Our New 'Old' Problem - Pricing Longevity Risk in Australia

Prepared by Patricia Berry, Lawrence Tsui and Gavin Jones

Presented to the 6th International Longevity Risk and Capital Markets Solutions Conference
Sydney
9-10 September 2010

Abstract

Longevity risk is a growing risk across the developed world as populations age. Australian post-retirement life expectancy, already some of the longest in the world, is growing rapidly (by around 2.5 months a year for males and 2 months a year for females). Associated with this growth in an older population is a social need to develop products to allow individuals to secure lifetime income, and a business need to attract capital to support this risk class.

This paper sets out factors for consideration in valuing longevity risk including variation in base mortality, and mortality improvements with estimates of the financial impact of this uncertainty applied to Australian data. It aims to highlight some key issues to actuaries who are faced with the task of setting longevity assumptions in an Australian context at the current time.

In particular we show that the key drivers of mortality experience include socio-economic class (benefit level, residence and occupation as proxies) and the self selection in annuity buying behaviour. The significant variation in current mortality rates, without allowance for future improvements, can be of the order of four or more years of life expectancy at age 65, equivalent to variation in annuity cost in excess of 10%. But current mortality rates are observable; future mortality improvements are that much more uncertain and require some modelling of mortality evolution. This paper applies three models to extrapolate Australian mortality improvement and evaluates the financial impact of these risks. It shows that the risk arising from uncertainty within a single model is of the same order as that arising from the model risk (due to choice of model). This level of uncertainty should be considered in setting capital requirements for longevity risk products and in pricing annuity and longevity swap transactions.

Finally, we discuss the implications of uncertainty in current and future mortality for the transfer of longevity risk with a focus on reinsurance solutions. In particular, we consider the implications of basis risk between annuitant / pensioner and population mortality, the need for risk transfer over extremely long durations and the diversification benefits of combining annuitant portfolios with long-term guaranteed mortality risks.

The authors would like to thank Kent Hopper for his peer review. Thanks also go to David Gott and Ian Lennox for their comments and assistance.

Substantial portions of this paper were prepared for and presented at the 5th Financial Services Forum of the Institute of Actuaries of Australia,13-14 May 2009, Sydney.

Keywords: longevity risk, annuity, mortality improvement, longevity modelling

Table of Contents

1. Introduction

In Australia the old-age dependency ratio is projected to climb from 20% in 2010 to 37% by 2050 (The 2010 Intergenerational Report, the Treasury, Australian Government) increasing the potential costs of financing old age for society as a whole. With the Baby Boomers entering retirement and post retirement savings expected to grow faster than most other superannuation segments, there has been much discussion about retirees' risks including longevity risk, i.e. the risk of outliving one's savings. We are currently seeing a groundswell of activities in post retirement product development e.g. the development of Variable Annuity products, traditional annuities back in vogue and industry funds busy expanding into the post retirement market.

To support the current market activities, actuaries may be required to price for longevity risk, which is becoming the new frontier in risk transfer. **The objectives of the paper are to provide a summary of selected experience data and modelling techniques currently available, to consider the various uncertainties inherent in longevity risk and finally to discuss the transfer or hedging of longevity risk with a focus on reinsurance solutions**. Given the limited data available in Australia, the paper draws on overseas experiences in similar markets and in lessons learnt from modelling future mortality improvement. The greatest depth of longevity experience studies and modelling techniques seem to lie in UK where private sector longevity risk is the most traded in the world.

The first section of the paper examines the population mortality experience at older ages in selected developed countries. The second section then considers the marked differences in sub-group experience within the population and the drivers of differential mortality. The third section provides a glimpse into the world of modelling for future mortality improvement and discusses some sources of uncertainty in the forecasting, fitting three selected models to the Australian population data and demonstrating the uncertainty in formulating a framework in trend forecasting. Lastly, in the final section, we examine the issues involved in the transfer or hedging of longevity risk as well as some of the determinants of longevity market pricing.

There are two basic components in longevity risk pricing and reserving. The underlying building block is represented by the current mortality assumptions. These are preferably set based on that particular portfolio's past experience, allowing for credibility. In the absence of direct experience, other similar portfolios' experiences might be drawn on. The current mortality assumptions are then overlain by assumptions in future mortality improvement. The latter tends to be based on more macro views of mortality and modelling techniques are typically employed.

Any published population life expectancy commonly refers to period life expectancy, i.e. the expected years of life remaining encapsulating the mortality rates of all ages above the reference age at that point in time and taking no account of future mortality improvement. While period life expectancy provides a good indicator of mortality experience at a point in time, it could be misleading for a particular cohort who are likely to experience mortality improvement. For example, the latest Australian Life Table 2006-2008 calculates a male life expectancy at age 65 to be 18.6 years. If future mortality improvements are assumed to be those projected by the Australian Government Actuary (25 year improvement, Australian Life Tables 2005-07), a 65 year old male in 2007 would be expected to live for another 20.5 years, the cohort life expectancy.

2. Population Experience by Country

Figure 1 shows the male period life expectancy at age 65 over the years in selected countries. The underlying central mortality rates have been taken from the Human Mortality Database which collects population mortality data from mostly industrialised countries e.g. Australian data from Australian Bureau of Statistics.

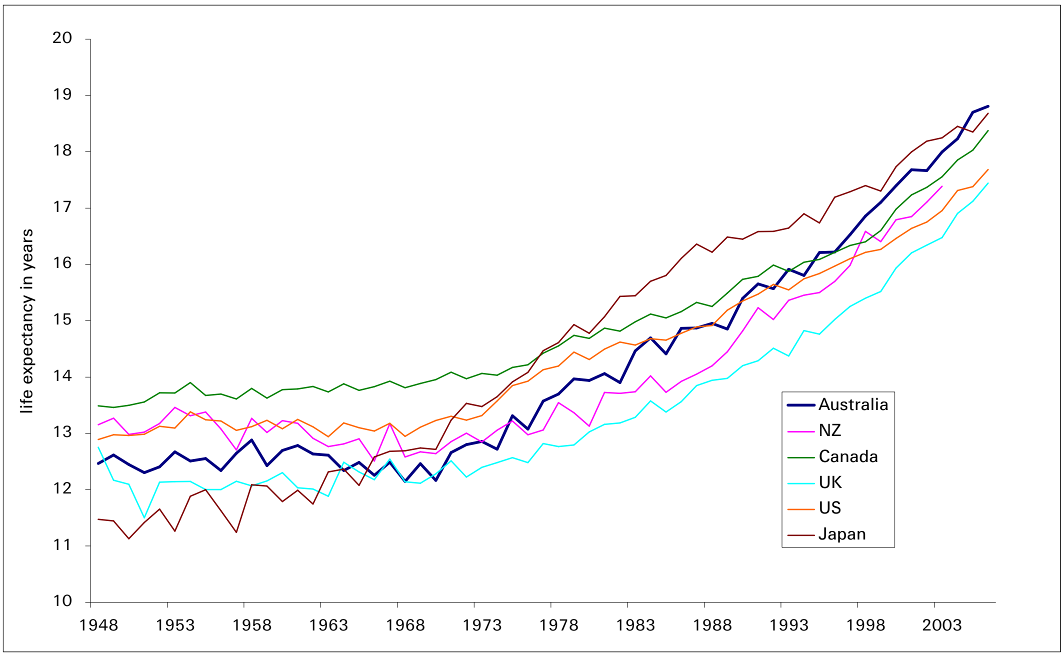

Figure 1: Male Period Life Expectancy At Age 65

There has been much common experience in developments of post World War II life expectancy. For males, the level of improvements were limited in the period up to the early 1970’s driven by the increase in cigarette consumption. Since then we have had persistent increases in life expectancy as smoking reduced and medical advances, particularly in cardio-vascular diseases, made a significant impact.

On average over the last 15 years the life expectancies at age 65 of these countries have improved by over two months per annum. For example, Australia’s life expectancy improved from 15.7 years in 1991 to 18.8 years in 2006, an average of 2.5 months per annum.

Within this broad picture of common experiences, there are many differences amongst countries with leaders becoming laggards and vice versa. As of 2006, the most recent year where data is available for most of these countries, Australia had the highest male life expectancy at age 65 (18.8 years), narrowly ahead of Japan (18.7 years). UK had the lowest life expectancy (17.4 years). In fact, Australia life expectancy improvement has been quite remarkable relative to others in the last 10 to 15 years. Prior to the 1970s, the relative performance of countries was quite different.

Figure 2 shows the equivalent female period life expectancy at age 65.

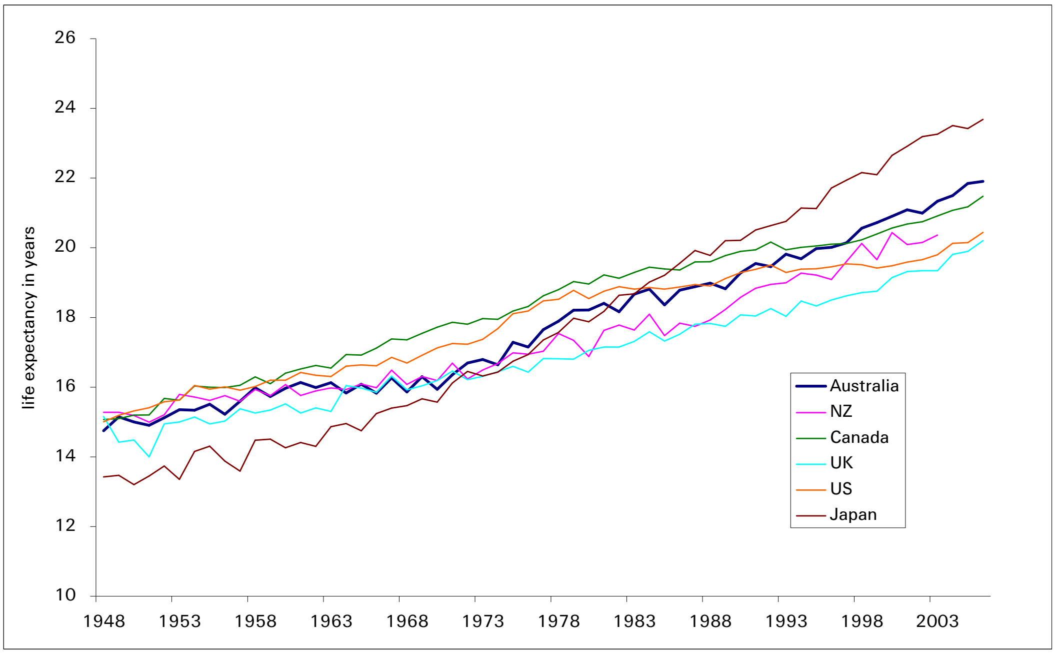

Figure 2: Female Period Life Expectancy At Age 65

With the exception of Japan, the female mortality improvements over the last 15 years have not fared as well as the male mortality improvements. For example, Australia's improvements on average have been around 1.9 months per annum (from 19.6 years in 1991 to 21.9 years in 2006). Smoking reduction and cardio-vascular medical advances having been relatively less significant for female mortality.

3. Sub-group Experience vs. Population

One would expect (voluntary) annuitant and pensioner mortality experience to be significantly better than that of the population. The higher socio-economic background of the group and, in the case of voluntary annuities, the self selection behaviour should lead to lighter mortality. The difficulty is to estimate the extent of the differentials and to identify the key drivers. This section of the paper examines some of the various pieces of data available, both overseas and in Australia, in order to gain some insights into the differences between select groups and the population.

3.1 Life Expectancy by Socio-economic Class in UK

The ONS (Office for National Statistics) Longitudinal Study shows estimates of life expectancy at birth and at age 65 by socio-economic class. The study tracks one percent sample of the England and Wales population from census to census and links the records to events such as births and deaths. Importantly, it categorises each sample member at entry into a social class based on data such as occupation relating to the member, the spouse or parents, according to priority rules.

Figures 3 and 4 show the life expectancy at age 65 by social class in UK over the last 30 years, males and females respectively. For both males and females the gaps between Professional and Unskilled are over four years in 2002-2005 (18.3 years vs. 14.1 years for males, 22.0 vs. 17.7 for females), potentially translating into a difference in annuity cost of 10% or more. It is a very significant gap.

For males there is a slight trend of widening differentials in life expectancy between social classes. One explanation is that improvements in cardio-vascular diseases have played a greater role in the higher socio-economic classes than in the lower classes (Willets et al, 2004).

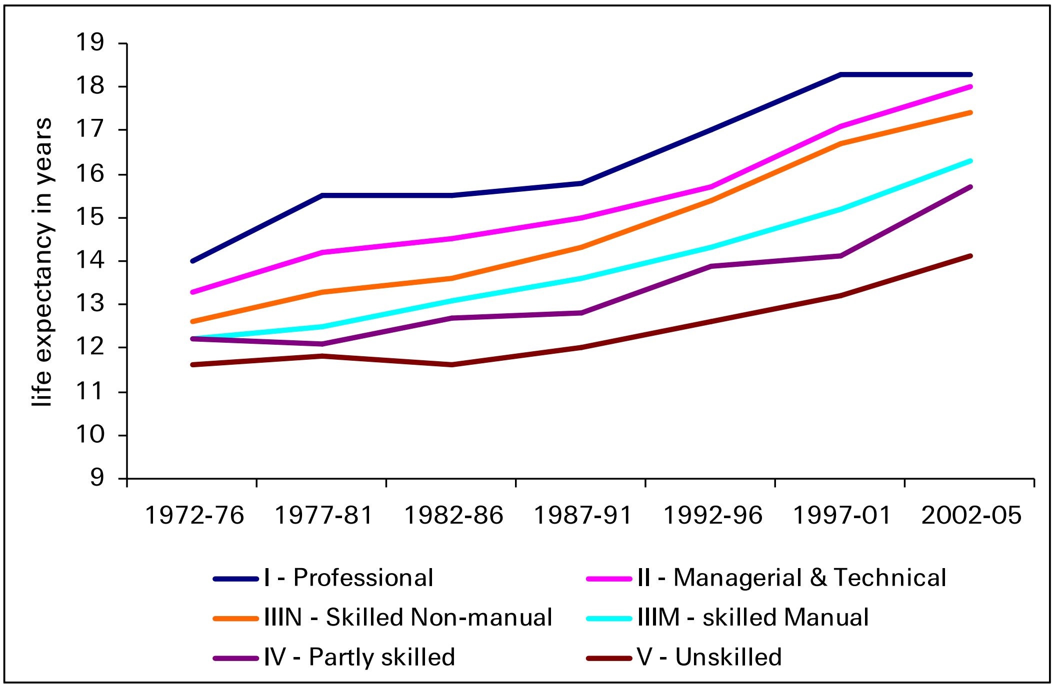

Figure 3: UK Life expectancy at 65 by social class, males

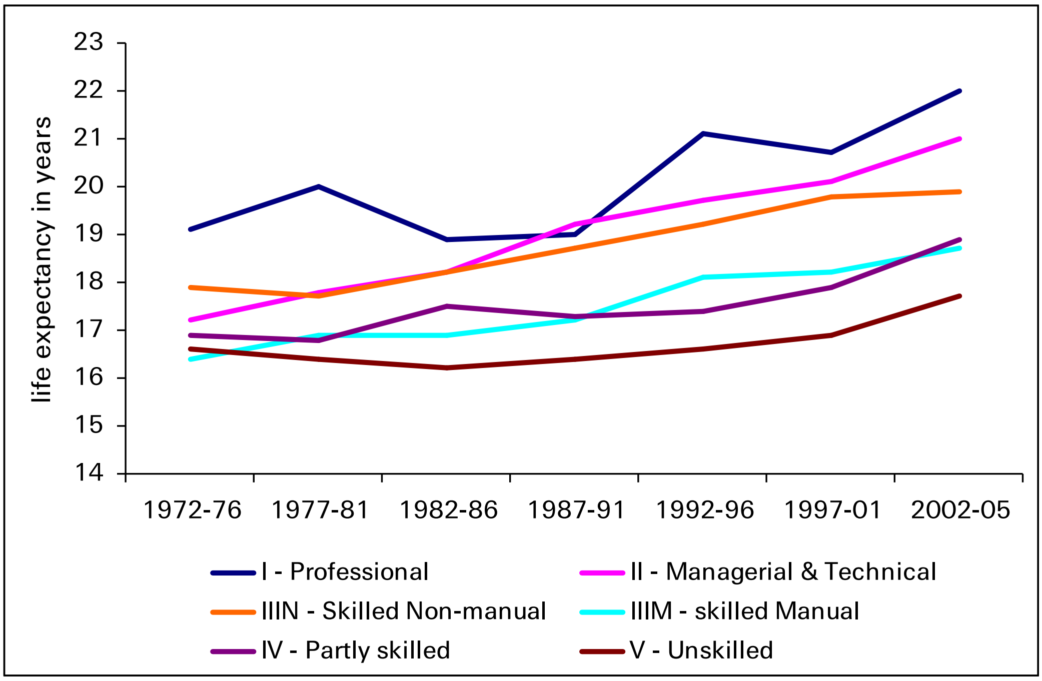

Figure 4: UK Life expectancy at 65 by social class, females

In a similar theme, the UK ONS also publishes life expectancy at age 65 by local authority. It highlights the geographical spread of life expectancy. For example, the life expectancy of males at age 65 for the period 2006-2008 varies from 13.8 years in the lowest ranking authority to 23.1 years in the highest ranking authority. A gap of 10 years might increase the value of an annuity by 30% or more. In fact, postcode is now routinely used in UK to price annuity business. For the same purchase amount the annuity benefit could vary by over 4% due to postcode differences. (Reference e.g. http://news.bbc.co.uk/1/hi/business/8441314.stm)

3.2 UK Pensioner Experience

Since the publication of a pilot study results in 2001, the Continuous Mortality Investigation (CMI) has been collecting data in respect of current pensioners of large Self Administered Pension Schemes (SAPS). Note that these are compulsory pensions in large corporate schemes, with diverse range of industries and socio-economic background.

