Lee, R.D. (1992). Modeling and forecasting US sex differentials in mortality. *International Journal of Forecasting*, 8, 393-411.

Chan, W.S., Li, J.S.H. & Li, J. (2014). The CBD mortality indexes: Modeling and applications. *North American Actuarial Journal*, 18(1), 38-58.

Collings, S. & White, G. (2001). *APRA Risk Margin Analysis*. The Institute of Actuaries of Australia XIII[th] General Insurance Seminar, Trowbridge Consulting.

Committee of European Insurance and Occupational Pensions Supervisors (CEIOPS). (2010). *QIS5 Technical Specifications*.

Coughlan, G.D., Epstein, D., Sinha, A. & Honig, P. (2007). *q-forwards: Derivatives for Transferring Longevity and Mortality Risk*. JPMorgan. www.lifemetrics.com

Coughlan, G.D., Khalaf-Allah, M., Ye, Y., Kumar, S., Cairns, A.J.G., Blake, D. & Dowd, K. (2011). Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness. *North American Actuarial Journal*, 15(2), 150-176.

Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of insurance and reinsurance (Solvency II). *Official Journal of the European Union*.

Dowd, K. (2003). Survivor bonds: A comment on Blake and Burrows. *Journal of Risk and Insurance*, 70(2), 339-348.

Dowd, K. & Blake, D. (2006). After VaR: The theory, estimation, and insurance applications of quantile-based risk measures. *Journal of Risk and Insurance*, 73(2), 193-229.

Dowd, K., Blake, D., Cairns, A.J.G. & Dawson, P. (2006). Survivor swaps. *Journal of Risk and Insurance*, 73(1), 1-17.

Haberman, S., Kaishev, V., Millossovich, P., Villegas, A., Baxter, S., Gaches, A., Gunnlaugsson, S. & Sison, M. (2014). *Longevity Basis Risk: A Methodology for Assessing Basis Risk*. Cass Business School and Hymans Robertson LLP. Institute and Faculty of Actuaries (IFoA) and Life and Longevity Markets Association (LLMA).

Hull, J.C. (2015). *Risk Management and Financial Institutions*. Fourth Edition. John Wiley & Sons, New Jersey.

Human Mortality Database (HMD). (2017). University of California, Berkeley (USA) and Max Planck Institute for Demographic Research (Germany). www.mortality.org

Kleinow, T. (2015). A common age effect model for the mortality of multiple populations. *Insurance: Mathematics and Economics*, 63, 147-152.

Koissi, M.C., Shapiro, A.F. & Högnäs, G. (2006). Evaluating and extending the Lee-Carter model for mortality forecasting: Bootstrap confidence interval. *Insurance: Mathematics and Economics*, 38, 1-20.

Lane Clark and Peacock LLP (LCP). (2012). *LCP Pension Buy-ins, Buy-outs and Longevity Swaps 2012*.

Lane Clark and Peacock LLP (LCP). (2016). *LCP Pensions De-Risking 2016*.

Lee, R.D. & Carter, L.R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, 87(419), 659-671.

Li, J. (2010). Projections of New Zealand mortality using the Lee-Carter model and its augmented common factor extension. *New Zealand Population Review*, 36, 27-53.

Li, J. (2014). A quantitative comparison of simulation strategies for mortality projection. *Annals of Actuarial Science*, 8(2), 281-297.

Li, J. & Haberman, S. (2015). On the effectiveness of natural hedging for insurance companies and pension plans. *Insurance: Mathematics and Economics*, 61, 286-297.

Li, J.S.H., Chan, W.S. & Zhou, R. (2015). Semicoherent multipopulation mortality modeling: The impact on longevity risk securitization. *Journal of Risk and Insurance*, DOI: 10.1111/jori.12135.

Li, J.S.H. & Hardy, M.R. (2011). Measuring basis risk in longevity hedges. *North American Actuarial Journal*, 15(2), 177-200.

Li, J.S.H. & Luo, A. (2012). Key q-duration: A framework for hedging longevity risk. *ASTIN Bulletin*, 42(2), 413-452.

Li, N. & Lee, R. (2005). Coherent mortality forecasts for a group of populations: An extension of the Lee-Carter method. *Demography*, 42(3), 575-594.

Life and Longevity Markets Association (LLMA). (2010). *Technical Note: The S-forward*. www.llma.org

Pension Insurance Corporation (PIC). (2015). *Buy-ins and Buyouts*.

Renshaw, A. & Haberman, S. (2003). Lee-Carter mortality forecasting: A parallel generalized linear modelling approach for England and Wales mortality projections. *Applied Statistics*, 52(1), 119-137.

Tan, C.I., Li, J., Li, J.S.H. & Balasooriya, U. (2014). Parametric mortality indexes: From index construction to hedging strategies. *Insurance: Mathematics and Economics*, 59, 285-299.

Tan, K.S., Blake, D. & MacMinn, R. (2015). Longevity risk and capital markets: The 2013-14 update. *Insurance: Mathematics and Economics*, 63, 1-11.

Tsai, J.T., Tzeng, L.Y. & Wang, J. L. (2011). Hedging longevity risk when interest rates are uncertain. *North American Actuarial Journal*, 15(2), 201-211.

Tsay, R.S. (2002). *Analysis of Financial Time Series*. John Wiley & Sons.

Zhou, R., Li, J.S.H. & Tan, K.S. (2013). Pricing standardized mortality securitizations: A two-population model with transitory jump effects. *Journal of Risk and Insurance*, 80(3), 733-774.

Acknowledgements

Appendix I Computation Algorithms

M7-M5 Model

The M7-M5 model has two major components:

$$\text{logit } q^R_{x,t} = \kappa^R_{t,1} + (x - \bar{x})\kappa^R_{t,2} + \left((x - \bar{x})^2 - \sigma^2_x\right)\kappa^R_{t,3} + \gamma^R_{t-x} = \eta^R_{x,t}, \tag{A1}$$

$$\text{logit } q^B_{x,t} - \text{logit } q^R_{x,t} = \kappa^B_{t,1} + (x - \bar{x})\kappa^B_{t,2} = \eta^B_{x,t}, \tag{A2}$$

in which for the first component, $q^R_{x,t}$ is the mortality rate of the reference population at age x in year t, $\kappa^R_{t,1}$, $\kappa^R_{t,2}$, and $\kappa^R_{t,3}$ refer to the level, slope, and curvature respectively of the mortality curve across age in year t, and γ^R_{t-x} represents the cohort effect of those lives born in year $t-x$. For the second component, the difference in the logit mortality rate between the book and reference populations, $\text{logit } q^B_{x,t} - \text{logit } q^R_{x,t}$, is modelled as a linear combination of $\kappa^B_{t,1}$ and $\kappa^B_{t,2}$, which are another two parameters for capturing the differences between the two populations in year t. The two model components are further denoted as $\eta^R_{x,t}$ and $\eta^B_{x,t}$. Note that $\bar{x} = \dfrac{1}{\text{no. of ages}}\sum_x x$ and $\sigma^2 = \dfrac{1}{\text{no. of ages}}\sum_x (x - \bar{x})^2$, and that there are three identifiability constraints $\sum_c \gamma^R_c = 0$, $\sum_c c\gamma^R_c = 0$, and $\sum_c c^2 \gamma^R_c = 0$. These constraints are set to ensure that the cohort parameters have valid solutions. The mortality rates can then be deduced as:

$$q^R_{x,t} = \frac{\exp\left(\eta^R_{x,t}\right)}{1 + \exp\left(\eta^R_{x,t}\right)} = \frac{1}{\exp\left(-\eta^R_{x,t}\right) + 1}, \tag{A3}$$

$$q^B_{x,t} = \frac{\exp\left(\eta^B_{x,t} + \eta^R_{x,t}\right)}{1 + \exp\left(\eta^B_{x,t} + \eta^R_{x,t}\right)} = \frac{1}{\exp\left(-\eta^B_{x,t} - \eta^R_{x,t}\right) + 1}. \tag{A4}$$

Assume that the force of mortality is constant within each age-time cell and so is equal to the central death rate. The random numbers of deaths at age x in year t of the reference and book populations are modelled as:

$$D^R_{x,t} \sim \text{Poisson}\left(e^R_{x,t}\, m^R_{x,t}\right), \tag{A5}$$

$$D^B_{x,t} \sim \text{Poisson}\left(e^B_{x,t}\, m^B_{x,t}\right), \tag{A6}$$

where $e^R_{x,t}$ and $e^B_{x,t}$ are the corresponding central exposed to risk measures, and $m^R_{x,t}$ and $m^B_{x,t}$ are the central death rates, which can be expressed as:

$$m^R_{x,t} = -\ln\left(1 - q^R_{x,t}\right) = \ln\left(1 + \exp\left(\eta^R_{x,t}\right)\right), \tag{A7}$$

$$m^B_{x,t} = -\ln\left(1 - q^B_{x,t}\right) = \ln\left(1 + \exp\left(\eta^B_{x,t} + \eta^R_{x,t}\right)\right). \tag{A8}$$

Their log likelihood functions are:

$$l^R = \sum_{x,t} \left(d_{x,t}^R \ln e_{x,t}^R + d_{x,t}^R \ln m_{x,t}^R - e_{x,t}^R \, m_{x,t}^R - \ln\left(d_{x,t}^R!\right) \right), \tag{A9}$$

$$l^B = \sum_{x,t} \left(d_{x,t}^B \ln e_{x,t}^B + d_{x,t}^B \ln m_{x,t}^B - e_{x,t}^B \, m_{x,t}^B - \ln\left(d_{x,t}^B!\right) \right), \tag{A10}$$

in which $d_{x,t}^R$ and $d_{x,t}^B$ are the observed numbers of deaths.

Each parameter in equation (A1) is estimated using the updating equation $\theta^* = \theta - \frac{\partial l}{\partial \theta} \big/ \frac{\partial^2 l}{\partial \theta^2}$ (Brouhns et al. 2002). The iterative updating scheme for estimating the parameters is summarised as follows:

(1)   Set all the initial parameter values ($\kappa_{t,1}^R$, $\kappa_{t,2}^R$, $\kappa_{t,3}^R$, and γ_c^R) as zero and calculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t, which are the fitted values computed by putting the parameter values into equations (A1) and (A7).

(2)   Update $\kappa_{t,1}^R$ for all t.

(3)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then update $\kappa_{t,2}^R$ for all t.

(4)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then update $\kappa_{t,3}^R$ for all t.

(5)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then update γ_c^R for all c.

(6)   Adjust $\kappa_{t,1}^R$, $\kappa_{t,2}^R$, $\kappa_{t,3}^R$, and γ_c^R for all t and c to incorporate the three identifiability constraints.

(7)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then calculate l^R.

(8)   Repeat steps (2) to (7) until the improvement in l^R is less than 10^{-11}.

Taking the estimated parameters $\hat{\kappa}_{t,1}^R$, $\hat{\kappa}_{t,2}^R$, $\hat{\kappa}_{t,3}^R$, and $\hat{\gamma}_c^R$ from above as given, the parameters in equation (A2) are estimated via the updating equation $\theta^* = \theta - \frac{\partial l}{\partial \theta} \big/ \frac{\partial^2 l}{\partial \theta^2}$ again and the iterative updating scheme below:

(1)   Set all the initial parameter values ($\kappa_{t,1}^B$ and $\kappa_{t,2}^B$) as zero and calculate $\hat{\eta}_{x,t}^B$ and $\hat{m}_{x,t}^B$ for all x and t, which are the fitted values computed by putting the parameter values into equations (A2) and (A8).

(2)   Update $\kappa_{t,1}^B$ for all t.

(3)   Recalculate $\hat{\eta}_{x,t}^B$ and $\hat{m}_{x,t}^B$ for all x and t and then update $\kappa_{t,2}^B$ for all t.

(4)   Recalculate $\hat{\eta}_{x,t}^B$ and $\hat{m}_{x,t}^B$ for all x and t and then calculate l^B.

(5)   Repeat steps (2) to (4) until the improvement in l^B is less than 10^{-11}.