In February 2009 the CMI produced the first set of graduated tables based on this database (CMI, 2009a). The data was for the period 2000-2006 with the graduated qx values applying around 1 September 2002. There are a series of tables separating male/female, "Ill-health"/"Normal Health" and pensioners / dependants. Interestingly, in order to highlight variability, there are "Light" and "Heavy" tables in addition to the overall tables. The "Light" tables represent subsets of the data where the pension amount exceeds a certain threshold. Similarly, the "Heavy" tables refer to subsets with pension amount below a certain threshold. Detail is shown below -

	Males			Females		
	Threshold p.a.	% of data represented		Threshold p.a.	% of data represented	
		Lives	Amounts		Lives	Amounts
Light	> £13,000	13%	43%	> £4,750	16%	49%
Heavy	< £1,500	20%	2%	< £750	25%	4%

Figures 5 and 6 highlight respectively the male and female experience of the various pension groups as a percentage of the UK population experience in the consistent time period. The graduated pensioner tables used in the comparisons are the ones representing all pensioners (normal health, ill health and unclassified, excluding dependants) and are weighted by amount.

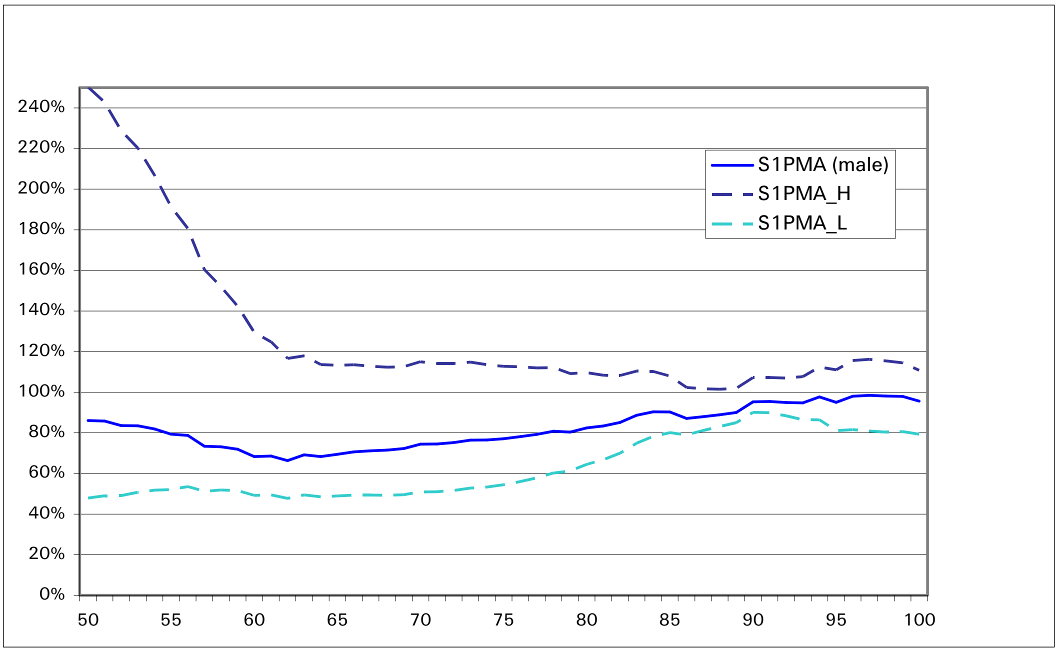

Figure 5: UK Pensioner Table Mortality as % of Population, Males

Male Pensioners

The overall male pensioner mortality as a percentage of population has the lowest values in ages 60-65, just below 70%. This coincides with the typical retirement age range and reflects the relatively good health of people having been in the workforce. Pensioners at younger ages are likely to be associated with retirement due to ill health and thus the wide range of experience. At the older ages, experience starts to converge towards the population as the 'selection' effect wears off.

The wide range of experience between Heavy and Light suggests that the level of pension amount, proxy for socio-economic class, is a key driver of mortality experience. This is not surprising. Translating the differential into financial impact, a mortality difference of 20% for instance could mean a difference of 5%-10% in annuity cost.

In a paper presented to the Institute of Actuaries in 2009, a mortality study based on a large dataset of pension schemes showed that the salary level at retirement would be an even better predictor of longevity than the pension amount (Madrigal et al, 2009).

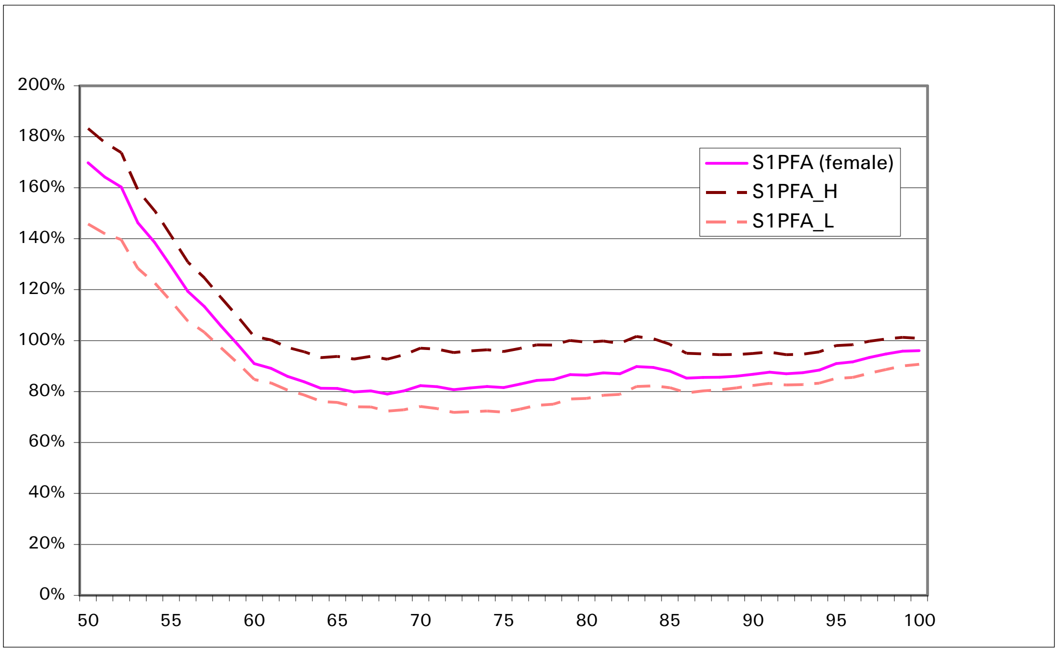

Figure 6: UK Pensioner Table Mortality as % of Population, Females

Female Pensioners

Compared with the male pensioner experience, the female ‘selection’ effect is not as pronounced as the male one but it persists into older ages. The early retirement ill health impact seems to be much greater. Overall the variability of experience is not as great as the one for males.

3.3 UK Life Office Annuitant / Pensioner Experience

The CMI Reports Number 23 is devoted to a series of graduated tables based on the 1999-2002 life office experience, known as the ‘00’ series. One such group of tables, PNMA00 and PNFA00, represents experience of Normal Life Office Pensioners by amount. Their experience relative to population experience is illustrated in Figure 7.

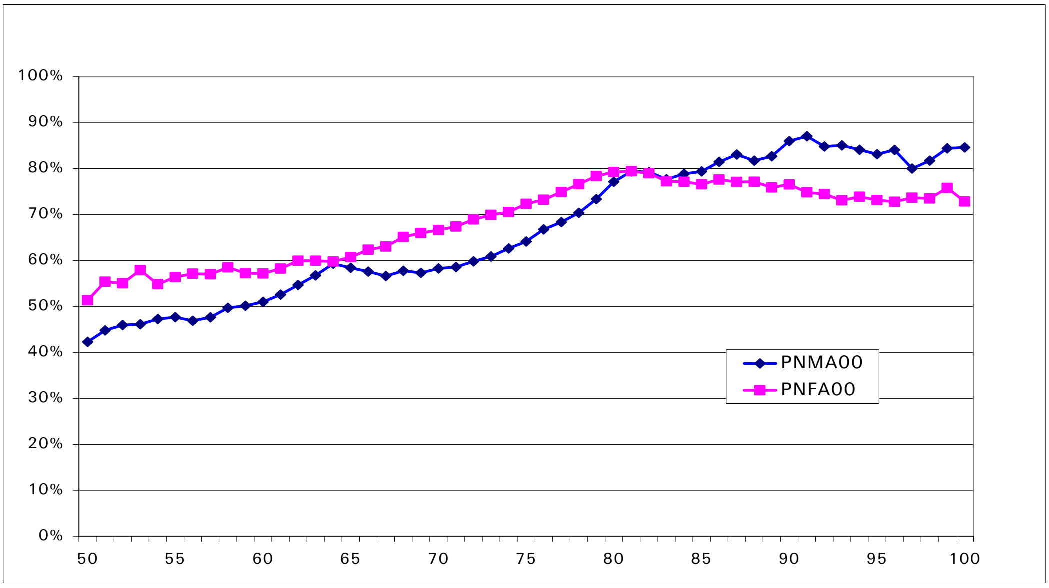

Figure 7: UK Annuitant Mortality as % of Population

Compared with the Self Administered Pension Schemes (SAPS) experience, the mortality as a percentage of population is lighter for both males and females. This could be explained by the more select nature of the group. It highlights the difficulties in pinpointing the experience of a particular group as there could be subtle differences in socio-economic background and purchasing behaviour. The exclusion of early retirement pensioners has removed the heavy experience at younger ages as observed in SAPS.

3.4 US Annuitant Experience

The Society of Actuaries has conducted an experience study of individual payout annuities for period 2000 – 2004. The study includes immediate annuities, annuitisations, and life settlement options of life insurance and annuity death claims. Pivot tables accompany the report enabling users to review experience by many variables, including age, gender, contract year, annuity type, tax class and income band.

Figures 8 and 9 show the annuitant experience as a percentage of US population mortality (US All 2002), split by refund and non refund. The refund type includes contract features such as period certain, cash refund and instalment refund.

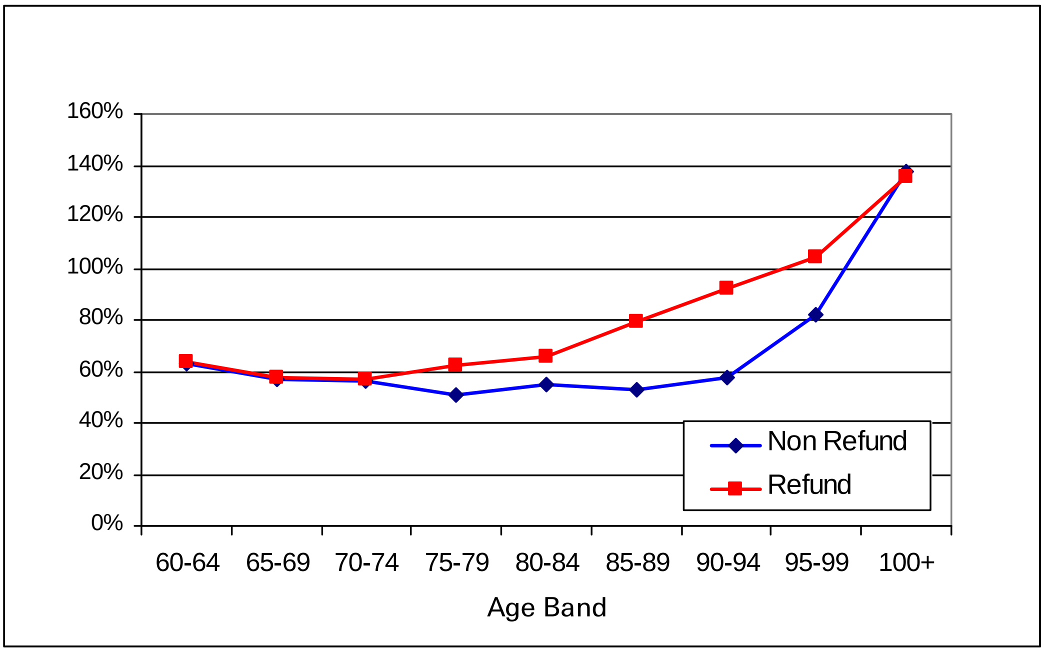

Figure 8: US Male Annuitant Mortality as % of Population

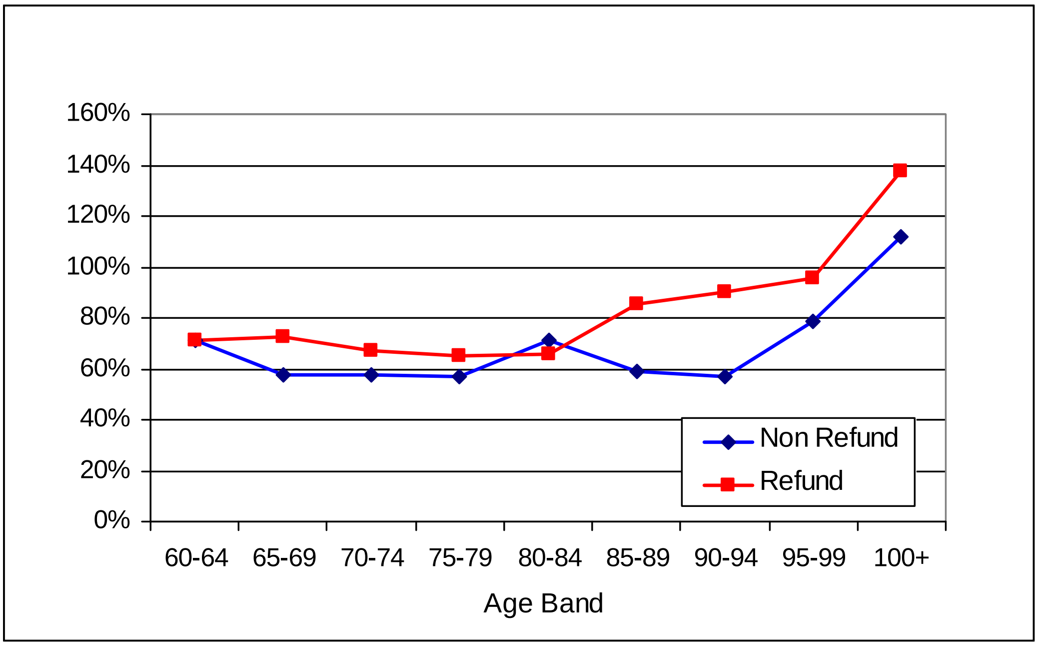

Figure 9: US Female Annuitant Mortality as % of Population

The annuitant experience is substantially lower than that of the population. The non refund experience is generally lower than the refund, highlighting the self selection of annuitants in a non-compulsory market.

Another interesting point from the SOA study is that the non refund annuities seem to have experienced greater improvement in mortality than the refund annuities in the 20 years leading to 2000-2004. When the results were compared with the ones of the 1976-1986 experience study for similar product categories, the A/E ratios were observed to have decreased significantly, particularly for the non refund annuities.

Figures 10 and 11 show the experience by annuity income band, males and females respectively. As expected, the variations reflect the significance of socio-economic background to annuitant mortality experience.

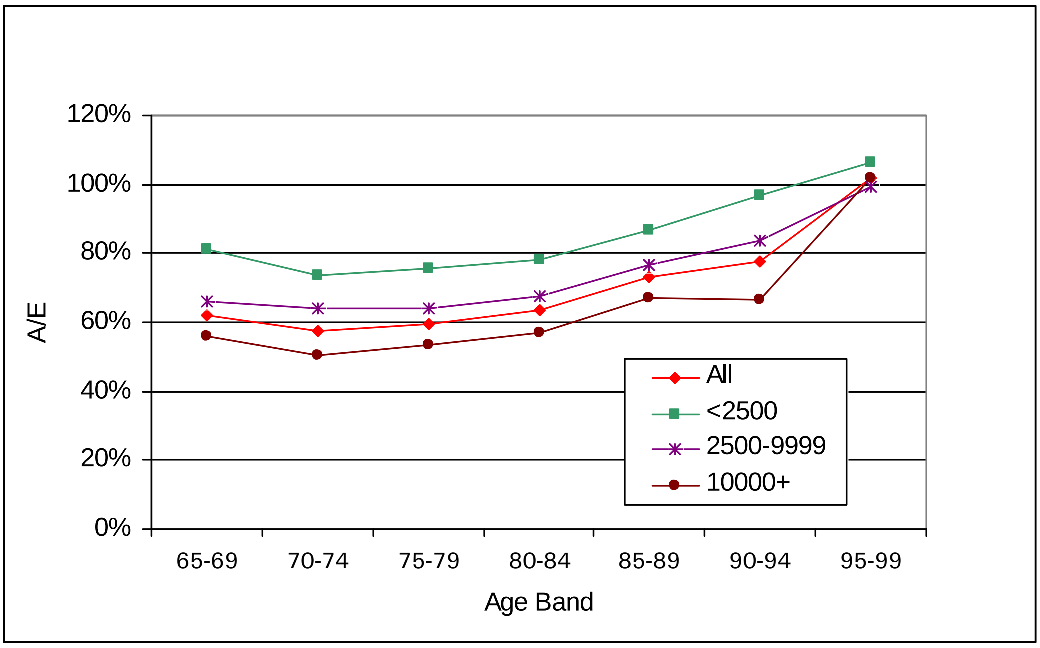

Figure 10: US Male Annuitant Mortality as % of Population, by Income Group

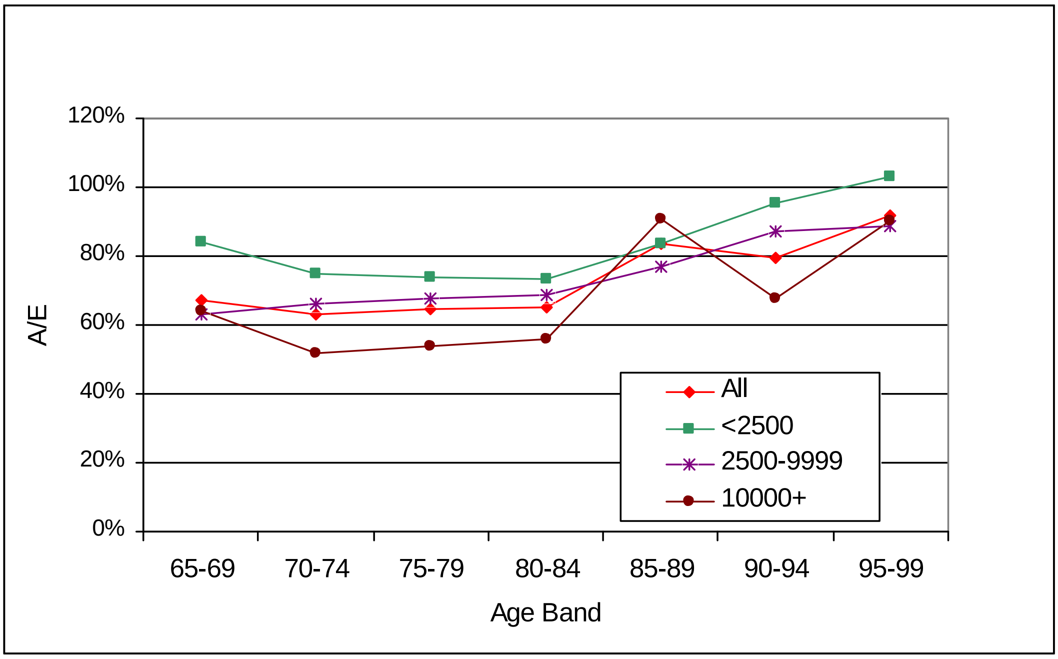

Figure 11: US Female Annuitant Mortality as % of Population, by Income Group

3.5 Australian Public Sector Scheme Pensioner Experience

The pensioner mortality of the major Australian public sector superannuation schemes has been the subject of two reports, with the last one covering period 2005-2007 (Stevenson and Wilson, 2008). It covered the large Commonwealth and state schemes, splitting by retirees, spouse and invalidity.

Figures A and B in the report illustrate the pensioner mortality experience relative to the projected population mortality. For both males and females up to the 85-89 age range, the relative experience looks remarkably similar to the relative experience of the UK Life Office Annuitants (PNMA00 and PNFA00 in section 3.3). Both cases exhibit similar 'shapes' when compared with their respective population mortality.

The similarity ceases from ages 85-89 onwards. Caution should be taken when drawing conclusions at older ages due to the much lower volume of data.

3.6 Australian Immediate Annuitant Experience

While there is a reasonable volume of in force immediate annuities in Australia, unfortunately the most recent published experience was for the period 1998-1999 (the Life Risk Insurance Committee, 2004). The total immediate annuity exposed to risk in that study was 41,806 live-years and total number of deaths was 903, a much smaller sample than the ones in the UK and US studies .

The experience in those two years has been re-expressed as a percentage of the 1998 population mortality (Human Mortality Database), as illustrated in Figure 12. The experience relative to the population was broadly consistent with the results we saw earlier in the UK and US studies.

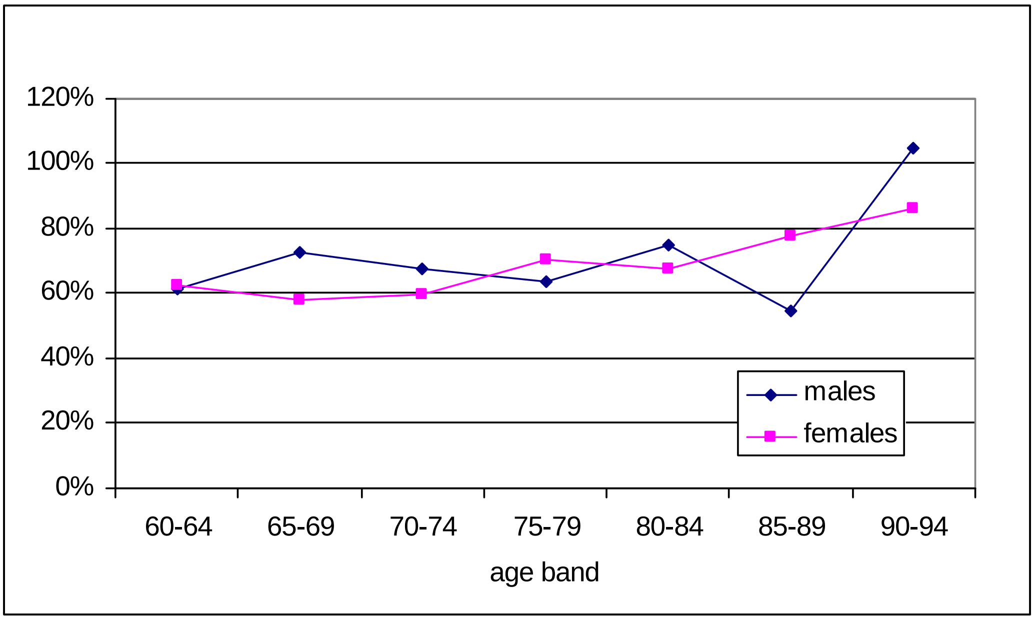

Figure 12: Australian Annuitants 1998-99 as % of Population

3.7 Summary of Observations

The key observations in this section can be summarised below –

- Socio-economic class is a strong predictor of longevity, proxies for this include postcode and benefit amount
- Other factors include health status, employment status and annuity buying behaviour, such as compulsory annuitisation and the choice between refund and non refund annuities
- There has been evidence of widening mortality gaps between the more healthy and the less healthy groups
- Assessing the impact of these drivers is complex and the effect is not uniform over time.

The financial impact of possible variation in base mortality may be greater than the expected profit margin on an annuity transaction. There is little evidential support for use of a single mortality table to price a variety of longevity business and in larger annuity markets pricing is likely to be segmented.

4. Mortality Modelling and Forecasting

In the previous sections, we have examined the uncertainty associated with setting *current* levels of mortality for annuitants. A further and greater challenge in the quantification of longevity risk is the uncertainty in *future* mortality rates. Can a history of improvements in mortality over time be translated into forecasts about future improvements in mortality?

In this section, we examine some of the approaches used in the forecasting of mortality improvements. We first look at some of the models that have been used in practice, discussing the relative strengths and weaknesses of different approaches in the context of some desirable features that a "good" mortality forecasting model should have, then we consider some sources of uncertainty in the forecasting of future mortality rates, some of which are introduced by the models themselves. Looking specifically at the Australian context, we will fit some of these models to Australian population data and demonstrate that seemingly valid mortality models can provide very different forecasts about future mortality.

We deliberately focus our attention on the practical aspects of mortality forecasting, specifically in the context of actuarial projections of future annuity and pension cash flows, concentrating on the modelling of mortality at post-retirement ages and in quantifying the financial uncertainty associated with longevity risk.

4.1 Mortality Forecasting Methods

Various authors (CMI (2004a), Coughlan, et al (2007), Booth & Tickle (2008)) have considered the various broad approaches to mortality forecasting. These can be characterised as:

- *Extrapolation* methods look primarily at historical levels of mortality and mortality improvement and extrapolating these trends into the future.

  This includes approaches that rely on simple percentage improvement factors being applied to a current mortality table, such as the mortality improvement factors currently prescribed in the Australian Solvency Standard:

  $q_{x.t} = q_{x.o} * RF(x.t)$
  where: $q_{x.t}$ is the "improved" q_x for age x at time t years after 1996
  $RF(x.t) = 0.975^t$ for $x \leq 60$
  $= (0.975 + 0.0005 * (x - 60))^t$ (max of 1) for $x > 60$

  Research on mortality forecasting has focused on extrapolation methods in recent years using regression, time-series and other statistical methods to fit a model of mortality rates (as a function of parameters that might include attained age, calendar year, birth year, etc) to historical mortality rates and then extrapolate the fitted model into the future.

  Naïve extrapolation on its own can lead to unrealistic forecasts of future experience so extrapolation methods should always be subject to tests for "reasonableness", often based on the two other broad methods described below.

- *Explanatory* or *Process-based* methods look at causative relationships between various risk factors and mortality rates, seeking to derive forecasts of future mortality by examining changes in some underlying risk factors or explanatory variables.

  An example of this approach is a model used by the World Health Organisation (described in Coughlan, *et al* (2007)) which seeks to estimate future mortality for 100 countries worldwide by relating mortality to forecast changes in average income per capita, average number of years of education, time (as a proxy for scientific/medical advancement) and tobacco consumption.

In actuarial circles, there has also been recent attempt to use explanatory methods as a validation or reasonableness check on the results of extrapolation models. In particular, Humble and Wilson (2008) look at changes in smoking pattern as a potential driver for the UK cohort effect, while the CMI (CMI 2009c and 2009d) validate their most recent mortality improvement model approach by considering projections of mortality by cause of death before re-aggregating into an overall mortality forecast. The latter can be characterised as a mixture of the extrapolative and explanatory approaches.

- *Expert Opinion* or *Expectation* is a third method that is implicitly built into many mortality forecasting approaches.

  Whereas extrapolation and explanation rely to a large extent on quantitative methods, qualitative input from the scientific and medical community shapes many of the underlying structures and parameters implicit in actuarial modelling of mortality. For example, some features of mortality forecasting models such as convergence towards long-term floors in improvement levels often reflect the consensus views of relevant experts from the perspective of "reasonableness" and simply "experience". Likewise, mortality forecasts in the absence of credible data (for example, at extreme old ages) will often rely primarily on expert opinion.

4.2 Features of Mortality Forecasting Models

At the most basic level, mortality models need to reflect the fundamental features of mortality rates:

- Age Effects
  While all models allow for mortality rates to increase with age, more structured models may assume that mortality rates follows a certain pre-determined shape, such as the Gompertz-Makeham family of exponential curves.

- Period / Calendar Year Effects
  With few exceptions (such as Russia or countries heavily affected by HIV/AIDS), mortality has generally been improving throughout the world. Extrapolative time series models (see section 4.4) fit a model to historical data and extrapolate past improvement rates into the future. Such models are usually set to extrapolate rates of mortality improvement that are constant in the long-run (as opposed to rates of improvement that accelerate or decelerate over time), although the rate of mortality improvement can fluctuate from the expected rate in a given year, and may be different across different ages or across different birth year cohorts.

- Age-Period Interactions
  As well as structuring a model for mortality by age and period, many mortality forecasting models will allow for some interaction between age and period. In short, this means that different rates of mortality improvements will be used at different ages. This is consistent with empirical observations for many developed countries (UK, US, Australia, etc) where the rate of mortality improvement tends to decrease at extremely old ages (for example, the observed rate of mortality improvement has usually been observed to be higher at age 65 than at ages 90+).

- Cohort Effects
  The existence of cohort effects, whereby lives in certain birth cohorts tend to exhibit persistently higher (or lower) rates of mortality improvement than surrounding birth cohorts has been extensively documented and discussed, especially within the UK actuarial profession (see Willets (1999) and (2004), CMI (2002a), Cairns, *et al* (2007) etc).

  While cohort effects are less prominent in some countries (such as the US), for those countries where a cohort effect features heavily (such as the UK), it is very difficult to fit a model that does not make some allowance for cohorts. Hence the basic Lee-Carter model (which models variation in mortality by age and age-period interactions only, without any allowance for birth cohorts) provides a very poor fit to UK mortality data. To take account for this, many recent mortality models have an explicit additional allowance for cohort effect, providing the flexibility for rates of mortality improvement to vary by birth year (eg Renshaw-Haberman M2, Currie APC M3, extensions to Cairns-Blake-Dowd M6-M8).

- Extensions and Further Enhancements
  Beyond the basic model features of age, period and cohort effects and their respective interactions with each other, there have been numerous enhancements to mortality models, mostly aimed at improving fit to historical data. A fairly comprehensive overview of various models that have been used for modelling mortality rates for pensioners can be found in Coughlan (2007), Cairns, *et al* (2007, 2008a, 2008b).

4.3 Mortality Improvement Patterns in Australia

It is useful to consider the features of historical mortality rates in Australia to compare the empirical data with some of the features of mortality forecasting models described in the previous section.