The various components in using the updating equations are derived as below (for $i = 1, 2, 3$):

$$\frac{\partial l^R}{\partial \kappa_{t,i}^R} = \sum_x \left(\frac{d_{x,t}^R}{m_{x,t}^R} - e_{x,t}^R \right) \frac{\partial m_{x,t}^R}{\partial \kappa_{t,i}^R}, \qquad \frac{\partial^2 l^R}{\partial \left(\kappa_{t,i}^R \right)^2} = \sum_x \left(d_{x,t}^R \frac{m_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\kappa_{t,i}^R \right)^2} - \left(\frac{\partial m_{x,t}^R}{\partial \kappa_{t,i}^R} \right)^2}{\left(m_{x,t}^R \right)^2} - e_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\kappa_{t,i}^R \right)^2} \right),$$

$$\frac{\partial l^R}{\partial \gamma_{t-x}^R} = \sum_{\substack{\text{cohort} \\ t-x}} \left(\frac{d_{x,t}^R}{m_{x,t}^R} - e_{x,t}^R \right) \frac{\partial m_{x,t}^R}{\partial \gamma_{t-x}^R}, \qquad \frac{\partial^2 l^R}{\partial \left(\gamma_{t-x}^R \right)^2} = \sum_{\substack{\text{cohort} \\ t-x}} \left(d_{x,t}^R \frac{m_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\gamma_{t-x}^R \right)^2} - \left(\frac{\partial m_{x,t}^R}{\partial \gamma_{t-x}^R} \right)^2}{\left(m_{x,t}^R \right)^2} - e_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\gamma_{t-x}^R \right)^2} \right),$$

$$\frac{\partial m_{x,t}^R}{\partial \kappa_{t,1}^R} = q_{x,t}^R, \qquad\qquad \frac{\partial^2 m_{x,t}^R}{\partial \left(\kappa_{t,1}^R \right)^2} = \exp\left(-\eta_{x,t}^R \right) \left(q_{x,t}^R \right)^2,$$

$$\frac{\partial m_{x,t}^R}{\partial \kappa_{t,2}^R} = (x - \bar{x}) q_{x,t}^R, \qquad\qquad \frac{\partial^2 m_{x,t}^R}{\partial \left(\kappa_{t,2}^R \right)^2} = (x - \bar{x})^2 \exp\left(-\eta_{x,t}^R \right) \left(q_{x,t}^R \right)^2,$$

$$\frac{\partial m_{x,t}^R}{\partial \kappa_{t,3}^R} = \left((x - \bar{x})^2 - \sigma^2 \right) q_{x,t}^R, \qquad \frac{\partial^2 m_{x,t}^R}{\partial \left(\kappa_{t,3}^R \right)^2} = \left((x - \bar{x})^2 - \sigma^2 \right)^2 \exp\left(-\eta_{x,t}^R \right) \left(q_{x,t}^R \right)^2,$$

$$\frac{\partial m_{x,t}^R}{\partial \gamma_{t-x}^R} = q_{x,t}^R, \qquad\qquad \frac{\partial^2 m_{x,t}^R}{\partial \left(\gamma_{t-x}^R \right)^2} = \exp\left(-\eta_{x,t}^R \right) \left(q_{x,t}^R \right)^2,$$

$$\frac{\partial l^B}{\partial \kappa_{t,i}^B} = \sum_x \left(\frac{d_{x,t}^B}{m_{x,t}^B} - e_{x,t}^B \right) \frac{\partial m_{x,t}^B}{\partial \kappa_{t,i}^B}, \qquad \frac{\partial^2 l^B}{\partial \left(\kappa_{t,i}^B \right)^2} = \sum_x \left(d_{x,t}^B \frac{m_{x,t}^B \frac{\partial^2 m_{x,t}^B}{\partial \left(\kappa_{t,i}^B \right)^2} - \left(\frac{\partial m_{x,t}^B}{\partial \kappa_{t,i}^B} \right)^2}{\left(m_{x,t}^B \right)^2} - e_{x,t}^B \frac{\partial^2 m_{x,t}^B}{\partial \left(\kappa_{t,i}^B \right)^2} \right),$$

$$\frac{\partial m_{x,t}^B}{\partial \kappa_{t,1}^B} = q_{x,t}^B, \qquad\qquad \frac{\partial^2 m_{x,t}^B}{\partial \left(\kappa_{t,1}^B \right)^2} = \exp\left(-\eta_{x,t}^B - \eta_{x,t}^R \right) \left(q_{x,t}^B \right)^2,$$

$$\frac{\partial m_{x,t}^B}{\partial \kappa_{t,2}^B} = (x - \bar{x}) q_{x,t}^B, \qquad\qquad \frac{\partial^2 m_{x,t}^B}{\partial \left(\kappa_{t,2}^B \right)^2} = (x - \bar{x})^2 \exp\left(-\eta_{x,t}^B - \eta_{x,t}^R \right) \left(q_{x,t}^B \right)^2.$$

The three constraints are imposed by finding the values of λ_1, λ_2, and λ_3 and then using them to adjust the parameters $\kappa_{t,1}^R$, $\kappa_{t,2}^R$, $\kappa_{t,3}^R$, and γ_c^R with the following equations:

$$\begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix} = \begin{pmatrix} \sum_c 1 & \sum_c c & \sum_c c^2 \\ \sum_c c & \sum_c c^2 & \sum_c c^3 \\ \sum_c c^2 & \sum_c c^3 & \sum_c c^4 \end{pmatrix}^{-1} \begin{pmatrix} \sum_c \gamma_c^R \\ \sum_c c\, \gamma_c^R \\ \sum_c c^2 \gamma_c^R \end{pmatrix}, \tag{A11}$$

adjusted $\kappa_{t,1}^R = \kappa_{t,1}^R + \lambda_1 + \lambda_2 (t - \bar{x}) + \lambda_3 \left((t - \bar{x})^2 + \sigma^2 \right),$ \hfill (A12)

adjusted $\kappa_{t,2}^R = \kappa_{t,2}^R - \lambda_2 - 2\lambda_3(t - \bar{x})$, $\hspace{3cm}$ (A13)

adjusted $\kappa_{t,3}^R = \kappa_{t,3}^R + \lambda_3$, $\hspace{4.5cm}$ (A14)

adjusted $\gamma_{t-x}^R = \gamma_{t-x}^R - \lambda_1 - \lambda_2(t - x) - \lambda_3(t - x)^2$. $\hspace{2cm}$ (A15)

Similar computation algorithms can readily be derived if the binomial distribution is used rather than the Poisson distribution. For instance, it may be assumed approximately that $D_{x,t}^R \sim \text{Binomial}\left(e_{x,t}^R, q_{x,t}^R\right)$ and $D_{x,t}^B \sim \text{Binomial}\left(e_{x,t}^B, q_{x,t}^B\right)$, where $e_{x,t}^R$ and $e_{x,t}^B$ represent the initial exposed to risk measures instead, although these exposures are usually non-integers.

CAE+Cohorts Model

The CAE+Cohorts model also has two main components:

$$\text{logit } q_{x,t}^R = \alpha_x^R + \beta_x^R \kappa_t^R + \gamma_{t-x}^R = \eta_{x,t}^R \hspace{2cm} \text{(A16)}$$

$$\text{logit } q_{x,t}^B - \text{logit } q_{x,t}^R = \alpha_x^B + \beta_x^R \kappa_t^B = \eta_{x,t}^B \hspace{2cm} \text{(A17)}$$

in which for the first component, $q_{x,t}^R$ is the mortality rate of the reference population at age x in year t, α_x^R describes the mortality schedule over age x, κ_t^R is called the mortality index which reflects the overall mortality improvement over time t, with β_x^R as the age-specific sensitivity measure, and γ_{t-x}^R refers to the cohort effect of those lives born in year $t - x$. For the second component, the difference in the logit mortality rate between the book and reference populations, $\text{logit } q_{x,t}^B - \text{logit } q_{x,t}^R$, is modelled as another Lee-Carter structure with the parameters α_x^B, β_x^R, and κ_t^B. The two model components are then denoted as $\eta_{x,t}^R$ and $\eta_{x,t}^B$. Note that there are five identifiability constraints $\sum_x \beta_x^R = 1$, $\sum_t \kappa_t^R = 0$, $\sum_c \gamma_c^R = 0$,

$\sum_c (c - \bar{c}) \gamma_c^R = 0$, and $\sum_t \kappa_t^B = 0$, in which $\bar{c} = \dfrac{1}{\text{no. of cohorts}} \sum_c c$. These constraints are set to make sure that the parameters have unique solutions. The mortality rates, numbers of deaths, central death rates, and (Poisson) log likelihood functions are treated in the same way as in equations (A3) to (A10). Similar computation algorithms can readily be derived if the binomial distribution is assumed instead of the Poisson distribution.

Each parameter in equation (A16) is computed by the updating equation as above. The iterative updating scheme for parameter estimation is given below:

(1) Set the initial parameter values of all α_x^R as the average (over time) logit mortality rate observed at age x, β_x^R as $\frac{1}{\text{no. of ages}}$, κ_t^R as zero, and γ_c^R as zero. Then calculate the fitted values $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t by incorporating the parameter values into equations (A16) and (A7).

(2)   Update α_x^R for all x.

(3)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then update κ_t^R for all t.

(4)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then update β_x^R for all x.

(5)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then update γ_c^R for all c.

(6)   Adjust α_x^R, β_x^R, κ_t^R, and γ_c^R for all x, t, and c to incorporate the first four identifiability constraints.

(7)   Recalculate $\hat{\eta}_{x,t}^R$ and $\hat{m}_{x,t}^R$ for all x and t and then calculate l^R.

(8)   Repeat steps (2) to (7) until the improvement in l^R is less than 10^{-11}.

Treating the computed parameters $\hat{\alpha}_x^R$, $\hat{\beta}_x^R$, $\hat{\kappa}_t^R$, and $\hat{\gamma}_c^R$ from above as given, the parameters in equation (A17) are computed via the same updating equation, with the iterative updating scheme as stated in the following:

(1)   Set the initial parameter values of all α_x^B as the average (over time) logit mortality rate observed at age x minus $\hat{\alpha}_x^R$ and κ_t^B as zero. Then calculate the fitted values $\hat{\eta}_{x,t}^B$ and $\hat{m}_{x,t}^B$ for all x and t by incorporating the parameter values into equations (A17) and (A8).

(2)   Update α_x^B for all x.

(3)   Recalculate $\hat{\eta}_{x,t}^B$ and $\hat{m}_{x,t}^B$ for all x and t and then update κ_t^B for all t.

(4)   Adjust α_x^B and κ_t^B for all x and t to incorporate the last identifiability constraint.

(5)   Recalculate $\hat{\eta}_{x,t}^B$ and $\hat{m}_{x,t}^B$ for all x and t and then calculate l^B.

(6)   Repeat steps (2) to (5) until the improvement in l^B is less than 10^{-11}.

The different components of the updating equations are derived as follows:

$$\frac{\partial l^R}{\partial \alpha_x^R} = \sum_t \left(\frac{d_{x,t}^R}{m_{x,t}^R} - e_{x,t}^R \right) \frac{\partial m_{x,t}^R}{\partial \alpha_x^R}, \quad \frac{\partial^2 l^R}{\partial \left(\alpha_x^R\right)^2} = \sum_t \left(d_{x,t}^R \frac{m_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\alpha_x^R\right)^2} - \left(\frac{\partial m_{x,t}^R}{\partial \alpha_x^R} \right)^2}{\left(m_{x,t}^R\right)^2} - e_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\alpha_x^R\right)^2} \right),$$

$$\frac{\partial l^R}{\partial \kappa_t^R} = \sum_x \left(\frac{d_{x,t}^R}{m_{x,t}^R} - e_{x,t}^R \right) \frac{\partial m_{x,t}^R}{\partial \kappa_t^R}, \quad \frac{\partial^2 l^R}{\partial \left(\kappa_t^R\right)^2} = \sum_x \left(d_{x,t}^R \frac{m_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\kappa_t^R\right)^2} - \left(\frac{\partial m_{x,t}^R}{\partial \kappa_t^R} \right)^2}{\left(m_{x,t}^R\right)^2} - e_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial \left(\kappa_t^R\right)^2} \right),$$

$$\frac{\partial l^R}{\partial \beta_x^R} = \sum_t \left(\frac{d_{x,t}^R}{m_{x,t}^R} - e_{x,t}^R \right) \frac{\partial m_{x,t}^R}{\partial \beta_x^R}, \qquad \frac{\partial^2 l^R}{\partial (\beta_x^R)^2} = \sum_t \left(d_{x,t}^R \frac{m_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial (\beta_x^R)^2} - \left(\frac{\partial m_{x,t}^R}{\partial \beta_x^R} \right)^2}{(m_{x,t}^R)^2} - e_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial (\beta_x^R)^2} \right),$$

$$\frac{\partial l^R}{\partial \gamma_{t-x}^R} = \sum_{\substack{\text{cohort} \\ t-x}} \left(\frac{d_{x,t}^R}{m_{x,t}^R} - e_{x,t}^R \right) \frac{\partial m_{x,t}^R}{\partial \gamma_{t-x}^R}, \qquad \frac{\partial^2 l^R}{\partial (\gamma_{t-x}^R)^2} = \sum_{\substack{\text{cohort} \\ t-x}} \left(d_{x,t}^R \frac{m_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial (\gamma_{t-x}^R)^2} - \left(\frac{\partial m_{x,t}^R}{\partial \gamma_{t-x}^R} \right)^2}{(m_{x,t}^R)^2} - e_{x,t}^R \frac{\partial^2 m_{x,t}^R}{\partial (\gamma_{t-x}^R)^2} \right),$$

$$\frac{\partial m_{x,t}^R}{\partial \alpha_x^R} = q_{x,t}^R, \qquad\qquad \frac{\partial^2 m_{x,t}^R}{\partial (\alpha_x^R)^2} = \exp(-\eta_{x,t}^R)(q_{x,t}^R)^2,$$

$$\frac{\partial m_{x,t}^R}{\partial \kappa_t^R} = \beta_x^R \, q_{x,t}^R, \qquad\qquad \frac{\partial^2 m_{x,t}^R}{\partial (\kappa_t^R)^2} = (\beta_x^R)^2 \exp(-\eta_{x,t}^R)(q_{x,t}^R)^2,$$

$$\frac{\partial m_{x,t}^R}{\partial \beta_x^R} = \kappa_t^R \, q_{x,t}^R, \qquad\qquad \frac{\partial^2 m_{x,t}^R}{\partial (\beta_x^R)^2} = (\kappa_t^R)^2 \exp(-\eta_{x,t}^R)(q_{x,t}^R)^2,$$

$$\frac{\partial m_{x,t}^R}{\partial \gamma_{t-x}^R} = q_{x,t}^R, \qquad\qquad \frac{\partial^2 m_{x,t}^R}{\partial (\gamma_{t-x}^R)^2} = \exp(-\eta_{x,t}^R)(q_{x,t}^R)^2,$$

$$\frac{\partial l^B}{\partial \alpha_x^B} = \sum_t \left(\frac{d_{x,t}^B}{m_{x,t}^B} - e_{x,t}^B \right) \frac{\partial m_{x,t}^B}{\partial \alpha_x^B}, \qquad \frac{\partial^2 l^B}{\partial (\alpha_x^B)^2} = \sum_t \left(d_{x,t}^B \frac{m_{x,t}^B \frac{\partial^2 m_{x,t}^B}{\partial (\alpha_x^B)^2} - \left(\frac{\partial m_{x,t}^B}{\partial \alpha_x^B} \right)^2}{(m_{x,t}^B)^2} - e_{x,t}^B \frac{\partial^2 m_{x,t}^B}{\partial (\alpha_x^B)^2} \right),$$

$$\frac{\partial l^B}{\partial \kappa_t^B} = \sum_x \left(\frac{d_{x,t}^B}{m_{x,t}^B} - e_{x,t}^B \right) \frac{\partial m_{x,t}^B}{\partial \kappa_t^B}, \qquad \frac{\partial^2 l^B}{\partial (\kappa_t^B)^2} = \sum_x \left(d_{x,t}^B \frac{m_{x,t}^B \frac{\partial^2 m_{x,t}^B}{\partial (\kappa_t^B)^2} - \left(\frac{\partial m_{x,t}^B}{\partial \kappa_t^B} \right)^2}{(m_{x,t}^B)^2} - e_{x,t}^B \frac{\partial^2 m_{x,t}^B}{\partial (\kappa_t^B)^2} \right),$$

$$\frac{\partial m_{x,t}^B}{\partial \alpha_x^B} = q_{x,t}^B, \qquad\qquad \frac{\partial^2 m_{x,t}^B}{\partial (\alpha_x^B)^2} = \exp(-\eta_{x,t}^B - \eta_{x,t}^R)(q_{x,t}^B)^2,$$

$$\frac{\partial m_{x,t}^B}{\partial \kappa_t^B} = \beta_x^R \, q_{x,t}^B, \qquad\qquad \frac{\partial^2 m_{x,t}^B}{\partial (\kappa_t^B)^2} = (\beta_x^R)^2 \exp(-\eta_{x,t}^B - \eta_{x,t}^R)(q_{x,t}^B)^2.$$

The first four constraints are imposed by using the following equations:

$$\tilde{\alpha}_x^R = \alpha_x^R + \beta_x^R \frac{1}{\text{no. of years for ref}} \sum_t \kappa_t^R + \frac{1}{\text{no. of cohorts for ref}} \sum_c \gamma_c^R, \qquad (A18)$$

$$\tilde{\beta}_x^R = \frac{1}{\sum_x \beta_x^R}\beta_x^R,$$ (A19)

$$\tilde{\kappa}_t^R = \sum_x \beta_x^R\left(\kappa_t^R - \frac{1}{\text{no. of years for ref}}\sum_t \kappa_t^R\right),$$ (A20)

$$\tilde{\gamma}_{t-x}^R = \gamma_{t-x}^R - \frac{1}{\text{no. of cohorts for ref}}\sum_c \gamma_c^R,$$ (A21)

$$\tilde{\kappa}_t^R = g\left(t-\bar{t}\right)+\delta_t, \qquad (\; g = \frac{\sum_t \left(t-\bar{t}\right)\tilde{\kappa}_t^R}{\sum_t \left(t-\bar{t}\right)^2}\; ; \; \bar{t} = \frac{1}{\text{no. of years for ref}}\sum_t t\;)$$ (A22)

$$\text{adjusted } \alpha_x^R = \tilde{\alpha}_x^R + \frac{h}{\sum_x 1}\left(x-\bar{x}\right), \qquad (\; h = -\frac{\sum_x 1 \sum_c \left(c-\bar{c}\right)\tilde{\gamma}_c^R}{\sum_c \left(c-\bar{c}\right)^2}\;)$$ (A23)

$$\text{adjusted } \beta_x^R = \frac{g}{g-h}\tilde{\beta}_x^R - \frac{h}{\sum_x 1 (g-h)},$$ (A24)

$$\text{adjusted } \kappa_t^R = \frac{g-h}{g}\tilde{\kappa}_t^R,$$ (A25)

$$\text{adjusted } \gamma_{t-x}^R = \tilde{\gamma}_{t-x}^R + \frac{h}{\sum_x 1}\left(t-x-\bar{t}+\bar{x}\right).$$ (A26)

The last constraint is imposed with the equations below:

$$\text{adjusted } \alpha_x^B = \alpha_x^B + \beta_x^R \frac{1}{\text{no. of years for book}}\sum_t \kappa_t^B,$$ (A27)

$$\text{adjusted } \kappa_t^B = \kappa_t^B - \frac{1}{\text{no. of years for book}}\sum_t \kappa_t^B.$$ (A28)

Residuals Bootstrapping

To include both process error (variability in the time series) and parameter error (uncertainty in parameter estimation) in simulating future mortality rates, the residuals bootstrapping method (Koissi et al. 2006; Li 2014) is applied here. The bootstrapping procedure is detailed as follows:

(1)   The residuals from fitting the M7-M5 model or CAE+Cohorts model to the actual data are resampled with replacement. The residuals are resampled for each age-time cell within all x and t. Note that the standardised deviance residuals are calculated by the formulae $r_{x,t}^R = \frac{1}{\sqrt{\hat{\phi}^R}}\text{sign}\left(d_{x,t}^R - e_{x,t}^R\,\hat{m}_{x,t}^R\right)\sqrt{2\left(d_{x,t}^R \ln\left(d_{x,t}^R\big/\left(e_{x,t}^R\,\hat{m}_{x,t}^R\right)\right)-d_{x,t}^R + e_{x,t}^R\,\hat{m}_{x,t}^R\right)}$

and $r_{x,t}^B = \frac{1}{\sqrt{\hat{\phi}^B}}\text{sign}\left(d_{x,t}^B - e_{x,t}^B\,\hat{m}_{x,t}^B\right)\sqrt{2\left(d_{x,t}^B \ln\left(d_{x,t}^B\big/\left(e_{x,t}^B\,\hat{m}_{x,t}^B\right)\right)-d_{x,t}^B + e_{x,t}^B\,\hat{m}_{x,t}^B\right)}$ for   the

reference and book populations respectively. The dispersion parameters are estimated

by the two formulae $\hat{\phi}^R = \frac{1}{n_d^R - n_p^R} \sum_{x,t} 2\left(d_{x,t}^R \ln\left(d_{x,t}^R / \left(e_{x,t}^R \hat{m}_{x,t}^R\right)\right) - d_{x,t}^R + e_{x,t}^R \hat{m}_{x,t}^R\right)$ and

$\hat{\phi}^B = \frac{1}{n_d^B - n_p^B} \sum_{x,t} 2\left(d_{x,t}^B \ln\left(d_{x,t}^B / \left(e_{x,t}^B \hat{m}_{x,t}^B\right)\right) - d_{x,t}^B + e_{x,t}^B \hat{m}_{x,t}^B\right)$, in which n_d^R (n_d^B) and n_p^R (n_p^B) are the number of data points and the number of effective parameters of the reference (book) population.

(2)   The inverse functions of the residuals formulae above are used to turn the resampled residuals into a pseudo sample of the number of deaths, $d_{x,t}^{R\,(i)}$ and $d_{x,t}^{B\,(i)}$, for all x and t, where (i) denotes the i^{th} iteration or scenario.

(3)   The M7-M5 model or CAE+Cohorts model is fitted to the pseudo data sample from step (2) and the corresponding model parameters ($\kappa_{t,1}^{R\,(i)}$, $\kappa_{t,2}^{R\,(i)}$, $\kappa_{t,3}^{R\,(i)}$, $\gamma_{t-x}^{R\,(i)}$, $\kappa_{t,1}^{B\,(i)}$, $\kappa_{t,2}^{B\,(i)}$, or $\alpha_x^{R\,(i)}$, $\beta_x^{R\,(i)}$, $\kappa_t^{R\,(i)}$, $\gamma_{t-x}^{R\,(i)}$, $\alpha_x^{B\,(i)}$, $\kappa_t^{B\,(i)}$) are computed using the iterative updating schemes noted earlier.

(4)   The time series processes are fitted to the temporal model parameters based on the pseudo data sample ($\kappa_{t,1}^{R\,(i)}$, $\kappa_{t,2}^{R\,(i)}$, $\kappa_{t,3}^{R\,(i)}$, $\gamma_{t-x}^{R\,(i)}$, $\kappa_{t,1}^{B\,(i)}$, $\kappa_{t,2}^{B\,(i)}$, or $\kappa_t^{R\,(i)}$, $\gamma_{t-x}^{R\,(i)}$, $\kappa_t^{B\,(i)}$) from step (3) to simulate their future values over time.

(5)   Samples of future mortality rates, $q_{x,t}^{R\,(i)}$ and $q_{x,t}^{B\,(i)}$, for all x and future t, are generated from incorporating the computed parameters and simulated values based on the pseudo data sample from steps (3) and (4) into equations (A1) and (A2) or (A16) and (A17). This set of future mortality rates represents one random future scenario.

(6)   Steps (1) to (5) are repeated to produce 5,000 random future scenarios.

(7)   For each random scenario, the future number of survivors in the pension portfolio over time is simulated as $l_{x+1,t+1}^{B\,(i)} \sim \text{Binomial}\left(l_{x,t}^{B\,(i)}, 1 - q_{x,t}^{B\,(i)}\right)$.