The charts below are “heat maps” of mortality improvements observed in the older age Australian population (for males and females respectively) over the years 1961 to 2007. Raw information on numbers of deaths and central mortality rates have been taken from the Human Mortality Database (based on data from the Australian Bureau of Statistics). The different colours in the chart reflect differences in the rate of mortality improvement, with some smoothing of the data (using simple moving average methods) in an attempt to remove the purely statistical fluctuations present in the raw data:

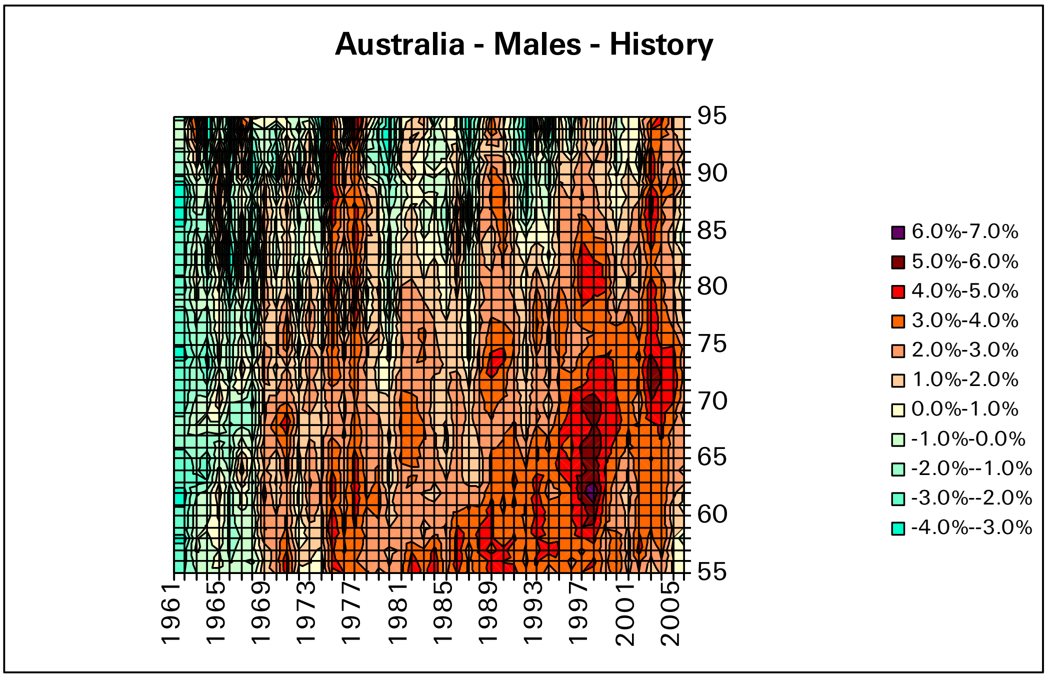

Figure 13: Australia – Males – History (1961-2007)

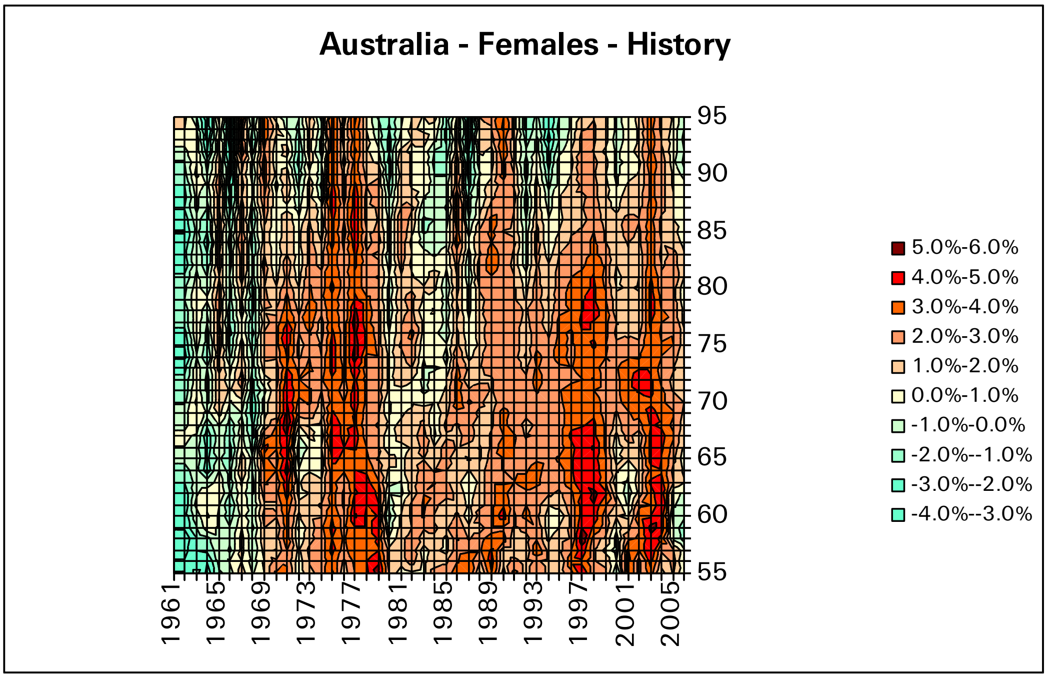

Figure 14: Australia – Females – History (1961-2007)

It is instructive to compare the Australian "heat maps" with similar charts from the UK and US, which help to identify some of the key features of present in historical rates of mortality improvement.

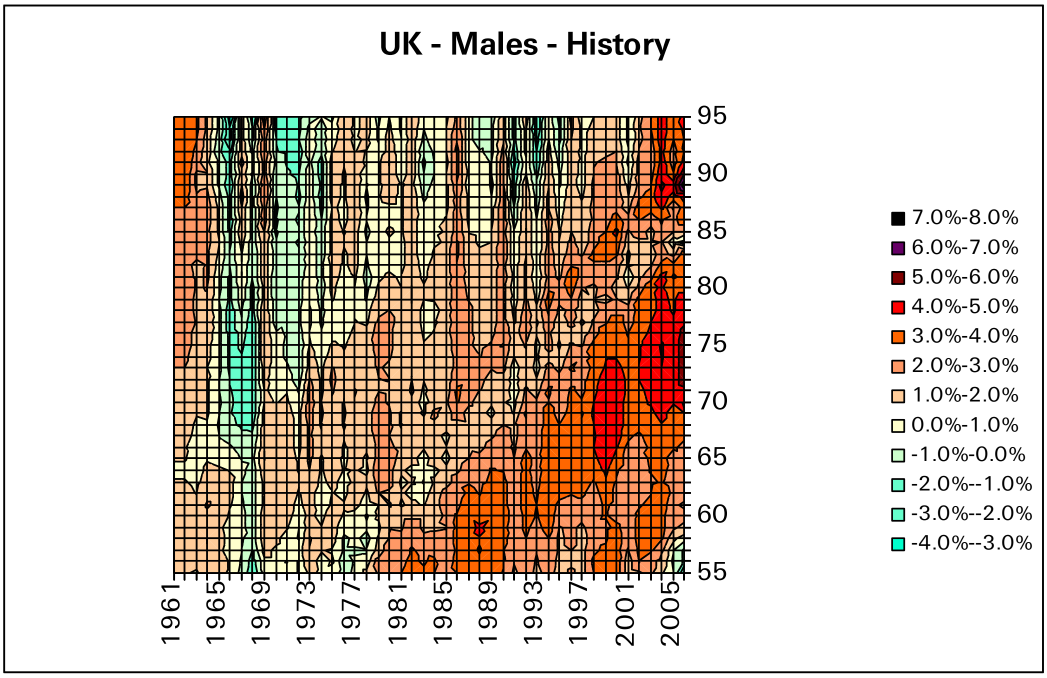

Figure 15: UK – Males – History (1961-2007)

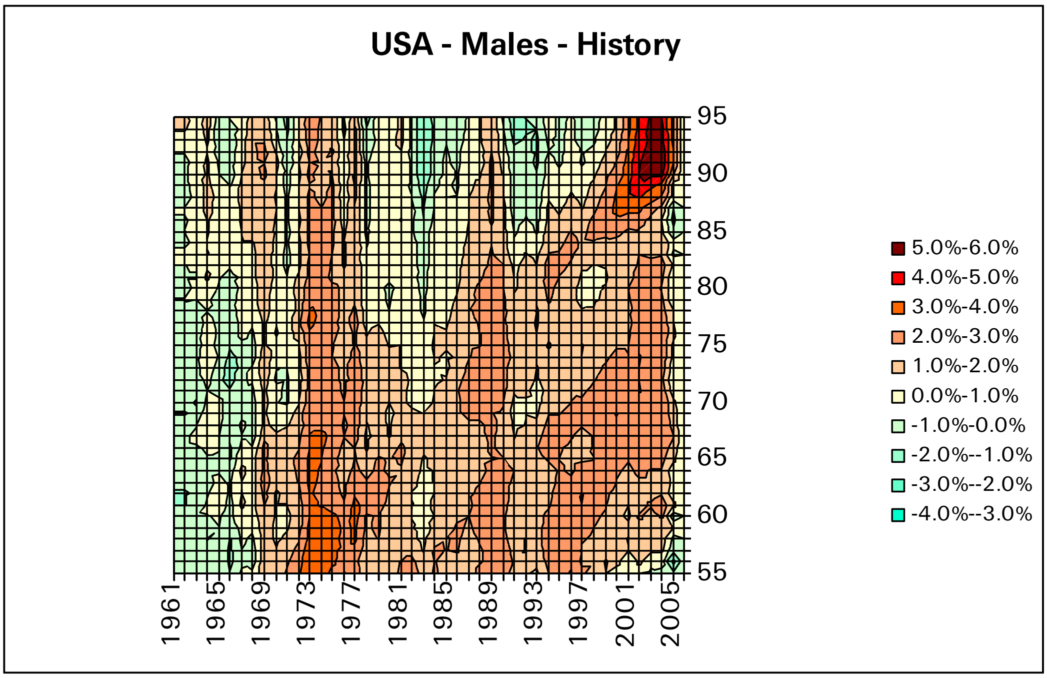

Figure 16: USA – Males – History (1961-2007)

Period Effects

At the most basic level, it is clear that for all three countries examined, mortality rates have been improving over time, with a tendency for high rates of mortality improvement in more recent years compared to the more distant past.

The Australian data appears to show certain period of time during which mortality improvement was higher or lower than the long-term average. These are the lighter or darker vertical bars that can be seen in the chart. For example, higher than average mortality improvement can be seen:

- in 1973-77 for males and slightly later in 1975-80 for females,
- in 1995-99 for both males and females; and
- in 2001-03 for both males and females.

Similar vertical bars are also visible in both the US and UK data.

Note also that low observed rates of mortality improvement in the most recent years should be treated with caution (i.e. low improvement rates for the period 2004-06). The rates of improvement at the edges of the chart are subject to greater volatility in the moving average smoothing process. It may also be noted that observed rates of mortality improvement tend to show negative autocorrelation across time (for example, see Coughlan (2007). That is, a period of low mortality improvement tends to be followed by a period of high mortality improvement.

Cohort Effects

Another feature that is very prominent in the UK Male chart, clearly visible in the Australian Male chart, but less noticeable in both the Australian Female and US Male charts is a cohort effect.

In the Australian and UK male charts, a cohort effect can be seen in the higher peak mortality improvement rates roughly following an upward sloping 45-degree diagonal line in the heat map. For example, in the UK Male chart, there is a clear cohort effect centred around lives born in 1927-1937 (aged 43-53 in 1980 and aged 70-80 in 2007). It is notable that almost exactly the same birth cohort has experience higher than average mortality improvements in the Australian male chart, centred around lives born in 1925-1935 (aged 45-55 in 1980 and aged 72-82 in 2007). There is some evidence of a similar cohort effect in the US data, centred around the same generation, but it appears to be much weaker.

4.4 Sample Model Fits for Australian Male Population Data

In the interests of illustrating the performance of some mortality projection models to Australian data, we have fitted three time series models to Australian male population mortality data. The various models have been fitted using LifeMetrics software. Each of these models is described briefly below.

(1). Lee-Carter Model (LifeMetrics M1)

This is one of the simplest of the time series models used for mortality projections, having been used extensively for modelling population mortality in a number of countries, especially the US. The Lee-Carter model has been described extensively in actuarial literature (Booth and Tickle (2008), CMI (2005 and 2007b), Coughlan, *et al* (2007), etc).

In its simplest form, the Lee-Carter model assumes that the expected rate of change at each attained age is in fixed proportion to the overall rate of improvement. For example, if overall mortality is expected to improve by 2% in a given year, and the model is fitted such that in that same year mortality at age 60 is expected to improve by 1% (50% of overall expected improvement) and mortality at age 80 is expected to improve by 3% (150% of overall expected improvement), then in a year when overall mortality is expected to improve by 3%, mortality will be expected to improve by 1.5% at age 60 and by 4.5% at age 80.

Formulaically, the function can be represented as follows:
$\log \mu_{x,t} = a_x + b_x\, p_t + \epsilon_{x,t}$

where a_x is average (over time) log-mortality at age x (age-effect),
b_x represents the fixed relationship between mortality changes at age x and overall changes in mortality (age-period interaction),
p_t represents the overall level of mortality in year t (period-effect), and
$\epsilon_{x,t}$ is a term for random statistical variation.

For the purposes of modelling mortality, one of the key points about the basic Lee-Carter model is that it does not model any cohort effects. This is clear in the following chart where we have fitted the Lee-Carter model to Australian male population mortality data from 1961 to 2007:

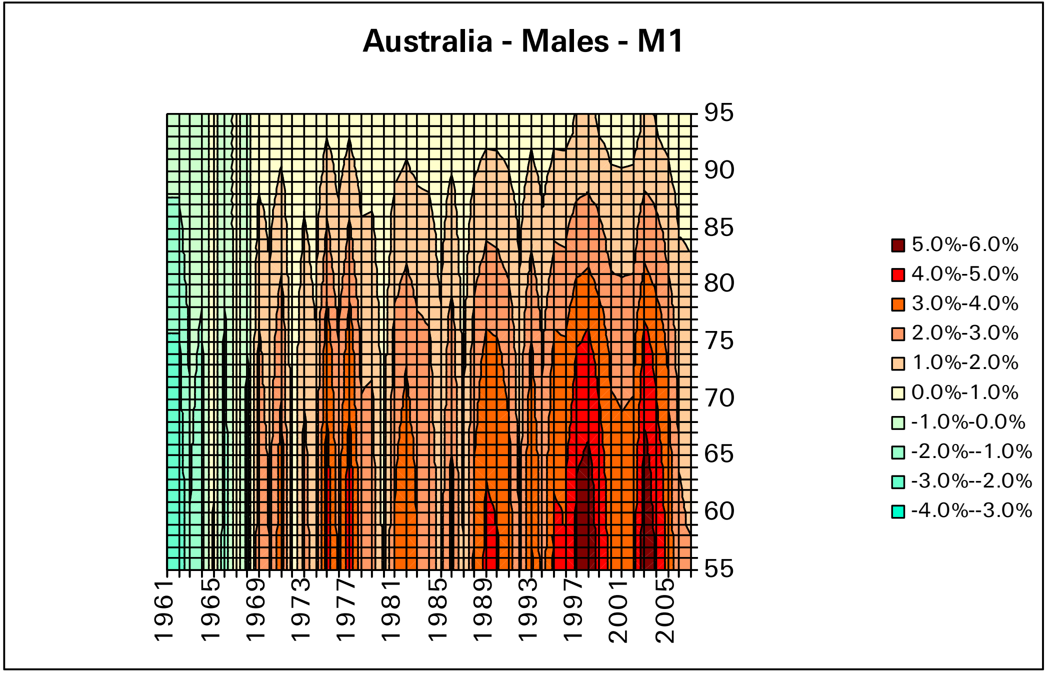

Figure 17: Australia – Males – M1 (Fitted 1961-2007)

While period effects (vertical bars) in mortality rates are fairly well represented in the above diagram, matching the high improvement periods in 1973-77, 1995-99 and 2001-03 observed in the original "heat map", there is a complete absence of any cohort effect (no upward sloping diagonals).

When extrapolating using the model to forecast rates of mortality improvement into the future, the model essentially takes average mortality improvement rates observed at each attained age over the sample period (1961 to 2007 in this case) and extrapolates this as an expected rate of mortality improvement into the future. The diagram below shows the sample period, plus extrapolated mortality rates to 2025.

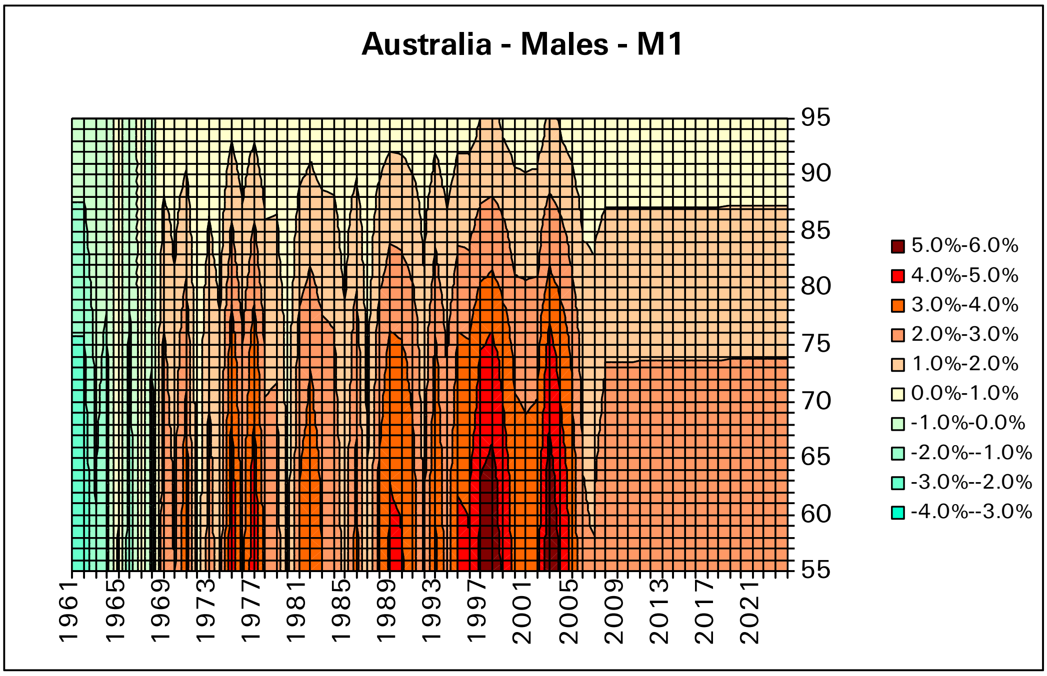

Figure 18: Australia – Males – M1 (Fitted 1961-2007, Projected 2008-2025)

(2). Currie Age-Period-Cohort (APC) Model (LifeMetrics M3)

One alternative model which is also reasonably simple and which seeks to address the cohort effect is the Currie APC model (described in Coughlan, *et al* (2007)). This model assumes that age effects, period effect and cohort effects have completely independent effects on mortality.