Approximation Methods

Using Excel VBA, the computation time of residuals bootstrapping in generating 5,000 random future scenarios can be up to a day or two. If the software R is used instead, the computation time is even longer, being about threefold in general. One possible way to speed up the process is by replacing some codes of the iterative updating schemes in Appendix II with certain matrix operations using in-built matrix functions in R. An alternative is to apply some approximate parameter estimation methods to replace the more precise iterative updating schemes. For instance, the M7-M5 model can be fitted first via a simple linear regression for each year t, like $\text{logit}\, q_{x,t}^R = \kappa_{t,1}^R + (x - \bar{x})\kappa_{t,2}^R + \left((x - \bar{x})^2 - \sigma_x^2\right)\kappa_{t,3}^R$ without the cohort parameter. Then taking the computed $\hat{\kappa}_{t,1}^R$, $\hat{\kappa}_{t,2}^R$, and $\hat{\kappa}_{t,3}^R$ as given, the residuals for all x and t are fitted with another simple linear regression as $\text{logit}\, q_{x,t}^R - \hat{\kappa}_{t,1}^R - (x - \bar{x})\hat{\kappa}_{t,2}^R - \left((x - \bar{x})^2 - \sigma_x^2\right)\hat{\kappa}_{t,3}^R = \sum_c I_c \gamma_c^R$, where I_c is an indicator variable which is one if $c = t - x$ and zero otherwise. Afterwards, the book component's parameters for each t are estimated

via the equation logit $q_{x,t}^B - \hat{\eta}_{x,t}^R = \kappa_{t,1}^B + (x - \bar{x})\kappa_{t,2}^B$, which is again a simple linear regression. Note that the central exposed to risk measures can be used as weights in the regression (i.e. weighted least squares) to improve the estimation. Furthermore, the semi-parametric bootstrapping method (Brouhns et al. 2005) can be adopted to simulate a pseudo sample of the number of deaths from the Poisson distribution with the observed number of deaths as the mean. Consequently, steps (1) to (3) of the previous residuals bootstrapping method are much simplified, and it is observed that the computation time can be shortened by more than half. It is also noted that this approximation method often gives rise to similar parameter values and risk reduction estimates to those produced by the more detailed original procedure. In the same way, for the CAE+Cohorts model, one may apply the singular value decomposition (SVD) (Lee and Carter 1992) or the principal component analysis (PCA) (Bell 1997) rather than the maximum likelihood to shorten the time in calculating the model parameters.

Tail Risk Measures

The value-at-risk (VaR) of a pension portfolio at confidence level α can be defined as:

$$\text{VaR}_\alpha(X) = F_X^{-1}(\alpha), \tag{A29}$$

in which X is either the present value of future cash outflows or the individual cash outflow in a particular future year, and $F_X^{-1}(\bullet)$ is the inverse cumulative distribution function of X. For instance, the 99.5% VaR can be estimated as the sample 99.5[th] percentile in the simulation. The expected shortfall (conditional VaR) of a pension portfolio at confidence level α can then be stated as:

$$\text{ES}_\alpha(X) = \frac{1}{1-\alpha}\int_\alpha^1 \text{VaR}_p(X)\,dp. \tag{A30}$$

Accordingly, the 99.5% expected shortfall is estimated as the sample mean of all the simulated values which exceed the sample 99.5[th] percentile.

As noted in Dowd and Blake (2006), there are a number of reasons behind the wide acceptance of the VaR. Firstly, it can serve as a common risk measure for different types of risk positions, unlike traditional methods such as the duration, Greek letters, and portfolio theory. The VaR also allows one to aggregate the risk positions and focus on the whole portfolio. Moreover, it is probabilistic and provides a simple concept which can easily be understood. However, the VaR provides no information about the worst possible losses, and it can potentially cause moral hazard problems when traders try to 'game' around the VaR requirement. Furthermore, theoretically it is not subadditive (and so not coherent), which means that it does not take into account diversification benefits properly. In fact, the VaR is not associated with any set of risk measure axioms, and it implies risk-loving under the expected utility downside risk framework.

In contrast, the expected shortfall (conditional VaR) has more information about the extreme events, better theoretical properties, and more reasonable implications. First, it is coherent and subadditive, and so it allows for diversification benefits. It is also in line with some common applications in insurance such as the excess-of-loss reinsurance, which is a familiar concept to actuaries. From a practical perspective, it can easily be estimated from a large number of simulated scenarios, and at the same confidence level, it generally provides a more conservative estimate than the VaR does. Recently, banking regulations are gradually moving towards the use of the expected shortfall in calculating regulatory capital for market risk (Hull 2015). Finally, besides the expected shortfall, there are also other decent alternatives like spectral risk measures and distortion risk measures (Dowd and Blake 2006).

Numerical Optimisation

A simple numerical optimisation procedure can be adopted to minimise longevity risk (and also reduce structural basis risk) and so maximise hedge effectiveness under the simulated environment. The numerical algorithms used by different software may vary and lead to slightly different hedging results, and the choice depends on the user's own preference. For the longevity swap in Subsection 4.1, the required notional amount of the swap is estimated from numerical optimisation with an objective to minimise the 99.5% VaR (minus the mean) of the present value of the aggregate pension portfolio position. As shown in the graph below, it is interesting to note that although the objective is based on the overall present value, it turns out that the risk reduction effect is still positive for the cash outflow of the aggregate position in each year, which is often not the case when using key q-duration (KQD) matching. In practice, other risk measures (e.g. the standard deviation and 99.5% expected shortfall as in Subsection 5.5) or a mix of different objectives (e.g. risk minimisation with a desired profitability level) may also be implemented, depending on the purpose of the analysis.

Level of longevity risk reduction (in % of initial longevity risk) of individual cash flows
100,000 UK male pensioners (basic materials; normal retirement; lower pension; M7-M5)
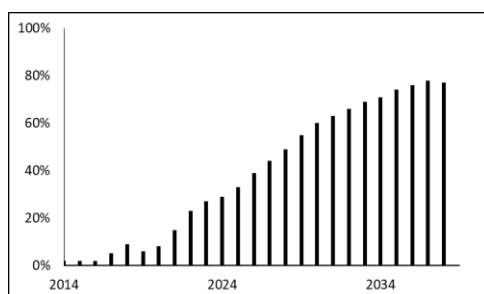
For the two longevity swaps in Subsection 4.2, the required notional amounts of the two swaps are estimated in the same way. In this optimisation exercise, there are then two unknown quantities to determine in order to achieve a single objective. If longevity swaps for more cohorts and with different maturities are also available in the life market and being

included in the hedge, the numerical optimisation procedure can readily be extended to find the required notional amounts of using multiple swaps simultaneously. However, when the number of unknown quantities to be estimated increases, there may be a higher chance of getting a local optimum rather than the true global optimum. Some simple methods in practice to alleviate this problem include using other sensible starting values, applying alternative numerical algorithms, simplifying the model settings where possible, and rerunning the optimisation procedure with the previous solutions as the new starting values.

The q-forwards and S-forwards in Subsection 5.4 have all maturities from 1 year to 25 years. In contrast to above, the forward contract of each maturity is calibrated separately by numerical optimisation with an objective to minimise the 99.5% VaR (minus the mean) of the cash outflow of the aggregate position in that particular year. Effectively, there are a total of 25 notional amounts to be computed for the different maturities of the forward contracts. In addition, these forwards may further be combined with certain longevity swaps in the hedge, such that the forwards positions are catered more for the individual cash flows while the swaps positions are structured mainly for the overall present value. This mix of different derivatives and objectives has the potential to yield better hedge effectiveness.

Analytical Formulae for Optimisation

We derive some analytical formulae below to find the required notional amounts for the S-forwards, q-forwards, and longevity swap when the objective is to minimise the variance but not the VaR of the aggregate position. Referring to the case study in Subsection 4.1, considering an infinitely large portfolio size, and using S-forwards instead, the cash outflow of the net position at each time $t = 1, 2, \dots, 25$ per pensioner at $t = 0$ can be expressed as:

$$\text{CF}_t\left(\text{hedged position}\right) = {}_t p_{65}^B - w_t \left({}_t p_{65}^R - {}_t p_{65}^{R;\text{forward}}\right).$$

The term ${}_t p_{65}^B = \left(1 - q_{65,0}^B\right)\left(1 - q_{66,1}^B\right)\dots\left(1 - q_{65+t-1,t-1}^B\right)$ is simulated as in Subsection 3.4 for the book population, under the assumption that the sampling variability in the future simulations is insignificant for a very large portfolio. So the variance of that cash outflow before hedging is $\text{Var}\left({}_t p_{65}^B\right)$, and after hedging it becomes:

$$\text{Var}\left({}_t p_{65}^B - w_t \, {}_t p_{65}^R\right) = \text{Var}\left({}_t p_{65}^B\right) - 2w_t \, \text{Cov}\left({}_t p_{65}^B, {}_t p_{65}^R\right) + w_t^2 \, \text{Var}\left({}_t p_{65}^R\right).$$

The weight w_t of the S-forward that minimises the variance is found by differentiating the equation above and then setting it to zero, which gives $w_t = \text{Cov}\left({}_t p_{65}^B, {}_t p_{65}^R\right) / \text{Var}\left({}_t p_{65}^R\right)$ and a longevity risk (variance) reduction of $\left[\text{Cor}\left({}_t p_{65}^B, {}_t p_{65}^R\right)\right]^2$. This analytical result has two important implications. First, the larger the correlation (in magnitude) between ${}_t p_{65}^B$ and ${}_t p_{65}^R$, the higher the level of longevity risk (variance) reduction. Second, this value of w_t can be used as a starting value for numerical optimisation when the objective is based on other risk measures and a similar solution cannot be derived analytically.

Alternatively, for a q-forward, the weight is $w_t = -\mathrm{Cov}\left(_t p_{65}^B, q_{65+t-1,t-1}^R\right) \Big/ \mathrm{Var}\left(q_{65+t-1,t-1}^R\right)$ and the risk reduction is $\left[\mathrm{Cor}\left(_t p_{65}^B, q_{65+t-1,t-1}^R\right)\right]^2$. If a longevity swap is used as in Subsection 4.1, the present value of the aggregate position per pensioner at $t=0$ is equal to:

$$\mathrm{PV}\left(\text{hedged portfolio}\right) = \sum_{t=1}^{25} {}_t p_{65}^B \left(1+i\right)^{-t} - w \sum_{t=1}^{25}\left({}_t p_{65}^R - {}_t p_{65}^{R;\text{forward}}\right)\left(1+i\right)^{-t}.$$

Similarly, the weight w of the swap that minimises the variance is derived as $w = \mathrm{Cov}\left(\sum_{t=1}^{25} {}_t p_{65}^B \left(1+i\right)^{-t}, \sum_{t=1}^{25} {}_t p_{65}^R \left(1+i\right)^{-t}\right) \Big/ \mathrm{Var}\left(\sum_{t=1}^{25} {}_t p_{65}^R \left(1+i\right)^{-t}\right)$, which results in a longevity risk (variance) reduction of $\left[\mathrm{Cor}\left(\sum_{t=1}^{25} {}_t p_{65}^B \left(1+i\right)^{-t}, \sum_{t=1}^{25} {}_t p_{65}^R \left(1+i\right)^{-t}\right)\right]^2$. Note that all the variance and covariance terms can be estimated directly from the simulated samples.

Key q-Duration Matching

The key q-duration (KQD) is defined as (Li and Luo 2012):

$$\mathrm{KQD} = \frac{\Delta \text{ in best estimate of present value}}{\Delta \text{ in best estimate of a particular future mortality rate}}. \tag{A31}$$

Considering the case study in Subsection 4.1, the random present value of future cash outflows of the pension portfolio for the next 25 years is expressed as:

$$\mathrm{PV}\left(\text{unhedged portfolio}\right) = \sum_{t=1}^{25} l_{65+t,t}^B \left(1+i\right)^{-t}.$$

Suppose q-forwards are available at key ages of 70, 75, 80, and 85 (i.e. maturities of 6, 11, 16, and 21 years). The random present value of future payoff (as a fixed rate receiver; based on the reference population) of a q-forward linked to key age 70 for the same birth cohort as the pensioners is stated as:

$$\mathrm{PV}\left(\text{hedging instrument}\right) = \left(q_{70,5}^{R;\text{forward}} - q_{70,5}^R\right)\left(1+i\right)^{-6}.$$

The random present values of the q-forwards for the other key ages are defined similarly. Note that for demonstration purposes the forward mortality rates $q_{x,t}^{R;\text{forward}}$ are calculated from the best (central) estimates of future mortality rates, i.e. setting all the time series error terms to zero in the central projection.

Then the random present value of the aggregate position after taking the longevity hedge is:

$$\mathrm{PV}\left(\text{hedged portfolio}\right) = \sum_{t=1}^{25} l_{65+t,t}^B \left(1+i\right)^{-t} - \sum_{t=5,10,15,20} w_t \left(q_{65+t,t}^{R;\text{forward}} - q_{65+t,t}^R\right)\left(1+i\right)^{-t-1}.$$

The weights w_t are the notional amounts of the four q-forwards and are calculated by equating the KQD of each q-forward position and the KQD of the portfolio to be hedged at each key age in turn:

$$\mathrm{KQD} \text{ of portfolio at key age } 65+t = \frac{\Delta \text{ in best estimate of present value of portfolio}}{\Delta \text{ in best estimate of } q_{65+t,t}^B},$$

KQD of q-forward linked to key age $65 + t = \dfrac{\Delta \text{ in best estimate of present value of } q\text{-forward}}{\Delta \text{ in best estimate of } q^B_{65+t,t}}$,

$\therefore w_t = \dfrac{\text{KQD of portfolio at key age } 65 + t}{\text{KQD of } q\text{-forward linked to key age } 65 + t}$.

Note that the q-forward KQD with respect to $q^B_{65+t,t}$ above can be estimated by the q-forward KQD with respect to $q^R_{65+t,t}$ (i.e. $-(1+i)^{-t-1}$) multiplied with an adjustment factor. The adjustment factor can be approximately calculated as the slope of plotting the simulated $q^R_{65+t,t}$ values against the simulated $q^B_{65+t,t}$ values. Using these weights, the aggregate position after taking the longevity hedge has an overall KQD of zero at each key age. That is, roughly speaking, a small deviation in a future mortality rate from expected would have little or no impact on the hedged portfolio. Even when the deviation is large, the longevity hedge would still reduce the resulting impact to some extent, though the precise effects would vary with the underlying conditions.

However, we find that while this KQD matching strategy can lead to reasonable hedge effectiveness regarding the present value of the aggregate position, it does not necessarily result in proper risk reduction for the individual cash flows. Based on our simulations, there are often very mixed results and rather offsetting effects between the cash flows in different years. More research is needed to inspect further on the practical limitations of this method.

Time Series Processes

Consider a multivariate time series of $k_{t,1}$, $k_{t,2}$, $k_{t,3}$, ... , $k_{t,n}$ over time t. A multivariate random walk with drift (MRWD) for the n-dimensional time series is defined as:

$$\begin{pmatrix} k_{t,1} \\ k_{t,2} \\ \vdots \end{pmatrix} = \begin{pmatrix} d_1 \\ d_2 \\ \vdots \end{pmatrix} + \begin{pmatrix} k_{t-1,1} \\ k_{t-1,2} \\ \vdots \end{pmatrix} + \begin{pmatrix} \varepsilon_{t,1} \\ \varepsilon_{t,2} \\ \vdots \end{pmatrix},$$

$$\begin{pmatrix} \varepsilon_{t,1} \\ \varepsilon_{t,2} \\ \vdots \end{pmatrix} \sim N(0, \Sigma), \qquad \Sigma = \begin{pmatrix} \sigma^2_{\varepsilon,1} & \sigma_{\varepsilon,1,2} & \cdots \\ \sigma_{\varepsilon,2,1} & \sigma^2_{\varepsilon,2} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}. \tag{A32}$$

The parameters d_j for $j = 1, 2, 3, ..., n$ are the drift terms, and $\varepsilon_{t,j}$ are the multivariate normal error terms with mean zero and covariance matrix Σ (with variances $\sigma^2_{\varepsilon,j}$ and covariances $\sigma_{\varepsilon,i,j}$ for $i, j = 1, 2, 3, ..., n$). The drift term d_j and variance $\sigma^2_{\varepsilon,j}$ are estimated as the sample mean and sample variance of the first differences $\Delta k_{t,j} = k_{t,j} - k_{t-1,j}$. The covariance $\sigma_{\varepsilon,i,j}$ is estimated as the sample covariance between $\Delta k_{t,i}$ and $\Delta k_{t,j}$. (Alternatively, to avoid the drift term being solely determined by the starting and ending observed values, a multivariate linear

regression may be applied to $k_{t,1}$, $k_{t,2}$, $k_{t,3}$, ..., $k_{t,n}$ against time and the slopes are taken as the drift parameters.) When $n = 1$, the MRWD reduces to a univariate random walk with drift (RWD), and if further $d_1 = 0$, it becomes a random walk *without* drift (RW). When all $\sigma_{\varepsilon,i,j} = 0$, the MRWD reduces to a set of independent univariate RWD processes.

Furthermore, a vector autoregressive integrated moving average process, VARIMA(1,1,0), is defined as:

$$\begin{pmatrix} \Delta k_{t,1} \\ \Delta k_{t,2} \\ \vdots \end{pmatrix} = \begin{pmatrix} \varphi_{1,0} \\ \varphi_{2,0} \\ \vdots \end{pmatrix} + \begin{pmatrix} \varphi_{1,1} & \varphi_{1,2} & \cdots \\ \varphi_{2,1} & \varphi_{2,2} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} \Delta k_{t-1,1} \\ \Delta k_{t-1,2} \\ \vdots \end{pmatrix} + \begin{pmatrix} \xi_{t,1} \\ \xi_{t,2} \\ \vdots \end{pmatrix},$$

$$\begin{pmatrix} \xi_{t,1} \\ \xi_{t,2} \\ \vdots \end{pmatrix} \sim N(0, \Psi), \qquad \Psi = \begin{pmatrix} \sigma_{\xi,1}^2 & \sigma_{\xi,1,2} & \cdots \\ \sigma_{\xi,2,1} & \sigma_{\xi,2}^2 & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}. \tag{A33}$$

The parameters $\varphi_{i,0}$ for $i = 1, 2, 3, ..., n$ are the constant terms, $\varphi_{i,j}$ for $i, j = 1, 2, 3, ..., n$ are the autoregressive parameters, and $\xi_{t,i}$ are the multivariate normal error terms with mean zero and covariance matrix Ψ (with variances $\sigma_{\xi,i}^2$ and covariances $\sigma_{\xi,i,j}$ for $i, j = 1, 2, 3, ..., n$). The model parameters can be computed by maximising the conditional likelihood function (Tsay 2002). In practice, for convenience, the parameters can also be estimated from performing a simple linear regression of $\Delta k_{t,i}$ on the vector $(\Delta k_{t-1,1}, \Delta k_{t-1,2}, ...)$ for each $i = 1, 2, 3, ..., n$ in turn. The parameters $\varphi_{i,0}$ and $\varphi_{i,1}$, $\varphi_{i,2}$, ... are then taken as the calculated intercept and coefficients of the regression line, and $\sigma_{\xi,i}^2$ is estimated as the variance of the regression error term. Then the covariance $\sigma_{\xi,i,j}$ can simply be estimated as the sample covariance between the residuals from the regressions of $\Delta k_{t,i}$ and $\Delta k_{t,j}$.

When $n = 1$, the VARIMA(1,1,0) reduces to a univariate autoregressive integrated moving average process, ARIMA(1,1,0). If the first differencing is not adopted, i.e. $k_{t,i}$ is used instead of $\Delta k_{t,i}$ in the model structure, the process becomes a vector autoregressive process of order one, VAR(1), and if further $n = 1$, it reduces to an autoregressive process of order one, AR(1). In addition, an autoregressive process of order p, AR(p), is a generalisation of the AR(1), in which the univariate k_t is expressed as a linear function of k_{t-1}, k_{t-2}, ..., k_{t-p}. In theory, the order p can be determined by inspecting the Akaike information criterion (AIC), Bayesian information criterion (BIC), partial autocorrelation function (PACF), and residuals' randomness. But since the data length of the book population is short and there is a need to ensure weak stationarity and so the coherence (Subsection 6.3), some arbitrary judgement would be required to choose the order p when the AR(p) is used for the book.

Appendix II   Excel VBA Coding

The following is a demonstration of how the key computation algorithms in Appendix I can be programmed with Excel VBA. Suppose the age range is 60 to 89, the data period of the reference population is 1980 to 2013, and the data period of the book population is 2000 to 2013.

M7-M5 Model

Let dref(x,t), eref(x,t), mref(x,t), qref(x,t), etaref(x,t), resref(x,t), k1ref(t), k2ref(t), k3ref(t), and gref(c) denote $d_{x,t}^{R}$, $e_{x,t}^{R}$, $m_{x,t}^{R}$, $q_{x,t}^{R}$, $\eta_{x,t}^{R}$, $r_{x,t}^{R}$, $\kappa_{t,1}^{R}$, $\kappa_{t,2}^{R}$, $\kappa_{t,3}^{R}$, and γ_{c}^{R} respectively for the reference component, and let dbook(x,t), ebook(x,t), mbook(x,t), qbook(x,t), etabook(x,t), resbook(x,t), k1book(t), and k2book(t) refer to $d_{x,t}^{B}$, $e_{x,t}^{B}$, $m_{x,t}^{B}$, $q_{x,t}^{B}$, $\eta_{x,t}^{B}$, $r_{x,t}^{B}$, $\kappa_{t,1}^{B}$, and $\kappa_{t,2}^{B}$ for the book component. The Excel VBA codes for fitting the M7-M5 model are given below.

```
' set initial values of reference kappa1, kappa2, kappa3, gamma

For t = 1980 To 2013
k1ref(t) = 0
k2ref(t) = 0
k3ref(t) = 0
Next t


For c = 1891 To 1953
gref(c) = 0
Next c

' iteratively update values of reference kappa1, kappa2, kappa3, gamma

lnLref = -10000000
lnLpreref = -11000000
Do Until lnLref - lnLpreref < 0.00000000001

' update reference kappa1

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = qref(x, t)
d2m(x, t) = Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x
```

```
For t = 1980 To 2013
dl = 0
d2l = 0
For x = 60 To 89
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
Next x
k1ref(t) = k1ref(t) - dl / d2l
Next t

' update reference kappa2

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = (x - 74.5) * qref(x, t)
d2m(x, t) = (x - 74.5) ^ 2 * Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x


For t = 1980 To 2013
dl = 0
d2l = 0
For x = 60 To 89
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
Next x
k2ref(t) = k2ref(t) - dl / d2l
Next t

' update reference kappa3

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = ((x - 74.5) ^ 2 - 2247.5 / 30) * qref(x, t)
d2m(x, t) = ((x - 74.5) ^ 2 - 2247.5 / 30) ^ 2 * Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x


For t = 1980 To 2013
dl = 0
d2l = 0
For x = 60 To 89
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
Next x
```

```
k3ref(t) = k3ref(t) - dl / d2l
Next t

' update reference gamma

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = qref(x, t)
d2m(x, t) = Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x

For c = 1891 To 1953
dl = 0
d2l = 0
For x = 60 To 89
For t = 1980 To 2013
If c = t - x Then
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
End If
Next t
Next x
gref(c) = gref(c) - dl / d2l
Next c

' incorporate reference cohort constraints

sumc = 0
sumc1 = 0
sumc2 = 0
sumc3 = 0
sumc4 = 0
sumg = 0
sumcg = 0
sumc2g = 0
For c = 1891 To 1953
sumc = sumc + 1
sumc1 = sumc1 + c
sumc2 = sumc2 + c ^ 2
sumc3 = sumc3 + c ^ 3
sumc4 = sumc4 + c ^ 4
sumg = sumg + gref(c)
sumcg = sumcg + c * gref(c)
sumc2g = sumc2g + c ^ 2 * gref(c)
Next c
```

```vba
cmatrix(1, 1) = sumc
cmatrix(1, 2) = sumc1
cmatrix(1, 3) = sumc2
cmatrix(2, 1) = sumc1
cmatrix(2, 2) = sumc2
cmatrix(2, 3) = sumc3
cmatrix(3, 1) = sumc2
cmatrix(3, 2) = sumc3
cmatrix(3, 3) = sumc4

gvector(1, 1) = sumg
gvector(2, 1) = sumcg
gvector(3, 1) = sumc2g

lambda = WorksheetFunction.MMult(WorksheetFunction.MInverse(cmatrix), gvector)

For t = 1980 To 2013
k1ref(t) = k1ref(t) + lambda(1, 1) + lambda(2, 1) * (t - 74.5) + lambda(3, 1) * ((t - 74.5) ^ 2 + 2247.5 / 30)
k2ref(t) = k2ref(t) - lambda(2, 1) - 2 * lambda(3, 1) * (t - 74.5)
k3ref(t) = k3ref(t) + lambda(3, 1)
Next t

For c = 1891 To 1953
gref(c) = gref(c) - lambda(1, 1) - lambda(2, 1) * c - lambda(3, 1) * c ^ 2
Next c

' calculate reference log likelihood

lnLpreref = lnLref
lnLref = 0
For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
lnLref = lnLref + dref(x, t) * WorksheetFunction.Ln(eref(x, t)) + dref(x, t) * WorksheetFunction.Ln(mref(x, t)) -
eref(x, t) * mref(x, t) - WorksheetFunction.GammaLn(dref(x, t) + 1)
Next t
Next x

Loop

' set initial values of book kappa1, kappa2

For t = 2000 To 2013
k1book(t) = 0
k2book(t) = 0
Next t
```

```
' iteratively update values of book kappa1, kappa2

lnLbook = -10000000
lnLprebook = -11000000
Do Until lnLbook - lnLprebook < 0.00000000001

' update book kappa1

For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
etabook(x, t) = k1book(t) + (x - 74.5) * k2book(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
dm(x, t) = qbook(x, t)
d2m(x, t) = Exp(-etabook(x, t) - etaref(x, t)) * qbook(x, t) ^ 2
Next t
Next x

For t = 2000 To 2013
dl = 0
d2l = 0
For x = 60 To 89
dl = dl + (dbook(x, t) / mbook(x, t) - ebook(x, t)) * dm(x, t)
d2l = d2l + dbook(x, t) * (mbook(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mbook(x, t) ^ 2 - ebook(x, t) * d2m(x, t)
Next x
k1book(t) = k1book(t) - dl / d2l
Next t

' update book kappa2

For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
etabook(x, t) = k1book(t) + (x - 74.5) * k2book(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
dm(x, t) = (x - 74.5) * qbook(x, t)
d2m(x, t) = (x - 74.5) ^ 2 * Exp(-etabook(x, t) - etaref(x, t)) * qbook(x, t) ^ 2
Next t
Next x

For t = 2000 To 2013
dl = 0
d2l = 0
For x = 60 To 89
dl = dl + (dbook(x, t) / mbook(x, t) - ebook(x, t)) * dm(x, t)
d2l = d2l + dbook(x, t) * (mbook(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mbook(x, t) ^ 2 - ebook(x, t) * d2m(x, t)
Next x
k2book(t) = k2book(t) - dl / d2l
Next t
```

```vba
' calculate book log likelihood

lnLprebook = lnLbook
lnLbook = 0
For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
etabook(x, t) = k1book(t) + (x - 74.5) * k2book(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
lnLbook = lnLbook + dbook(x, t) * WorksheetFunction.Ln(ebook(x, t)) + dbook(x, t) *
WorksheetFunction.Ln(mbook(x, t)) - ebook(x, t) * mbook(x, t) - WorksheetFunction.GammaLn(dbook(x, t) + 1)
Next t
Next x


Loop

' calculate reference standardised residuals

devref = 0
For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
resref(x, t) = Sgn(dref(x, t) - eref(x, t) * mref(x, t)) * (2 * (dref(x, t) * WorksheetFunction.Ln(dref(x, t) / eref(x, t)
/ mref(x, t)) - dref(x, t) + eref(x, t) * mref(x, t))) ^ 0.5
devref = devref + resref(x, t) ^ 2
Next t
Next x
disref = devref / (34 * 30 - (34 * 3 + 63 - 3))

For x = 60 To 89
For t = 1980 To 2013
resref(x, t) = resref(x, t) / disref ^ 0.5
Next t
Next x

' calculate book standardised residuals

devbook = 0
For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
etabook(x, t) = k1book(t) + (x - 74.5) * k2book(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
resbook(x, t) = Sgn(dbook(x, t) - ebook(x, t) * mbook(x, t)) * (2 * (dbook(x, t) * WorksheetFunction.Ln(dbook(x,
t) / ebook(x, t) / mbook(x, t)) - dbook(x, t) + ebook(x, t) * mbook(x, t))) ^ 0.5
devbook = devbook + resbook(x, t) ^ 2
Next t
```

```
Next x
disbook = devbook / (14 * 30 - 14 * 2)

For x = 60 To 89
For t = 2000 To 2013
resbook(x, t) = resbook(x, t) / disbook ^ 0.5
Next t
Next x
```