Formulaically, the function can be represented as follows:

$\log \mu_{x,t} = a_x + p_t + c_{t-x} + \epsilon_{x,t}$

where a_x is the element of mortality that varies by age (age-effect),
p_t represents the element of mortality that varies over time (period-effect),
c_{t-x} represents the impact on mortality of birth year t-x (cohort-effect)
$\epsilon_{x,t}$ is a term for random statistical variation.

Fitting this model to the Australian male mortality population mortality data from 1961 to 2007, produced the fitted "heat-map" shown below.

From a purely visual perspective, the Currie APC model appears to fit the Australian data better than the Lee-Carter model. Again the periods of high mortality improvement in 1973-77, 1995-99 and 2001-03 are apparent, but the fitted model also clearly shows a cohort effect for birth years 1925-1935 (aged 45-55 in 1980 and aged 72-82 in 2007).

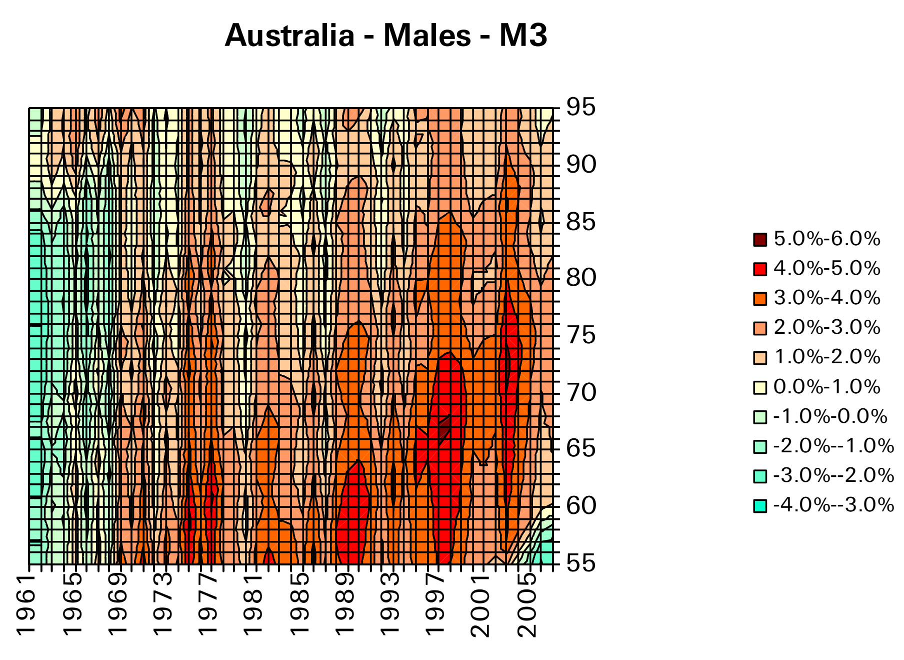

Figure 19: Australia – Males – M3 (Fitted 1961-2007)

When extrapolating this model to forecast expected future levels of mortality improvement, a clear cohort effect can be seen projected into the future as shown by the upward diagonal bars showing higher levels of expected improvement for the 1925-35 birth cohort (aged 90-100 in 2025). The Currie APC model also projects mortality deterioration for the 1950-53 birth cohort (aged 72-75 in 2025).

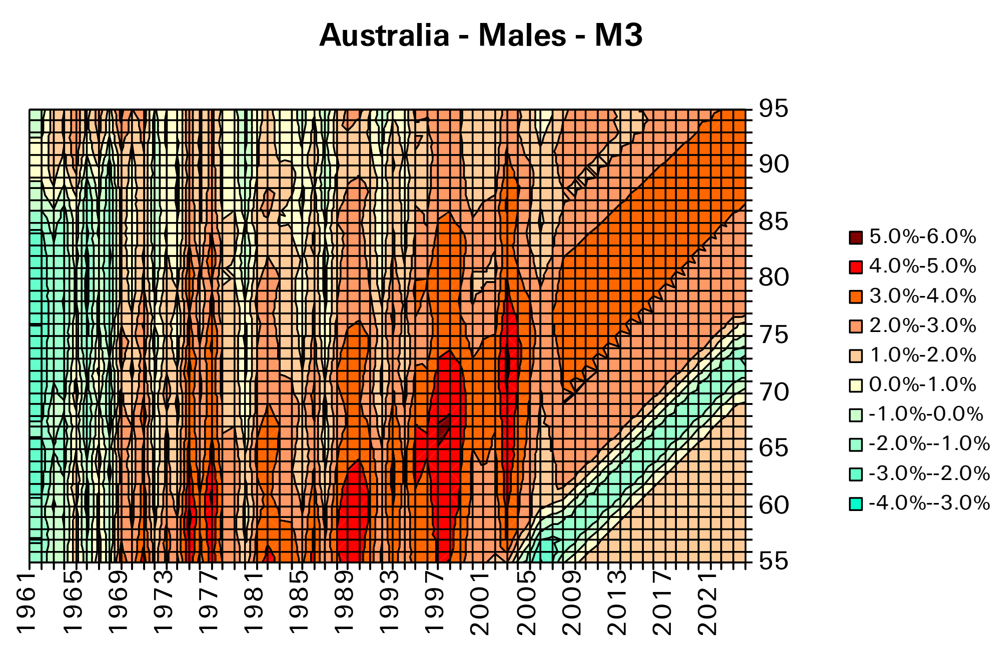

Figure 20: Australia – Males – M3 (Fitted 1961-2007, Projected 2008-2025)

Some might however question the plausibility of the cohort effect projected in this model persisting for many years into the future, especially the notion that mortality improvements for the 1925-35 birth cohort will continue at 3-4% per annum into the foreseeable future. For example, the UK actuarial profession has tended to assume some grading down of the cohort effect with time, such that the cohorts experiencing very high rates of mortality improvement in the past converge towards a lower long-term rate of improvement in the more distant future.

(3). Cairns-Blake-Dowd (CBD) Model with Cohort Term (LifeMetrics M6)

The CBD family of models (described in Coughlan, *et al* (2007) and Cairns, *et al* (2007)) use a logarithmic or logit rather than a log function for mortality rates (with a rigid pre-defined age structure for mortality rates, as opposed to the fitted age structure used for the previous models) and assume a relationship between period and age effects that depends on the difference between a particular age and the average age in the fitted sample range. While the basic CBD model does not contain a cohort term, one simple extension of the model that does include a cohort term is described below:

$$\text{logit } q_{x,t} = p_t + r_t\,(x - \overline{x}) + c_{t-x} + \epsilon_{x,t}$$

where p_t is one component of mortality that varies with time (period-effect),
r_t is a second component of mortality that varies with time (period-effect), but which has a functional relationship by age,
$(x - \overline{x})$ is the difference between a specific age x and the mean age in the sample $\overline{x}$
c_{t-x} represents the impact on mortality of lives born in year t-x (cohort-effect)
$\epsilon_{x,t}$ is a term for random statistical variation, and

$$\text{logit } q_x = \log\left(\frac{q_x}{1 - q_x}\right)$$

Again from a visual perspective, this model appears to represent the sample data reasonably well, with matching period and cohort effects. The "heat map" below is actually very similar to the one produced by the previous Currie APC model, except that there are some small differences in the intensity of different peaks and troughs in mortality improvement.

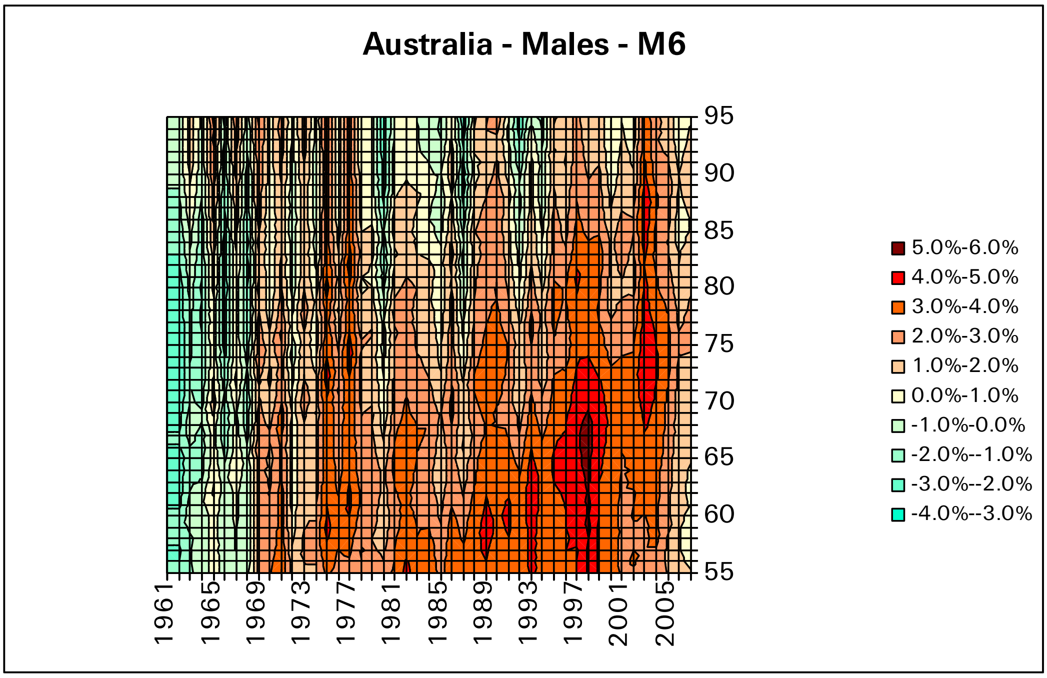

Figure 21: Australia – Males – M6 (Fitted 1961-2007)

However, when this model is used to project rates of mortality improvement into the future, the results are quite different from the APC model. Whereas the Currie APC model projected rates of mortality improvements that continued at 3-4% pa for an indefinite period, this version of the CBD model projects more moderate rates of improvement for the 1925-35 birth cohort at 2-3% which seem to grade off into lower rates of improvement in more distant years. The grading off is likely to be due to the other terms in the CBD with Cohort model which allocate more of the peak mortality improvements observed for the 1925-35 birth cohort in 1995-99 and 2001-03 to period or age-period effects, rather than to the cohort effect.

The CBD with Cohort model also projects the 1950-53 cohort with low levels of mortality improvement (0-1%) rather than the mortality deterioration projected by the Currie APC model.

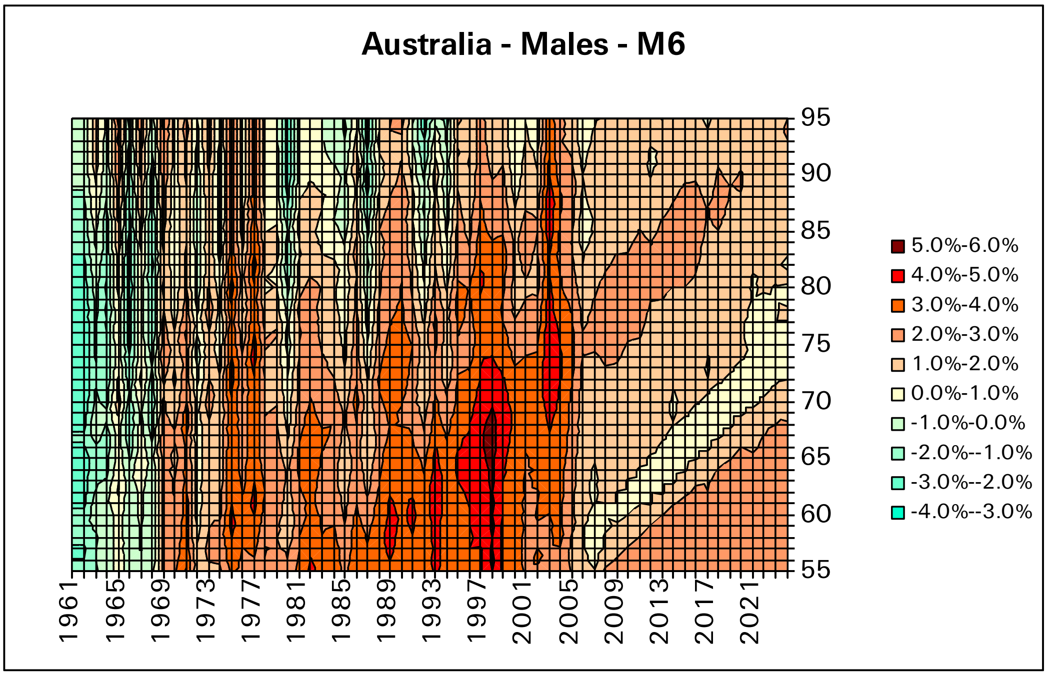

Figure 22: Australia – Males – M6 (Fitted 1961-2007, Projected 2008-2025)

4.5 Desirable Features for Mortality Projection Models

Having described a number of different mortality models in the previous section and fitted Australian male population mortality to a number of these models, it is worthwhile pausing for a moment to consider the desirable features of a mortality projection model. This issue has been considered extensively in CMI (2005, 2007b) as well as in Cairns, *et al* (2007, 2008a).

- *Ease of Implementation*
  Fitting any of the models described in the previous section requires some programming, but among mortality projection models, the first 3 described above (Lee-Carter, Currie APC, CBD with Cohort) are generally regarded as reasonably easy to implement.

- *Parsimony*
  This refers to the general characteristic of simplicity, that simple explanations are generally preferred over more complex ones. From a modelling perspective, a parsimonious model is one which minimises the number of parameters or variables requires to explain the behaviour of the underlying process. Among the models examined above, the Lee-Carter model is more parsimonious than the APC or M6.

- *Transparency of Structure*
  Each of the first 3 models described above (Lee-Carter, Currie APC, CBD with Cohort) have reasonably transparent structures, with parameters and structures that are fairly simple to explain.

- *Goodness of Fit / Cohort Effect*
  Specifically considering the Australian context, it is clear from the previous section that the Lee-Carter model provides a poor fit to the Australian data. With explicit cohort terms, both the Currie APC and the CBD with Cohort model provide a better fit to the underlying data.

  Goodness-of-fit can also be measured statistically by chi-square and other least squares measures, as well as likelihood functions. Cairns, *et al* (2007) also use the Bayes Information Criteria (which combines goodness-of-fit with parsimony in taking account of the number of parameters used in the fit – the fewer the better) as well as examining the distribution of "residuals" (ie the variation in mortality rates not explained by the fitted mortality models) to look for patterns representing systematic errors in the modelling.

- *Robust / Satisfactory Back-Testing*
  Robustness can be considered in terms of the extent to which minor changes to the underlying data can lead to significant changes in the model fit. A further element of robustness is "back-testing", where a model can be fitted to less recent historical data to see whether the model would predict the results observed in the more recent historical data. Cairns, *et al* (2008a) find that simpler models with a reasonable fit, such as the Currie APC and CBD with Cohort tend to be more robust and perform reasonably well upon back-testing of results.

- *Plausible Parameters / Plausible Projections and Variability*
  Cairns, *et al* (2008b) introduce a number of additional desirable features of mortality models which mostly can be summarised in terms of plausibility of parameters fitted to the model, as well as plausibility of projections and variability envisaged by the model. When used for forecasting, both the CMI (2007b) as well as Cairns, *et al* (2008b) found that the Lee-Carter model produces forecasts that are not sufficiently variable, especially at older ages. However, Cairns, *et al* (2008b) found that the Currie APC and CBD with Cohort models generally performed reasonably well on plausibility grounds. Interestingly, some of the more complex models tested by Cairns, *et al* (2008b) which showed a better goodness-of-fit against historical data, performed far worse in terms of plausibility of projections, providing a strong warning against introducing too many parameters and potentially over-fitting models to historical data.

- *Ability to Model Parameter Volatility / Confidence Intervals*
  All of the models described in the previous section are able to recognise volatility in the fitting of parameters by considering the potential for misestimation of the model parameters. By estimating volatility in model parameters, these models are also able to estimate confidence intervals for the model fit, effectively providing a range of possible projections for future mortality outcomes.

4.6 Quantifying Uncertainty in Future Mortality

One of the key insights from the construction of mortality projection models is some quantification of uncertainty in future mortality rates. Such insights are especially important in the context of modern solvency regimes that require solvency capital to be held at a level consistent with, for example, a 1-in-200 year event or 99.5th percentile.

Once such a model has been constructed, there are a number of sources of uncertainty in modelling of future mortality rates, discussed in CMI (2004a) and Booth and Tickle (2008):

- *Data Uncertainty*
  Especially when using data which goes back many years, the data used to fit the model can itself be a source of uncertainty in future mortality projections. For example, Cairns, *et al* (2007), have observed some discontinuities and anomalies in US male population age-specific exposures in the late 1970s, while the CMI (2009b) have reported data anomalies in UK assured life data, possibly due to relatively small sample sizes. Inadequate historical data, in terms of both volume and accuracy, is almost universally reported for extreme old ages (90+) where the precise age for exposures and deaths are often unknowable for some data elements, even in developed counties. The models that we have fitted in section 4.4 deliberately exclude ages above 90 for this reason.

- *Heterogeneity / Sampling Uncertainty*
  Closely related to data uncertainty, this is the uncertainty that arises from the possibility that mortality trends and variances may differ between different sub-segments of a heterogenous population. In the context of actuarial projections, the key uncertainty relates to potential differences between trends in the mortality of (insured) annuitants or pensioners and trends in the mortality of the overall population. While earlier work in the UK (CMI (2002a)) suggested that mortality improvements in the assured and pensioner populations might be somewhat higher than

in the general UK population, with additional evidence of higher historical mortality improvements in higher socio-economic classes than in lower socio-economic classes (CMI (2009d)), more recent evidence suggests that these differentials have narrowed in more recent periods. In the most recent UK work (CMI (2009c, 2009d and 2009e)), the decision has been taken to base mortality projections for annuitants and pensioners on UK population data, with the advantage of a much larger and more consistent dataset.

- *Model Uncertainty*
  The very act of selecting a model automatically reduces the scope for future mortality projections to vary. For example, we have seen in section 4.4 that the selection of the simple Lee-Carter model restricts mortality variation to age-period effects, with no scope to incorporate mortality variations by cohort.

  Likewise, we have seen from section 4.4, that two models (the Currie APC and the CBD with Cohort models), both quite plausible in their underlying structure and parameterisation and both providing for a reasonable fit to historical Australian male population mortality data, produce quite different median estimates for future mortality improvement. The financial significance of these model differences are further examined below.

- *Parameter Uncertainty*
  This is the risk that the parameters for the model might be misestimated. This is largely a function of volatility in the underlying data as well as the precise data periods chosen for the fitting of the model For example, fitting the same model to a different period can result in different forecast of future mortality.

  To illustrate the potential impact of differences in the fitted parameters, we have fitted the Currie APC model to a shorter time period. Whereas in section 4.4, we fitted the model to Australian male population data for 1961 to 2007, the fit below looks only at data from 1992 to 2007.

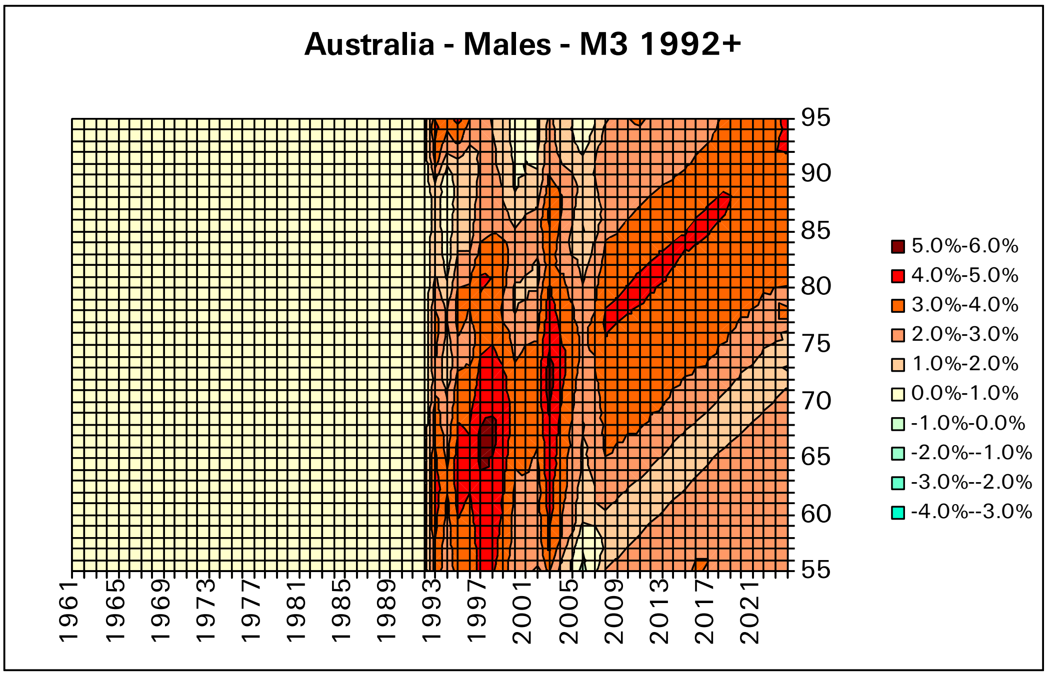

Figure 23: Australia – Males – M3 (Fitted 1992-2007, Projected 2008-2025)

Mortality improvement rates have generally been higher in the period 1992 to 2007 than over the period from 1961 to 2007, so it is not surprising that this fit of the model projects higher rates of future improvement than the previous fit. Whereas the fit in section 4.4, projected peak improvement of 3-4% pa for the 1925-1935 birth cohort, this fit projects 3-4% pa improvement for a wider birth cohort (approximately birth years 1920-1940), with a higher peak of 4-5% pa improvement for the middle of the cohort (approximately birth years 1930-

31). This fitting of the model also projects modest mortality improvements for the 1950-53 birth cohort, rather than mortality deterioration.

- *Statistical Uncertainty*

  This is generally the simplest element of mortality uncertainty to model and is a feature that is automatically included in any time series models. Each of the time series models described in section 4.4 is capable of producing sample paths of future mortality rates with varying levels of probability, within the limitations imposed by the model.

  As an example, the charts below show the 1st and 99th percentiles (out of 1,000 simulations) taken from the Currie APC model.[1] The level of volatility built into the model is based on an ARIMA process with volatility based on that observed in the historical sample data used to fit the model (we allow for period volatility only with age and cohort effects held constant).

  The 1st percentile immediately below, showing the top 1% of simulations forecasting the highest levels of future mortality improvement is shown above. With allowance for statistical volatility only, it is clear that the highest levels of improvement (up to 7-8% pa) are forecast for the immediate future (from 2007) in the projection, with mortality improvement levels moving closer to long term average levels (peak of 4-5% compared to peak of 3-4% in the median projection) in more distant periods.

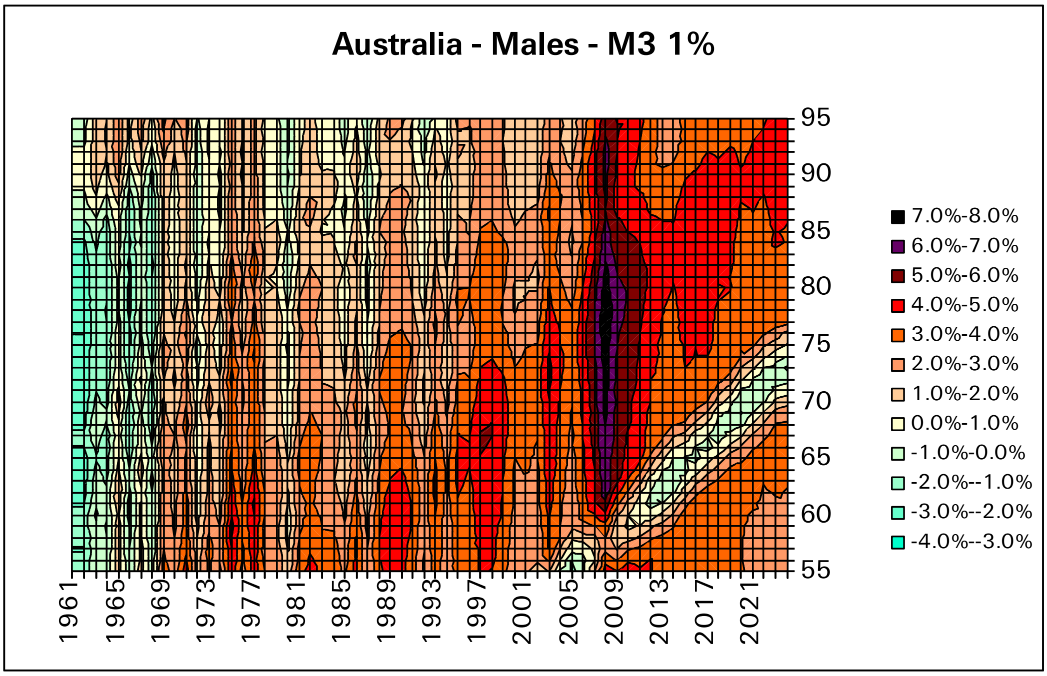

Figure 24: Australia – Males – M3 1% (Fitted 1961-2007, Projected 2008-2025)

The 99th percentile shows a similar pattern but at the opposite end of the spectrum with mortality deterioration projected for the near future, converging back to the long term average (peak of 2-3% compared to peak of 3-4% in the median projection) in more distant periods.

---

[1] We also projected 1st and 99th percentiles for the CBD with Cohort model, which gave slightly wider confidence intervals than the Currie APC model. The CBD with Cohort model tends to allocate more volatility to period effects and less to cohort effects, so that projected stochastic period volatility tends to be higher.

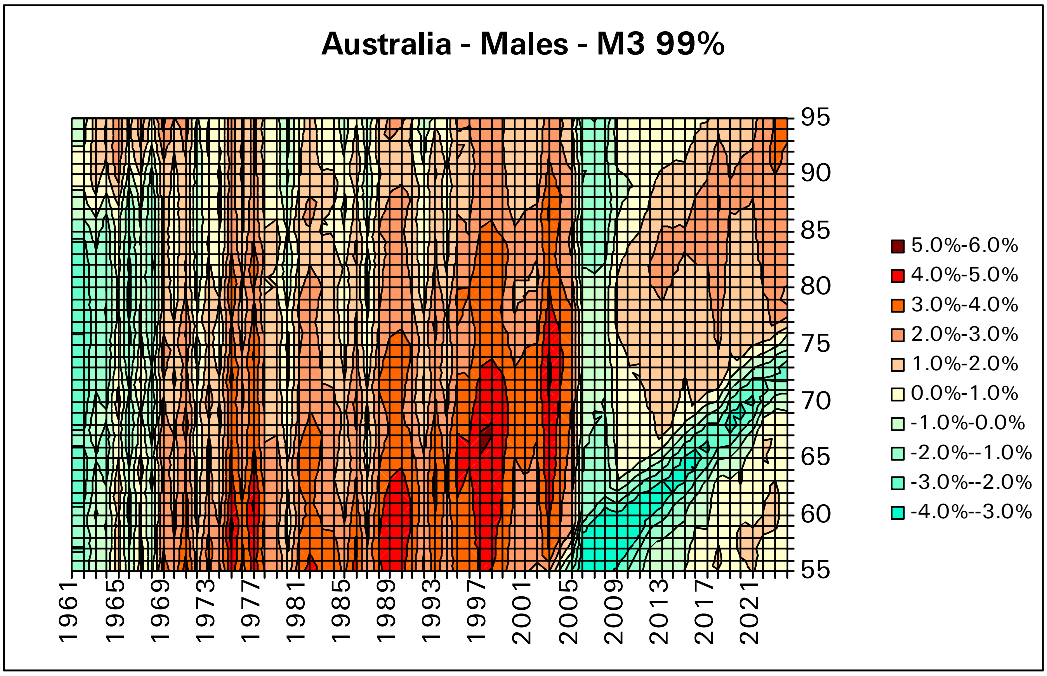

Figure 25: Australia – Males – M3 99% (Fitted 1961-2007, Projected 2008-2025)

- *Extrapolation Uncertainty*
  A final form of uncertainty with extrapolation methods is the risk of a future event that is not captured in any of the historical data and which is therefore impossible to predict with extrapolation methods. Examples include future pandemics or future advances in technology that might lead to a paradigm shift in future mortality rates. By their very nature, many such events can be characterised as "unknown unknowns", but come closest to being predictable via expert opinion, particularly consultation with the scientific and medical community concerned with the study of ageing.

A further understanding of the impact of the above types of uncertainty is illustrated in the charts and tables below.

The first two charts below show for the various mortality models, (a) projected mortality for Australian Males in 2015 as a percentage of mortality in 2007(the last year of data used in fitting the various models), and (b) average projected mortality improvement per annum for Australian Males by attained age from 2007 to 2025):

- The Lee-Carter Model (M1) projects the lowest rates of future mortality improvement. As there is no cohort effect in the Lee-Carter model, the rate of improvement is lower at older attained ages. It is also notable that the best estimate level of projected mortality for the Lee-Carter model is often outside the 1st percentile for the Currie APC model (M3).

- The Currie APC Model (M3) projects the largest cohort effect of the three models examined and the highest level of future mortality improvement. Projected mortality levels are slightly different for the same model fitted to a different sample period (M3 1992+). By 2015 (8 years after the end of the sample period 2007), there is a range of about 30% (or 3-4% improvement per annum) between the 1st and 99th percentiles based purely on statistical fluctuations in mortality improvement.

- The CBD with Cohort Model (M6) projects lower levels of improvement than the Lee-Carter and Currie APC models at younger ages, but improvement levels that are between the other two models at older ages. This reflects the presence of a cohort effect that is weaker than that fitted for the Currie APC model.

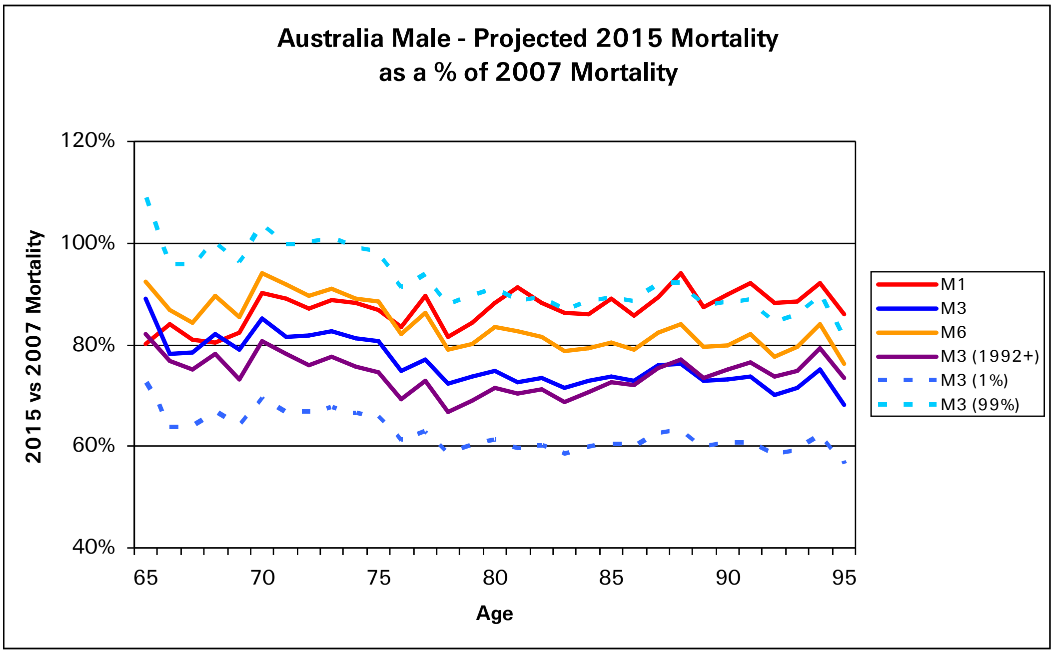

Figure 26: Australia Male – Projected 2015 Mortality as a % of 2007 Mortality

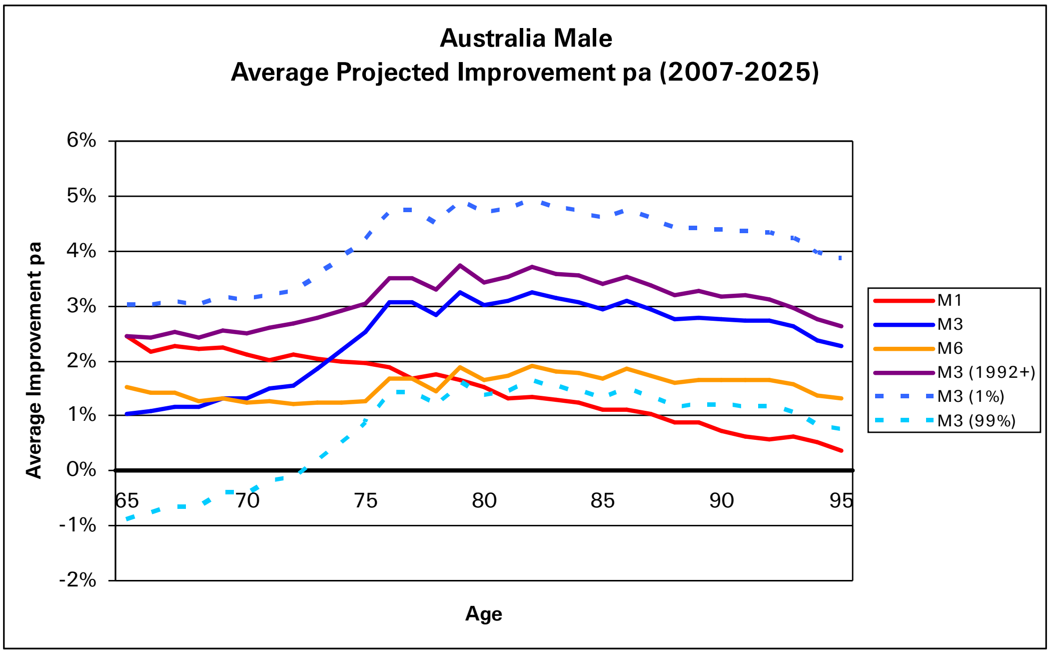

Figure 27: Australia Male – Average Projected Improvement pa (2007-2025)

The next two charts show projected mortality in future calendar years for age 65 and age 80:

- For age 65, the fluctuations over the years 2010 to 2018 mostly reflect the impact of some minor birth cohorts. For the Lee-Carter (M1) model, there is no cohort effect models, so mortality at age 65 simply follows and downward sloping line. However, for the other models, where allowance is made for cohort effects, a period of higher improvement can be observed in 2008-10 (birth years 1943-45), followed by a short period of deterioration in 2012 (birth years 1947). There is also another cohort of higher improvements in 2013-14 (birth years 1948-49), followed by deterioration from the years 2014-18 (birth years 1949-53). It is notable than both the Currie APC (M3) and CBD with Cohort (M6) models recognise the same cohort effects for these birth years.