CAE+Cohorts Model

Let dref(x,t), eref(x,t), mref(x,t), qref(x,t), etaref(x,t), resref(x,t), aref(x), bref(x), kref(t), and gref(c) refer to $d_{x,t}^{R}$, $e_{x,t}^{R}$, $m_{x,t}^{R}$, $q_{x,t}^{R}$, $\eta_{x,t}^{R}$, $r_{x,t}^{R}$, α_{x}^{R}, β_{x}^{R}, κ_{t}^{R}, and γ_{c}^{R} for the reference component, and let dbook(x,t), ebook(x,t), mbook(x,t), qbook(x,t), etabook(x,t), resbook(x,t), abook(x), and kbook(t) denote $d_{x,t}^{B}$, $e_{x,t}^{B}$, $m_{x,t}^{B}$, $q_{x,t}^{B}$, $\eta_{x,t}^{B}$, $r_{x,t}^{B}$, α_{x}^{B}, and κ_{t}^{B} for the book component. The Excel VBA codes for fitting the CAE+Cohorts model are stated as follows.

```
' set initial values of reference alpha, beta, kappa, gamma

For x = 60 To 89
aref(x) = 0
bref(x) = 1 / 30
For t = 1980 To 2013
aref(x) = aref(x) + WorksheetFunction.Ln(dref(x, t) / (eref(x, t) + 0.5 * dref(x, t)) / (1 - dref(x, t) / (eref(x, t) + 0.5
* dref(x, t))))
Next t
aref(x) = aref(x) / 34
Next x

For t = 1980 To 2013
kref(t) = 0
Next t

For c = 1891 To 1953
gref(c) = 0
Next c

' iteratively update values of reference alpha, beta, kappa, gamma

lnLref = -10000000
lnLpreref = -11000000
Do Until lnLref - lnLpreref < 0.00000000001

' update reference alpha

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
```

112

```
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = qref(x, t)
d2m(x, t) = Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x


For x = 60 To 89
dl = 0
d2l = 0
For t = 1980 To 2013
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
Next t
aref(x) = aref(x) - dl / d2l
Next x

' update reference kappa

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = bref(x) * qref(x, t)
d2m(x, t) = bref(x) ^ 2 * Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x


For t = 1980 To 2013
dl = 0
d2l = 0
For x = 60 To 89
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
Next x
kref(t) = kref(t) - dl / d2l
Next t

' update reference beta

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = kref(t) * qref(x, t)
d2m(x, t) = kref(t) ^ 2 * Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x
```

```
For x = 60 To 89
dl = 0
d2l = 0
For t = 1980 To 2013
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
Next t
bref(x) = bref(x) - dl / d2l
Next x

' update reference gamma

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
dm(x, t) = qref(x, t)
d2m(x, t) = Exp(-etaref(x, t)) * qref(x, t) ^ 2
Next t
Next x

For c = 1891 To 1953
dl = 0
d2l = 0
For x = 60 To 89
For t = 1980 To 2013
If c = t - x Then
dl = dl + (dref(x, t) / mref(x, t) - eref(x, t)) * dm(x, t)
d2l = d2l + dref(x, t) * (mref(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mref(x, t) ^ 2 - eref(x, t) * d2m(x, t)
End If
Next t
Next x
gref(c) = gref(c) - dl / d2l
Next c

' incorporate reference constraints

sumbref = 0
For x = 60 To 89
sumbref = sumbref + bref(x)
Next x

sumkref = 0
For t = 1980 To 2013
sumkref = sumkref + kref(t)
Next t

sumgref = 0
For c = 1891 To 1953
sumgref = sumgref + gref(c)
```

```
Next c

For x = 60 To 89
aref(x) = aref(x) + bref(x) * sumkref / 34 + sumgref / 63
bref(x) = bref(x) / sumbref
Next x

For t = 1980 To 2013
kref(t) = sumbref * (kref(t) - sumkref / 34)
Next t

For c = 1891 To 1953
gref(c) = gref(c) - sumgref / 63
Next c

sum1 = 0
sum2 = 0
For t = 1980 To 2013
sum1 = sum1 + (t - 1996.5) * kref(t)
sum2 = sum2 + (t - 1996.5) ^ 2
Next t
g = sum1 / sum2

sum1 = 0
sum2 = 0
For c = 1891 To 1953
sum1 = sum1 + (c - 1922) * gref(c)
sum2 = sum2 + (c - 1922) ^ 2
Next c
h = -30 * sum1 / sum2

For x = 60 To 89
aref(x) = aref(x) + h / 30 * (x - 74.5)
bref(x) = g / (g - h) * bref(x) - h / 30 / (g - h)
Next x

For t = 1980 To 2013
kref(t) = (g - h) / g * kref(t)
Next t

For c = 1891 To 1953
gref(c) = gref(c) + h / 30 * (c - 1922)
Next c

' calculate reference log likelihood

lnLpreref = lnLref
lnLref = 0
For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
```

```vbnet
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
lnLref = lnLref + dref(x, t) * WorksheetFunction.Ln(eref(x, t)) + dref(x, t) * WorksheetFunction.Ln(mref(x, t)) -
eref(x, t) * mref(x, t) - WorksheetFunction.GammaLn(dref(x, t) + 1)
Next t
Next x


Loop


' set initial values of book alpha, kappa


For x = 60 To 89
abook(x) = 0
For t = 2000 To 2013
abook(x) = abook(x) + WorksheetFunction.Ln(dbook(x, t) / (ebook(x, t) + 0.5 * dbook(x, t)) / (1 - dbook(x, t) /
(ebook(x, t) + 0.5 * dbook(x, t))))
Next t
abook(x) = abook(x) / 14 - aref(x)
Next x


For t = 2000 To 2013
kbook(t) = 0
Next t


' iteratively update values of book alpha, kappa


lnLbook = -10000000
lnLprebook = -11000000
Do Until lnLbook - lnLprebook < 0.00000000001


' update book alpha


For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
etabook(x, t) = abook(x) + bref(x) * kbook(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
dm(x, t) = qbook(x, t)
d2m(x, t) = Exp(-etabook(x, t) - etaref(x, t)) * qbook(x, t) ^ 2
Next t
Next x


For x = 60 To 89
dl = 0
d2l = 0
For t = 2000 To 2013
dl = dl + (dbook(x, t) / mbook(x, t) - ebook(x, t)) * dm(x, t)
d2l = d2l + dbook(x, t) * (mbook(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mbook(x, t) ^ 2 - ebook(x, t) * d2m(x, t)
Next t
abook(x) = abook(x) - dl / d2l
```

116

```
Next x

' update book kappa

For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
etabook(x, t) = abook(x) + bref(x) * kbook(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
dm(x, t) = bref(x) * qbook(x, t)
d2m(x, t) = bref(x) ^ 2 * Exp(-etabook(x, t) - etaref(x, t)) * qbook(x, t) ^ 2
Next t
Next x

For t = 2000 To 2013
dl = 0
d2l = 0
For x = 60 To 89
dl = dl + (dbook(x, t) / mbook(x, t) - ebook(x, t)) * dm(x, t)
d2l = d2l + dbook(x, t) * (mbook(x, t) * d2m(x, t) - dm(x, t) ^ 2) / mbook(x, t) ^ 2 - ebook(x, t) * d2m(x, t)
Next x
kbook(t) = kbook(t) - dl / d2l
Next t

' incorporate book constraint

sumkbook = 0
For t = 2000 To 2013
sumkbook = sumkbook + kbook(t)
Next t

For x = 60 To 89
abook(x) = abook(x) + bref(x) * sumkbook / 14
Next x

For t = 2000 To 2013
kbook(t) = kbook(t) – sumkbook / 14
Next t

' calculate book log likelihood

lnLprebook = lnLbook
lnLbook = 0
For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
etabook(x, t) = abook(x) + bref(x) * kbook(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
```

```
lnLbook = lnLbook + dbook(x, t) * WorksheetFunction.Ln(ebook(x, t)) + dbook(x, t) *
WorksheetFunction.Ln(mbook(x, t)) - ebook(x, t) * mbook(x, t) - WorksheetFunction.GammaLn(dbook(x, t) + 1)
Next t
Next x


Loop

' calculate reference standardised residuals

devref = 0
For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mref(x, t) = -WorksheetFunction.Ln(1 - qref(x, t))
resref(x, t) = Sgn(dref(x, t) - eref(x, t) * mref(x, t)) * (2 * (dref(x, t) * WorksheetFunction.Ln(dref(x, t) / eref(x, t)
/ mref(x, t)) - dref(x, t) + eref(x, t) * mref(x, t))) ^ 0.5
devref = devref + resref(x, t) ^ 2
Next t
Next x
disref = devref / (34 * 30 - (34 + 30 * 2 + 63 - 4))

For x = 60 To 89
For t = 1980 To 2013
resref(x, t) = resref(x, t) / disref ^ 0.5
Next t
Next x

' calculate book standardised residuals

devbook = 0
For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
etabook(x, t) = abook(x) + bref(x) * kbook(t)
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mbook(x, t) = -WorksheetFunction.Ln(1 - qbook(x, t))
resbook(x, t) = Sgn(dbook(x, t) - ebook(x, t) * mbook(x, t)) * (2 * (dbook(x, t) * WorksheetFunction.Ln(dbook(x,
t) / ebook(x, t) / mbook(x, t)) - dbook(x, t) + ebook(x, t) * mbook(x, t))) ^ 0.5
devbook = devbook + resbook(x, t) ^ 2
Next t
Next x
disbook = devbook / (14 * 30 - (14 + 30 - 1))

For x = 60 To 89
For t = 2000 To 2013
resbook(x, t) = resbook(x, t) / disbook ^ 0.5
Next t
Next x
```

Time Series Processes

The following Excel VBA codes are written for fitting the time series processes under the M7-M5 model, in which simply the method of moments and the ordinary least squares are used. More sophisticated methods like the conditional maximum likelihood may also be adopted. Note that some values in the intermediate calculation steps are temporarily recorded in the worksheet for convenience. The terms below are rather self-explanatory; for example, k1driftref, k1sdref, and cor12ref refer to d_1, $\sigma_{\varepsilon,1}$, and $\sigma_{\varepsilon,1,2}$ of the multivariate random walk with drift (MRWD); gconref, garref, and gsdref represent ϕ_0, ϕ_1, and σ_ω of the autoregressive integrated moving average process, ARIMA(1,1,0); and k1conbook, k1ar1book, k1ar2book, k1sdbook, and corbook denote $\varphi_{1,0}$, $\varphi_{1,1}$, $\varphi_{1,2}$, $\sigma_{\xi,1}$, and $\sigma_{\xi,1,2}$ of the vector autoregressive process of order one, VAR(1).

```
' fit MRWD to reference kappa1, kappa2, kappa3

For t = 1981 To 2013
Cells(t - 1980, 1) = k1ref(t) - k1ref(t - 1)
Cells(t - 1980, 2) = k2ref(t) - k2ref(t - 1)
Cells(t - 1980, 3) = k3ref(t) - k3ref(t - 1)
Next t

dk1ref = Range("A1:A33")
dk2ref = Range("B1:B33")
dk3ref = Range("C1:C33")

k1driftref = WorksheetFunction.Average(dk1ref)
k2driftref = WorksheetFunction.Average(dk2ref)
k3driftref = WorksheetFunction.Average(dk3ref)
k1sdref = WorksheetFunction.StDev(dk1ref)
k2sdref = WorksheetFunction.StDev(dk2ref)
k3sdref = WorksheetFunction.StDev(dk3ref)
cor12ref = WorksheetFunction.Correl(dk1ref, dk2ref)
cor13ref = WorksheetFunction.Correl(dk1ref, dk3ref)
cor23ref = WorksheetFunction.Correl(dk2ref, dk3ref)

' fit ARIMA(1,1,0) to reference gamma

For c = 1892 To 1953
Cells(c - 1891, 4) = gref(c) - gref(c - 1)
Next c

xref = Range("D1:D61")
yref = Range("D2:D62")

fitref = WorksheetFunction.LinEst(yref, xref, True, True)
gconref = fitref(1, 2)
garref = fitref(1, 1)
```

```
gsdref = fitref(3, 2)

' fit VAR(1) to book kappa1, kappa2

For t = 2000 To 2013
Cells(t - 1999, 5) = k1book(t)
Cells(t - 1999, 6) = k2book(t)
Next t

xbook = Range("E1:F13")
y1book = Range("E2:E14")
y2book = Range("F2:F14")

fit1book = WorksheetFunction.LinEst(y1book, xbook, True, True)
k1conbook = fit1book(1, 3)
k1ar1book = fit1book(1, 2)
k1ar2book = fit1book(1, 1)
k1sdbook = fit1book(3, 2)

fit2book = WorksheetFunction.LinEst(y2book, xbook, True, True)
k2conbook = fit2book(1, 3)
k2ar1book = fit2book(1, 2)
k2ar2book = fit2book(1, 1)
k2sdbook = fit2book(3, 2)

corbook = 0
For t = 2001 To 2013
corbook = corbook + (k1book(t) - k1conbook - k1ar1book * k1book(t - 1) - k1ar2book * k2book(t - 1)) *
(k2book(t) - k2conbook - k2ar1book * k1book(t - 1) - k2ar2book * k2book(t - 1))
Next t
corbook = corbook / (13 - 2 * 1 - 1) / k1sdbook / k2sdbook
```

The Excel VBA codes below are written for fitting the time series processes under the CAE+Cohorts model. The notation is again self-explanatory.

```
' fit RWD to reference kappa

For t = 1981 To 2013
Cells(t - 1980, 1) = kref(t) - kref(t - 1)
Next t

dkref = Range("A1:A33")

kdriftref = WorksheetFunction.Average(dkref)
ksdref = WorksheetFunction.StDev(dkref)

' fit ARIMA(1,1,0) to reference gamma

For c = 1892 To 1953
Cells(c - 1891, 2) = gref(c) - gref(c - 1)
Next c
```

120

```
xref = Range("B1:B61")
yref = Range("B2:B62")

fitref = WorksheetFunction.LinEst(yref, xref, True, True)
gconref = fitref(1, 2)
garref = fitref(1, 1)
gsdref = fitref(3, 2)

' fit AR(1) to book kappa

For t = 2000 To 2013
Cells(t - 1999, 3) = kbook(t)
Next t

xbook = Range("C1:C13")
ybook = Range("C2:C14")

fitbook = WorksheetFunction.LinEst(ybook, xbook, True, True)
kconbook = fitbook(1, 2)
karbook = fitbook(1, 1)
ksdbook = fitbook(3, 2)
```

Central Projections

The following Excel VBA codes are used to calculate the best (central) estimates of future mortality rates under the M7-M5 model.

```
' calculate central projection estimates

For t = 2014 To 2043
k1ref(t) = k1driftref + k1ref(t - 1)
k2ref(t) = k2driftref + k2ref(t - 1)
k3ref(t) = k3driftref + k3ref(t - 1)
Next t

For c = 1892 To 1953
gdiffref(c) = gref(c) - gref(c - 1)
Next c

For c = 1954 To 1983
gdiffref(c) = gconref + garref * gdiffref(c - 1)
gref(c) = gdiffref(c) + gref(c - 1)
Next c

For t = 2014 To 2043
k1book(t) = k1conbook + k1ar1book * k1book(t - 1) + k1ar2book * k2book(t - 1)
k2book(t) = k2conbook + k2ar1book * k1book(t - 1) + k2ar2book * k2book(t - 1)
Next t
```

```
For x = 60 To 89
For t = 2014 To 2043
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
etabook(x, t) = k1book(t) + (x - 74.5) * k2book(t)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
Next t
Next x
```

The Excel VBA codes below are used to calculate the best (central) estimates of future mortality rates under the CAE+Cohorts model.

```
' calculate central projection estimates

For t = 2014 To 2043
kref(t) = kdriftref + kref(t - 1)
Next t


For c = 1892 To 1953
gdiffref(c) = gref(c) - gref(c - 1)
Next c


For c = 1954 To 1983
gdiffref(c) = gconref + garref * gdiffref(c - 1)
gref(c) = gdiffref(c) + gref(c - 1)
Next c


For t = 2014 To 2043
kbook(t) = kconbook + karbook * kbook(t - 1)
Next t


For x = 60 To 89
For t = 2014 To 2043
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
etabook(x, t) = abook(x) + bref(x) * kbook(t)
qref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
qbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
Next t
Next x
```

Residuals Bootstrapping

The pseudo data sample in one random scenario under the M7-M5 model is generated from the Excel VBA codes below. Note that the residuals formula for the reference component is typed manually in cells(5 ,1), using the values from cells(2, 1), cells(3,1), and cells(4,1), and that the difference between cells(5,1) and cells(1,1) is typed in cells(6,1), before running the codes. The residuals formula for the book component is treated similarly in the second column. The M7-M5 model or CAE+Cohorts model can then be fitted to the resulting pseudo

122

data sample, dref(x, t) and dbook(x,t), in which the time series processes are fitted to the computed model parameters based on the pseudo data sample.

```
' rearrange standardised residuals

l = 1
For x = 60 To 89
For t = 1980 To 2013
residualsref(l) = resref(x, t)
l = l + 1
Next t
Next x


l = 1
For x = 60 To 89
For t = 2000 To 2013
residualsbook(l) = resbook(x, t)
l = l + 1
Next t
Next x

' calculate fitted values

For x = 60 To 89
For t = 1980 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 – 2247.5 / 30) * k3ref(t) + gref(t - x)
qfittedref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
mfittedref(x, t) = -WorksheetFunction.Ln(1 - qfittedref(x, t))
Next t
Next x

For x = 60 To 89
For t = 2000 To 2013
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
etabook(x, t) = k1book(t) + (x - 74.5) * k2book(t)
qfittedbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
mfittedbook(x, t) = -WorksheetFunction.Ln(1 - qfittedbook(x, t))
Next t
Next x

' resample reference residuals, bootstrap reference number of deaths

For x = 60 To 89
For t = 1980 To 2013
u = Rnd() * 34 * 30
u = WorksheetFunction.RoundUp(u, 0)
Cells(1, 1) = residualsref(u)
Cells(2, 1) = 5000
Cells(3, 1) = eref(x, t) * mfittedref(x, t)
Cells(4, 1) = disref
```

```
Range("A6").GoalSeek Goal:=0, ChangingCell:=Range("A2")
dref(x, t) = Cells(2, 1)
Next t
Next x
```

```
' resample book residuals, bootstrap book number of deaths
```

```
For x = 60 To 89
For t = 2000 To 2013
u = Rnd() * 14 * 30
u = WorksheetFunction.RoundUp(u, 0)
Cells(1, 2) = residualsbook(u)
Cells(2, 2) = 500
Cells(3, 2) = ebook(x, t) * mfittedbook(x, t)
Cells(4, 2) = disbook
Range("B6").GoalSeek Goal:=0, ChangingCell:=Range("B2")
dbook(x, t) = Cells(2, 2)
Next t
Next x
```

For the CAE+Cohorts model, only three lines of the codes above need to be changed, as follows.

```
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
etabook(x, t) = abook(x) + bref(x) * kbook(t)
```

Future Simulations

The following Excel VBA codes are used to generate one random scenario of future mortality rates under the M7-M5 model. Note that all the parameters used here are computed from the pseudo data sample (not the original data) in the residuals bootstrapping process.

```
' calculate Cholesky matrix for reference kappa1, kappa2, kappa3
```

```
kcorref(1, 1) = 1
kcorref(1, 2) = cor12ref
kcorref(1, 3) = cor13ref
kcorref(2, 1) = cor12ref
kcorref(2, 2) = 1
kcorref(2, 3) = cor23ref
kcorref(3, 1) = cor13ref
kcorref(3, 2) = cor23ref
kcorref(3, 3) = 1
```

```
CTref = kcorref
```

```
For h = 1 To 2
For i = h + 1 To 3
ratio = CTref(i, h) / CTref(h, h)
```

```vba
For j = h To 3
If j = h Then
CTref(i, j) = 0
Else
CTref(i, j) = CTref(i, j) - ratio * CTref(h, j)
End If
Next j
Next i
Next h


For i = 1 To 3
ratio = CTref(i, i) ^ 0.5
For j = i To 3
CTref(i, j) = CTref(i, j) / ratio
Next j
Next i


Cref = WorksheetFunction.Transpose(CTref)

' calculate Cholesky matrix for book kappa1, kappa2

kcorbook(1, 1) = 1
kcorbook(1, 2) = corbook
kcorbook(2, 1) = corbook
kcorbook(2, 2) = 1

CTbook = kcorbook

For h = 1 To 1
For i = h + 1 To 2
ratio = CTbook(i, h) / CTbook(h, h)
For j = h To 2
If j = h Then
CTbook(i, j) = 0
Else
CTbook(i, j) = CTbook(i, j) - ratio * CTbook(h, j)
End If
Next j
Next i
Next h


For i = 1 To 2
ratio = CTbook(i, i) ^ 0.5
For j = i To 2
CTbook(i, j) = CTbook(i, j) / ratio
Next j
Next i


Cbook = WorksheetFunction.Transpose(CTbook)
```

```
For t = 2014 To 2043
For i = 1 To 3
Rref(i, 1) = WorksheetFunction.NormSInv(Rnd())
Next i
Zref = WorksheetFunction.MMult(Cref, Rref)
k1ref(t) = k1driftref + k1ref(t - 1) + k1sdref * Zref(1, 1)
k2ref(t) = k2driftref + k2ref(t - 1) + k2sdref * Zref(2, 1)
k3ref(t) = k3driftref + k3ref(t - 1) + k3sdref * Zref(3, 1)
Next t
```