- For age 80, the impact of cohorts is much less noticeable. Each model follows a smoother pattern of mortality improvement with the Lee Carter (M1) model projecting the lowest rates of improvement, the Currie APC (M3) model projecting the highest rates of improvement, and the CBD with Cohort (M6) model projecting something in between.

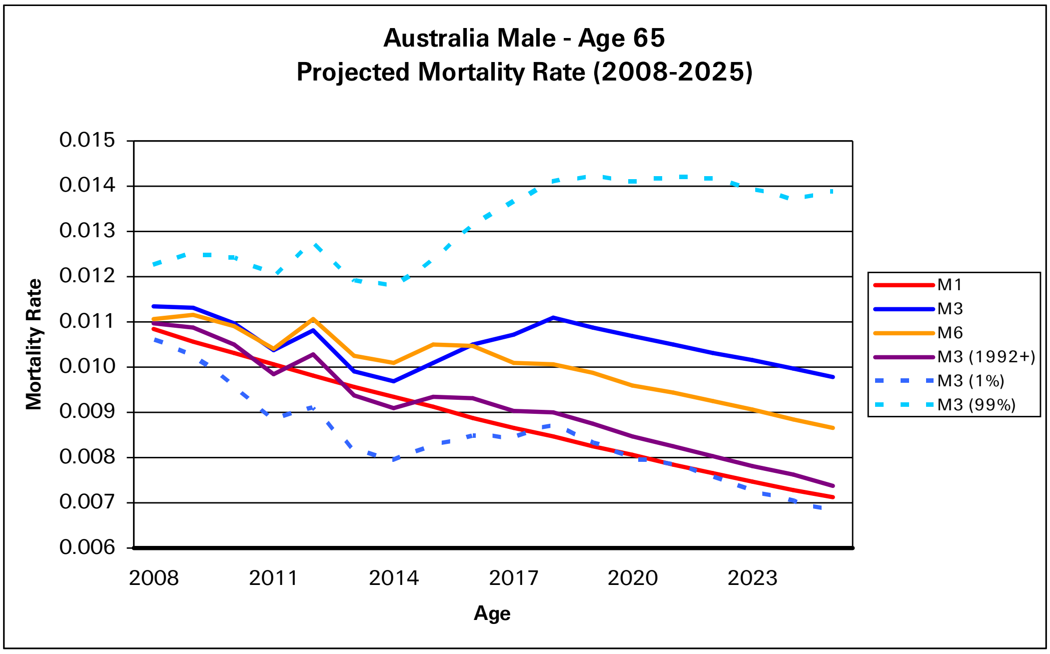

Figure 28: Australia – Male – Age 65 Projected Mortality Rate (2008-2025)

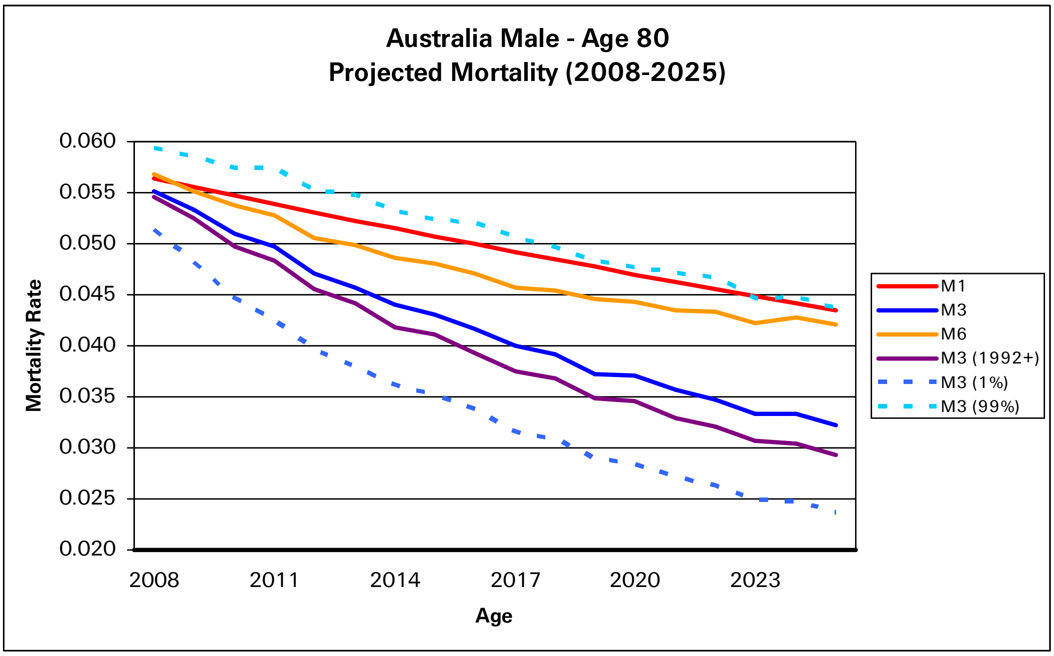

Figure 29: Australia Male – Age 80 Projected Mortality (2008-2025)

Finally, to gauge the financial impact of the various different mortality models for life insurance companies and defined benefit superannuation funds, we have calculated lifetime annuity values for a lives aged 65 and 80 in 2007 based on the various different models. In order to calculate these values, we used projected mortality rates from the various models up to age 95 with projections to the year 2040.[2]

---

[2] For ages beyond 95, we extrapolated mortality rates assuming that the rate of increase in mortality with age reduces linearly over time to 0 at age 120 and assumed $q_{120} = 1$. Beyond 2040, we assumed 1% pa mortality improvement.

Firstly, the charts below show the survival curves for these two ages for the various models:

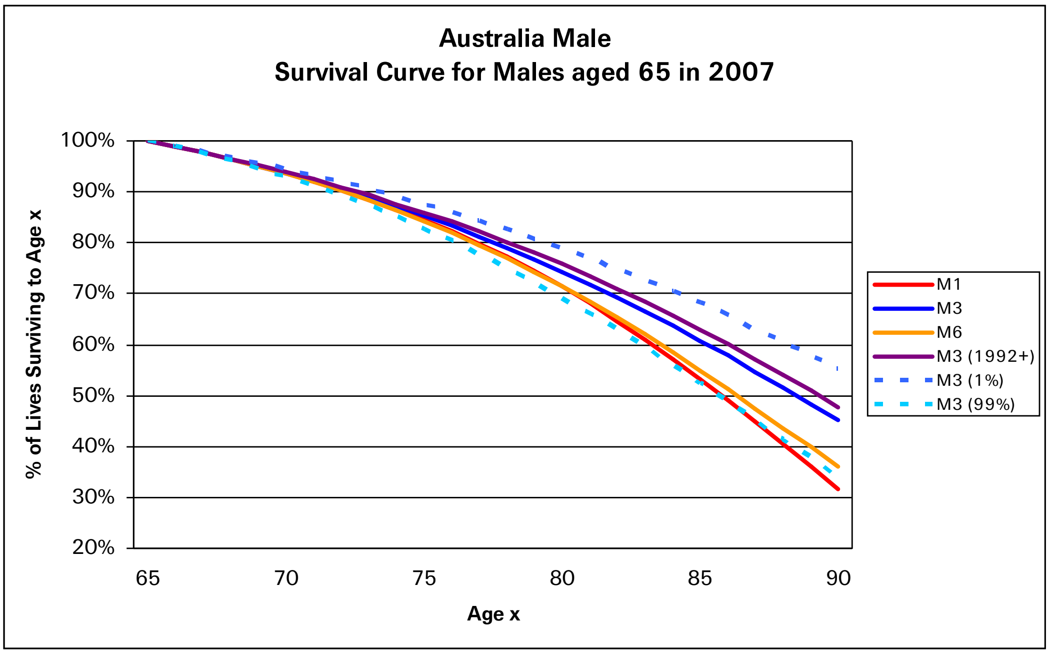

Figure 30: Australia Male – Survival Curve for Age 65 in 2007

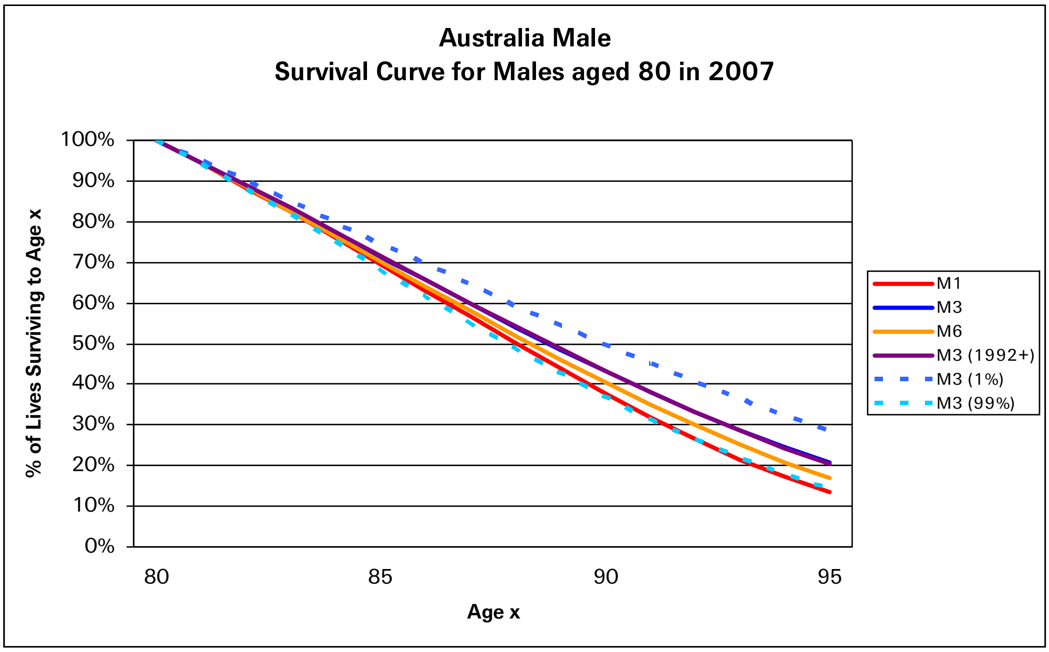

Figure 31: Australia Male – Survival Curve for Age 80 in 2007

Then, annuity values have been calculated at a few different interest rates:

- 0% representing a calculation of life expectancy
- 3% representing a long-term real discount rate (for indexed annuities)
- 6% representing a long-term nominal discount rate (for level annuities).

All annuity values are expressed as a percentage of the annuity value calculated for the M3 model to provide an idea of the range of outcomes. The resultant values are shown in the table below. As expected, differences in the assumed rate of future mortality improvement tends to have a larger impact at younger ages and at lower interest rates.

Differences between models (Lee-Carter M1, Currie APC M3, CBD with Cohort M6) accounts for differences of roughly 5-10% in annuity liabilities (and up to 15% in life expectancies). For the Currie APC (M3) model, differences between the sample period used to fit the model accounts for smaller differences, up to 2% in annuity liabilities.

Age	**Int Rate**	**M3**	**M1**	**M6**	**M3 1992+**
65	**0%**	100.0%	85.8%	89.5%	102.0%
	3%	100.0%	91.1%	93.4%	101.6%
	6%	100.0%	94.5%	95.8%	101.2%
80	**0%**	100.0%	87.8%	93.1%	98.9%
	3%	100.0%	90.6%	94.7%	99.4%
	6%	100.0%	92.6%	95.9%	99.7%

As shown in the next table, statistical volatility (looking at the 1st and 99th percentiles for the Currie APC M3 model) also accounts for differences of roughly 6-12% in annuity liabilities. We also note that this level of volatility is substantially greater than the deviations that would result from some typical "sensitivities" that might be considered for an annuity portfolio. In the table below, we considered:

- a 10% parallel shift in mortality for the Currie APC (M3) model, resulting in changes of 2-6% in annuity liabilities;
- a 20% parallel shift in mortality for the Currie APC (M3) model, resulting in changes of 5-13% in annuity liabilities;
- an additional 1% pa mortality improvement for the Currie APC (M3) model, resulting in changes of 3-5% in annuity liabilities.

Age	**Int Rate**	**M3**	**M3 1%**	**M3 99%**	**90% M3**	**80% M3**	**M3+1%pa**
65	**0%**	100.0%	114.6%	87.7%	104.9%	110.6%	109.2%
	3%	100.0%	109.0%	91.7%	103.4%	107.1%	105.2%
	6%	100.0%	105.7%	94.3%	102.4%	105.0%	103.1%
80	**0%**	100.0%	115.6%	88.0%	107.7%	116.8%	106.0%
	3%	100.0%	111.8%	90.3%	106.3%	113.5%	104.3%
	6%	100.0%	109.3%	92.0%	105.2%	111.5%	103.2%

Overall, the range of potential deviations in annuity liabilities implied by model, parameter and statistical uncertainty in relation to future mortality improvements alone (each of which imply that deviations of up to 10-12% in annuity liabilities is plausible) would be close to a 20% parallel shift in mortality rates. It should be noted that the risks under consideration relate to *trend* risks which are of a systematic nature that cannot be diversified within a portfolio of annuities, but which could be diversified across other risks within a company's overall portfolio of business.

This magnitude of deviation can be compared against the existing Australian Capital Adequacy Standard, for longevity risk, which specifies a "low margin" of 10% and a "high margin" of 20% for base mortality, plus a specified range of mortality improvement factors (which is consistent with those fitted by the Currie APC (M3) model). The above sensitivities suggest that deviations in annuity liabilities (with a combination of model, parameter and trend variability) could exceed the "high margin" deviation for a 1-in-100 year event, so that the level of capital implied by the Australian Capital Adequacy Standard could be inadequate for a monoline longevity writer. However, a diversified company should be able to take credit in its capital requirements for other risks in its portfolio, particularly mortality risk which should be negatively correlated with longevity. Nevertheless, in an Australian context, there may be limited diversification benefits between non-guaranteed yearly renewable term mortality risks at working ages and long-term guaranteed longevity risks at post-retirement ages.

4.7 Summary of Observations

The key observations in this section are summarised below –

- Various methods can be used to project future mortality rates, none of them ideal. A combination of extrapolation, explanation and expert opinion is necessary to ensure that projections are reasonable and take account of all relevant and available information.
- Historical rates of mortality improvement in Australia have exhibited significant variations by age, period and cohort, with particularly high rates of mortality improvement for the cohort of Australian males born in 1925-35. Mortality projections that ignore cohort effects may significantly underestimate future mortality improvements.
- Similarly plausible models for extrapolating mortality rates into the future can produce very different projections of future mortality and very different financial outcomes for writers of lifetime annuities. There is significant systematic variability associated with model selection and parameter fitting in addition to statistical volatility of future mortality improvement trends.
- Given the inherently high level of volatility, a single projection of future mortality improvements is unlikely to be sufficient for the proper quantification and management of longevity risk. For actuaries looking to price lifetime annuity business, an understanding of the range of potential outcomes will be especially important when assessing internal capital needs and profitability requirements.