' simulate reference gamma from ARIMA(1,1,0)

```
For c = 1954 To 1983
Zref = WorksheetFunction.NormSInv(Rnd())
gdiffref(c) = gconref + garref * gdiffref(c - 1) + gsdref * Zref
gref(c) = gdiffref(c) + gref(c - 1)
Next c
```

' simulate book kappa1, kappa2 from VAR(1)

```
For t = 2014 To 2043
For i = 1 To 2
Rbook(i, 1) = WorksheetFunction.NormSInv(Rnd())
Next i
Zbook = WorksheetFunction.MMult(Cbook, Rbook)
k1book(t) = k1conbook + k1ar1book * k1book(t - 1) + k1ar2book * k2book(t - 1) + k1sdbook * Zbook(1, 1)
k2book(t) = k2conbook + k2ar1book * k1book(t - 1) + k2ar2book * k2book(t - 1) + k2sdbook * Zbook(2, 1)
Next t
```

' simulate reference q, book q

```
For x = 60 To 89
For t = 2014 To 2043
etaref(x, t) = k1ref(t) + (x - 74.5) * k2ref(t) + ((x - 74.5) ^ 2 - 2247.5 / 30) * k3ref(t) + gref(t - x)
etabook(x, t) = k1book(t) + (x - 74.5) * k2book(t)
qsimref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
qsimbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
Next t
Next x
```

The Excel VBA codes below are used to produce one random scenario of future mortality rates under the CAE+Cohorts model. Note that all the parameters used here are calculated from the pseudo data sample (not the original data) in the residuals bootstrapping process.

' simulate reference kappa from RWD

```
For t = 2014 To 2043
Zref = WorksheetFunction.NormSInv(Rnd())
```

```
kref(t) = kdriftref + kref(t - 1) + ksdref * Zref
Next t

' simulate reference gamma from ARIMA(1,1,0)

For c = 1954 To 1983
Zref = WorksheetFunction.NormSInv(Rnd())
gdiffref(c) = gconref + garref * gdiffref(c - 1) + gsdref * Zref
gref(c) = gdiffref(c) + gref(c - 1)
Next c

' simulate book kappa from AR(1)

For t = 2014 To 2043
Zbook = WorksheetFunction.NormSInv(Rnd())
kbook(t) = kconbook + karbook * kbook(t - 1) + ksdbook * Zbook
Next t

' simulate reference q, book q

For x = 60 To 89
For t = 2014 To 2043
etaref(x, t) = aref(x) + bref(x) * kref(t) + gref(t - x)
etabook(x, t) = abook(x) + bref(x) * kbook(t)
qsimref(x, t) = Exp(etaref(x, t)) / (1 + Exp(etaref(x, t)))
qsimbook(x, t) = Exp(etabook(x, t) + etaref(x, t)) / (1 + Exp(etabook(x, t) + etaref(x, t)))
Next t
Next x
```

Numerical Optimisation

Consider the standardised longevity swap in Subsection 4.1. Suppose the cell A2 contains the notional amount of the swap and the cell A1 calculates the level of longevity risk reduction based on the simulations and the weight of the swap. The following Excel VBA codes can be used to find numerically the optimal weight of the swap in order to maximise the level of longevity risk reduction, i.e. minimise the longevity risk of the aggregate position.

```
SolverReset
SolverOk SetCell:="$A$1", MaxMinVal:=1, ValueOf:=0, ByChange:="$A$2", Engine:=1 _
    , EngineDesc:="GRG Nonlinear"
SolverSolve True
```

If there are two standardised longevity swaps as in Subsection 4.2, the optimal weights of the swaps (say, in the cells A2 and A3) can be found by simply adjusting the codes above to ByChange:= "$A$2:$A$3". This procedure can readily be extended to incorporate other objectives and hedging instruments.

Appendix III   Original Call for Proposals

1. Summary

This document is an invitation to tender for Phase 2 of a research project for the Longevity Basis Risk Working Group (LBRWG).

The aim of the overall project is:

to develop a readily-applicable methodology for quantifying the basis risk arising from the use of population-based mortality indices for managing the longevity risk inherent in specific blocks of pension benefits or annuitant liabilities.

We believe this project will offer the successful party an opportunity to use statistical knowledge and/or original research to produce a solution to a real industry problem. If the project were successful and facilitated the transfer of longevity risk between market participants, the work would be ground-breaking and very high-profile. We would expect that the methodology would use the indices published by the LLMA but be applicable in any territory world-wide subject to the availability of appropriate data.

The first phase of this project was successfully completed at the end of 2014. This phase focused on research into longevity basis risk and consequent development of a methodology that can be used to measure longevity basis risk.

The results from Phase 1 were presented at a sessional event at the Institute and Faculty of Actuaries (IFoA) on 8 December 2014. The research report on Phase 1 can be downloaded at:

http://www.actuaries.org.uk/events/one-day/sessional-research-event-longevity-basis-risk-methodology

Phase 2 of the project will focus on putting the work carried out in Phase 1 into practice.

The LBRWG has received a commitment to fund Phase 2 of the project from the IFoA and the Life and Longevity Markets Association (LLMA), subject to receipt of a satisfactory proposal and to achievement of interim project targets.

We expect Phase 2 of the project to last approximately 6 months from the time the project is awarded; further details of the timeline are set down below. However, credible proposals that could be completed in a shorter time frame would be considered.

We are seeking proposals from actuarial consultancies and academic institutions. Responses to the tender should be received by **5pm (GMT) on Monday 23 November 2015**. We anticipate that shortlisted candidates will be invited to present and discuss their responses during the week beginning 30 November 2015.

2. Background to the project

The LLMA began publishing indices linked to population mortality statistics in March 2012 with the goal of facilitating the hedging of longevity risk for pension funds and annuity books. The launch of the LLMA indices was an important milestone towards a longevity market where risk management can be carried out through transactions that are linked to

standardised population-level data. Index-based hedges have considerable potential to provide effective risk and capital management for all holders of longevity risk.

In addition to the mortality indices, the LLMA has also produced a significant body of work around possible derivative transactions that could reference mortality indices and offer 'standardised' longevity risk management tools (see http://www.llma.org/publications.html).

However the building blocks described above have not proved sufficient to develop a 'liquid market' in longevity and have not led to transactions based on these standardized measures. Indeed, both are underutilised relative to more traditional longevity transactions that occur in the market. Some institutions currently use risk management tools linked to indices – the concept is proven. Even so, we believe that a major obstacle to widespread use of longevity risk management tools that reference population-based mortality indices is the difficulty in quantifying, and hence managing, **longevity basis risk.**

There are two major considerations for longevity basis risk:

- The need to understand the nature of the risk and its impact in different circumstances, and
- The need to account for the basis risk underlying the transaction in reported results.

In December 2011 the LLMA and IFoA formed the LBRWG (see Appendix A for current membership). Its remit is straightforward: to investigate how to provide a market-friendly means of analysing longevity basis risk.

Having carefully considered the matter, we concluded that the task is beyond the scope of the working group by itself. The challenge is technically complex and time-consuming. Further research, or considerable work to synthesise existing research, is required before a solution can be developed. So we require the assistance of either a consultancy firm or an academic/research institution to perform that research.

Phase 1 was undertaken by a joint team from Hymans Robertson and Cass Business School, and we are now looking for a partner for Phase 2 of the project.

3. A description of Longevity Basis Risk

Longevity basis risk is the potential mismatch between the behaviour of a longevity hedge and the portfolio of pensioners or annuitants being hedged, in cases where the hedging transaction's cash flows are determined by reference to a mortality index and not directly linked to the actual pool of lives.

There are three major sources of basis risk between the pension fund/annuity book risk to be hedged and the value of the hedging tools employed to reduce that risk. These are:

- **Demographic risk:** the difference between μ_1 and μ_2, the underlying forces of mortality for the reference portfolio and the pension fund/annuity book, respectively, due to demographic or socio-economic differences. This difference may comprise two elements: the initial (current) level of mortality and the rates of future improvement.

- **Sampling risk:** the difference in the population sizes (exposures) and varying levels of annuity amounts, because any sub-population is a random sample of the large population, so the observed mortality rates in the two populations will not be the same, except by chance.
- **Structural risk** due to the payoff structure of the hedge. We could for example use a portfolio of S-Forward derivatives and compare how the value of that portfolio behaves versus the original liabilities being hedged (see the LLMA website for a description of S-Forward hedges). The pay out of the hedge is unlikely to exactly match the liabilities being hedged.

These three sources of basis risk all contribute to a longevity hedge being a less-than-perfect match to the portfolio being hedged. We believe that demographic risk and sampling risk are most usefully analysed through stochastic projections of mortality rates. Structural risk can be analysed relatively simply after the other two, because structural risk can be quantified in a straightforward fashion once scenarios of mortality rates have been projected for the different populations under consideration. Such quantification involves calculating the value of the hedge instrument under every scenario of mortality and then looking at the expected value of the result, either in isolation or relative to the pension or annuity portfolio value using a relevant metric. Therefore defining and optimising a hedge portfolio is a separate exercise from trying to estimate the relationship between the progression of mortality behaviours between $\mu1$ and $\mu2$.

4. The proposal

Throughout the project, the goal will be an outcome that is practically applicable to analysing basis risk arising from standard information available to a regular market participant. Original academic research may be required, but only in so far as it leads towards that goal.

Our proposal is for an overall project delivered in two phases:

Phase 1 focused on the demonstration of the feasibility of a methodology for determining the relationship between $\mu1$ and $\mu2$ in the future.

Phase 2 will be the practical application of the Phase 1 work to demonstrate the use of the initial research in practice.

The scope of Phase 2 would be:

1. Consideration and definition of the most relevant **metrics** for the measurement of basis risk for practitioners, and a demonstration of how the outputs from Phase 1 can be used to assess hedge effectiveness under the chosen metric(s);

2. The application of the approach from Phase 1 to **realistic worked examples**, based on the most appropriate available data. Where a suitable history of data is available, this should include examples where the predicted hedge effectiveness is back-tested against actual outcomes over a suitable timeframe. The majority of the worked examples should be able to be reproduced by practitioners using publicly available data;

3. Demonstration of how the outputs from Phase 1 might be used to **present a robust quantification of basis risk to third parties** such as regulators;

4. Investigate **potential limitations of the choice of time series** used in the Phase 1 approach, and suggest ways in which practitioners may be able to manage these limitations (including worked examples). There is a potential risk that the level of mean reversion inherent in the choice of time series under Phase 1 would result in an understatement of the potential divergence between the mortality rates of the two populations, and would like to investigate whether there are practical ways of addressing this potential limitation.

Beijing

14F China World Office 1 · 1 Jianwai Avenue · Beijing · China 100004
Tel: +86 (10) 6535 0248

Edinburgh

Level 2 · Exchange Crescent · 7 Conference Square · Edinburgh · EH3 8RA
Tel: +44 (0) 131 240 1300 · Fax: +44 (0) 131 240 1313

Hong Kong

1803 Tower One · Lippo Centre · 89 Queensway · Hong Kong
Tel: +852 2147 9418

London (registered office)

7th Floor · Holborn Gate · 326-330 High Holborn · London · WC1V 7PP
Tel: +44 (0) 20 7632 2100 · Fax: +44 (0) 20 7632 2111

Oxford

1st Floor · Park Central · 40/41 Park End Street · Oxford · OX1 1JD
Tel: +44 (0) 1865 268 200 · Fax: +44 (0) 1865 268 211

Singapore

163 Tras Street · #07-05 Lian Huat Building · Singapore 079024
Tel: +65 6717 2955

www.actuaries.org.uk