5. Longevity Risk Transfer and Hedging

So far in this paper, we have considered the various uncertainties inherent in longevity risk, both in respect of the current base mortality applicable to an annuitant or pensioner portfolio as well as the future mortality improvements that might be experienced by such a portfolio. In this section, we consider the implications of uncertainty in current and future mortality for organisations that hold significant amounts of longevity risk and discuss some of the options available for managing such risks.

5.1 Recognition of Longevity Risk

The nature of longevity risk is such that any deviation in longevity from expected levels tends to be recognised only very slowly. As the previous sections have highlighted, trend risk is a significant element of longevity risk and deviations from trend expectations are only typically recognised once the deviation has compounded over many years. For example, if a mortality improvement trend of 1% pa is anticipated, but the actual trend is 2% pa, it will usually be impossible to distinguish such a trend deviation from normal statistical fluctuations until at least 5-10 years (or more) have passed, by which time the impact of the trend will have already compounded to a 5-10% deviation in base mortality levels, in addition to the 1% pa deviation in future improvement trends. A writer of lifetime annuities who initially misestimates future mortality trends may find that it has written 5-10 years of business at inadequate annuity rates before discovering the mispricing error.

Likewise, the fact that new entrants to an annuity or pension portfolio tend to be concentrated around retirement age gives rise to potentially very long delays in recognising that experience has deviated from expectations. For example, a new writer of annuity business will have little mortality experience for lives aged 80+ until at least 15 to 20 years from commencement of the business.

Lags in the recognition of deviations in longevity from expectations magnify the risks associated with longevity business. Historical experience in the UK, where compulsory annuitisation of retirement savings has driven the formation of one of the largest worldwide markets for private sector lifetime annuities and pensions, provides ample evidence that deviations in longevity tend to be recognised in occasional bursts. Over the past decades, UK actuaries have consistently underestimated the level of mortality improvement, but the recognition of these errors has tended to manifest in the form of infrequent large shifts (as opposed to frequent small shifts) in mortality assumptions, usually coinciding with the release of a new industry mortality table or with the publication of a new set of future mortality improvement projections.

Recognition of increased longevity above the levels originally expected would affect major holders of longevity risk as follows:

- Insurance companies with a substantial portfolio of lifetime annuity business would need to increase policy reserves and solvency margins on in force annuity portfolios with adverse impacts on the profitability and solvency levels, as well as increasing prices for new annuities;
- Superannuation or pension funds offering significant defined benefit lifetime pensions would need to increase pension liabilities with corresponding fund surplus impacts , as well as increasing employer contributions; and
- Governments with obligations to pay lifetime pensions would need to reduce pensions, increase taxation revenues and/or decrease other forms of spending.

In order to mitigate these impacts, holders of significant amounts of longevity risk might look to hedge their longevity risk or transfer these risks to other organisations willing to take on longevity risk.

5.2 Requirements for Transferring Longevity Risk

In recognition of the way in that longevity risks emerge over time, there are a number of key requirements that need to be addressed in any proposal to hedge or transfer longevity risk:

- *Asset Risk and Longevity Risk*
  Insurance companies and superannuation funds holding significant amounts of longevity risk in the form of lifetime annuities and pensions also typically carry large amounts of asset risk on the same portfolios, with all future annuity or pension payments expected to be funded from existing reserves together with future investment income.

  Depending on the particular features of the benefits, there may be other risks embedded within the portfolio, including:
  - inflation risk, where the benefits are indexed with a variable rate of inflation;
  - dependent risks, where benefit continue to be paid to a surviving spouse or surviving dependents upon the death of the primary annuitant or pensioner (especially where data on the number and age of the dependents might be unavailable or unreliable); and
  - other "option" risks such as the ability to make lump sum withdrawals or surrender the annuity or pension for a cash value.

  For some products, such as "variable annuity" products offering lifetime guaranteed minimum withdrawal or income benefits, the asset risk and the longevity risk will be inextricably linked. The extent of any exposure to longevity risk will depend on asset returns, making it very difficult to separate asset and longevity risks.

  Holders of longevity risk may differ in terms of their desire to seek protection against the various risks embedded within the portfolio:
  - Some organisations may seek to protect against all risks, transferring the entire liability to another entity. For example, superannuation funds purchasing annuities from an insurance company for members with lifetime pension would fall into this category (buy-in or buy-out).
  - As an alternative, organisations could also separately obtain protection against asset and longevity risks with different counterparties. For example, asset risks could be managed via an interest rate swap, while longevity risks could be managed via a longevity swap (see section 5.3 below).
  - Other organisations may feel more comfortable with asset risk, but wish to protect against longevity risk via a longevity swap, noting that the conversion of uncertain future annuity cash flows into a stream of fixed future cash flows may facilitate the management of asset risks.
  - It should be noted that any inflation risks will usually be incorporated into asset risk solutions, while risks associated with dependents will typically be incorporated into longevity risk solutions.

  While the structure of the hedge or risk transfer solution should address the particular risks that the holder of the risk seeks to transfer, the remainder of this paper will focus on solutions that are designed to address longevity risk only.

- *Duration*
  Portfolios of lifetime annuity business usually have very long outstanding durations. With the oldest verified human lifespan at 122 years and 164 days (Ms Jeanne Calment of France), this means that a portfolio of annuitants currently aged 65 may have 57 years or more to run before the last annuitant dies. The effective duration of a portfolio may be further increased by the presence of indexation features, as well as the existence of benefits payable to reversionary spouses, who could be significantly younger then the primary annuitant.

  It is therefore important in any attempt to transfer longevity risk to ensure that longevity risks are transferred for a sufficient duration. Using the modelled results from section 4 of the paper, the following chart shows, for a male aged 65, the percentage deviation in the present value of a lifetime annuity (discounted at a net interest rate of 3% pa) for the 1st and 99th percentiles of the

Currie APC M3 model, split between successive 10 year periods following policy issue. The deviations are expressed as a percentage of the M3 central estimate annuity present value.

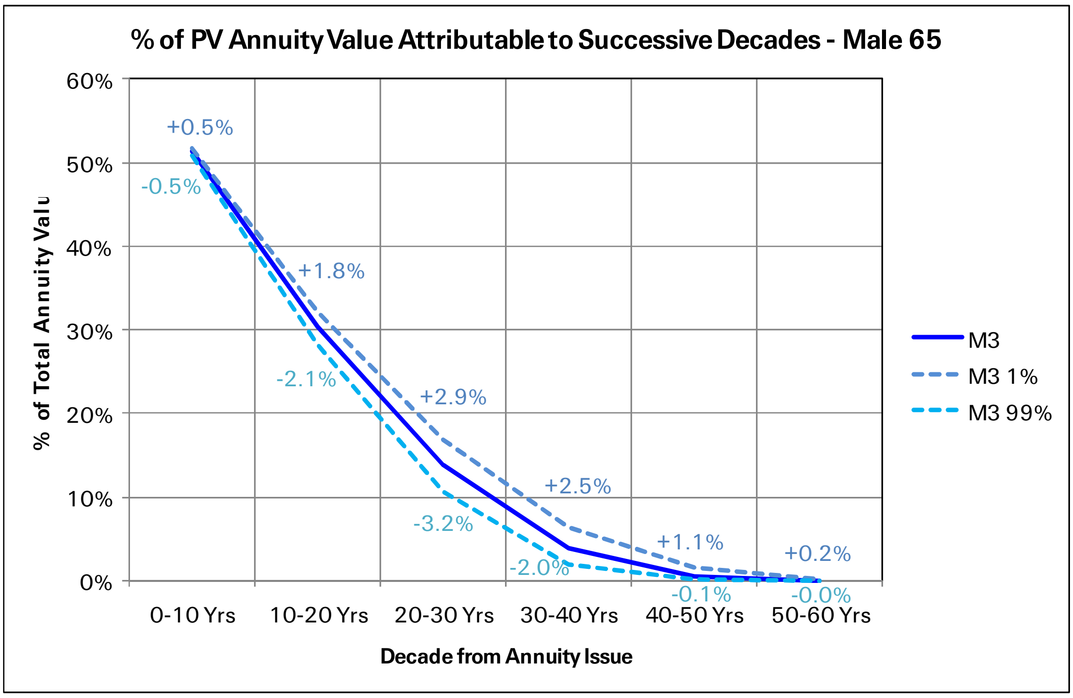

Figure 32: Australia Male 65 – Deviations in Annuity Value due to Model Shocks by Decade

Following the "M3" line, the chart shows that, under the central estimate for the Currie APC M3 model, 51.3% of the annuity present value relates to the first 10 years of the annuity, 30.3% of the annuity present value relates to next 10 years of the annuity, 13.9% relates to the period between 20 and 30 years, etc.

The 1st and 99th percentile shocks provide an indication of how little of the shock relates to the early years of the contract. For example, if we consider the 1st percentile fitting of the Currie APC M3 model (M3 1%), the total impact is to increase the annuity present value to 109.0% of the central estimate of the Currie APC M3 model. However:

- 0.5% of the total 9.0% impact relates to the first 10 years of the annuity
- 1.8% of the total 9.0% impact relates to the period from 10 to 20 years
- 2.9% of the total 9.0% impact relates to the period from 20 to 30 years
- 2.5% of the total 9.0% impact relates to the period from 30 to 40 years
- 1.1% of the total 9.0% impact relates to the period from 40 to 50 years
- 0.2% of the total 9.0% impact relates to the period beyond 50 years.

It should be noted that while the cash flow impact of deviations in longevity experience is heavily delayed, recognition of the deviation may take place earlier, with an adjustment in reserves to capitalise the expected future impact of the deviation. Nevertheless, the figures above demonstrate that a very long contract duration will usually be necessary to hedge or transfer substantial amounts of longevity risk. The availability of such long duration contracts may be difficult to find in capital markets that would typically have a time horizon of considerably less than 30 years.

- *Payment Trigger and Basis Risk*

  Some potential solutions for longevity protection make use of a "parametric" trigger to construct hedging instruments for longevity risk. These "parametric" triggers are based on the experience of a reference portfolio or parametric index (typically based on population mortality) rather than the actual mortality of the specific annuity or pension portfolio for which protection against longevity risk is sought.

  An example of longevity hedging solutions based on a parametric trigger are "q-Forward" derivatives providing for a payout equal to the difference between an agreed forward price, and the actual mortality rate experienced by a reference population (for a given gender and age band) at some agreed future point in time. The ability to construct a portfolio of q-Forwards and related longevity derivatives to replicate the expected cash flows under an annuity or pension portfolio would create a means for major holders of longevity risk to hedge their longevity risk. Organisations such as the Life & Longevity Markets Association (www.llma.org) are already pushing for the establishment of such a market. Ultimately, only the development of standardised longevity derivative instruments traded through the capital markets will be able to provide sufficiently deep and liquid markets to fully meet the demand for longevity risk protection.

While a detailed discussion of longevity derivatives is beyond the scope of this paper, we would highlight that hedging of longevity risk by using derivatives based on "parametric" triggers may leave significant "basis risk". That is, the mortality experience of the portfolio being hedged and the reference portfolio used to determine longevity derivative payouts may diverge materially, reducing the effectiveness of the hedge. While the use of a population-based index would provide a suitable hedge against macro longevity trends affecting the entire population (such as a cost-effective cure for cancer) and could be suitable for governments or other organisations with large and well-diversified longevity portfolios, basis risk could be significant for small portfolios and for portfolios with a high degree of concentration in particular sub-segments of the population (by socio-economic class, by industry, by occupation, etc).

From section 3 of the paper, we have already seen that old-age mortality varies significantly by socio-economic class, with evidence from the UK that life expectancy differentials by socio-economic class have been widening over time. That is, that mortality is improving faster among higher socio-economic classes than among lower socio-economic classes. Such divergences between the mortality of sub-segments of the population could be related to:
- differences in the underlying prevalence of smoking and differences in the rate of smoking cessation in different population segments;
- differences in the prevalence of certain infectious (eg HIV/AIDS) or chronic (eg diabetes) diseases;
- differences in attitudes towards or access to preventative healthcare and healthy lifestyle choices;
- etc.

Further problems with basis risk may arise from certain product features that may not be readily hedgeable with available longevity derivative instruments. These might include:
- indexation features that are not fixed in advance, such as CPI indexation and reversionary bonus additions, making it more difficult to quantify future longevity exposures;
- dependent risks associated with reversionary payments to spouses and/or dependents;
- annuities offering additional "options" such as lump sum withdrawals and surrender values;
- etc

By contrast, some longevity risk management solutions offer cover on an "indemnity" basis, whereby the payouts on the longevity hedge are based on the actual annuity or pension payouts made by the portfolio for which protection against longevity is sought. Indemnity solutions include many of the longevity risk transfer solutions offered by reinsurers as well as some capital market solutions placed with selected private and/or institutional investors. An aggregator with a large and diversified portfolio of longevity risks, such as a reinsurer, will generally be in a better position to take advantage of capital markets solutions based on parametric triggers.

- *Transaction Size and Risk Capacity*
  Organisations offering longevity protection will generally impose limits on the minimum and maximum amount of risk capacity that they are prepared to offer. Where an "indemnity" solution based on the experience of a specific portfolio is required, portfolios that are too small may not justify the amount of effort required to structure a longevity protection solution. This is especially the case for "indemnity" capital markets solutions for which the additional work required to structure the transfer to capital markets would usually require a larger minimum deal size than a private placement or reinsurance arrangement. By contrast, "parametric" solutions operate

independently of a specific portfolio and for the entity offering longevity protection, may be viable even for very small portfolios (provided the basis risk is tolerable).

For very large portfolios of longevity risk, other issues may arise. Organisations offering longevity protection will generally have a limited amount of capacity available and may restrict the amount of capacity available for a single transaction as well as in aggregate across a market and from a global perspective. For writers of longevity protection, it must be borne in mind that the very long contract durations required for effective longevity protection means that it will take a very long time before any utilised capacity can be released.

Hence, the size of the transaction and the risk capacity sought by the holder of longevity risk may influence the availability and relative efficiency of different longevity protection solutions.

- *Creditworthiness*
  Given the very long duration contracts required to effectively hedge or transfer longevity risk, the creditworthiness of the counterparty to any transfer of longevity risk is very important. The amount of exposure to counterparty credit risk will however depend on the structure of the longevity protection. A full transfer of asset and longevity risk will of course involve a much larger exposure to counterparty credit risk than a longevity swap structure where the net amount of counterparty exposure remains relatively small until there has been a very large deviation in longevity experience from original expectations.

  Counterparty credit risk can be managed in different ways:
  - As with any normal long-term insurance contract, organisations seeking to hedge their longevity risk may seek out counterparties that are set up and regulated to write contracts offering risk transfer over very long periods. Insurance and reinsurance companies are typically set up with this purpose in mind, holding prudent reserves and additional capital to back their long-term obligations, with regulatory frameworks in developed insurance markets designed to promote long term solvency.
  - Another mechanism for managing counterparty credit risk would be to spread the risk across multiple counterparties to limit exposure to any single counterparty at a tolerable level.
  - For less familiar or less creditworthy counterparties, arrangements for counterparties to post collateral based performance to date and expected future performance of the longevity hedge may be appropriate. Such arrangements are more common for longevity derivative contracts and for longevity hedges purchased via the capital markets, especially where the party providing the longevity protection intends to sell the risk on to a third party. It should however be borne in mind that collateral arrangements can create complexity and frictional costs for both parties to the transaction. In particular, if collateral obligations are symmetrical such that either party can be required to post collateral, the party seeking longevity protection may find that it needs to post collateral either at inception of the contract (as a longevity hedge will generally be out of the money at inception) or in case that longevity shortens (mortality increases) beyond initial expectations. Another point to note is that, much of the risk associated with a longevity contract relates to future mortality trends so that any collateral requirements will ultimately be based largely on a "mark-to-model" basis, introducing an element of model risk.

- *Administrative Simplicity and Flexibility*
  A final requirement not to be overlooked when organisations seek to hedge or transfer longevity risk is the complexity involved in the administration of hedging or risk transfer agreements. Solutions based on "parametric" triggers will generally be easier to administer as they are based on standardised instruments that make no reference to the underlying risk being hedged. However, for "indemnity" solutions, administrative complexity may arise because of the way that the business underlying the transaction is administered.

  As a rule, contracts involving single counterparties will usually be simpler and more flexible than those involving multiple counterparties or the capital markets and private investors. This may be a particular consideration for holders of longevity risk who experience problems in relation to the data underlying the portfolio of annuitants or pensioners. While capital markets solutions will generally be packaged in such a way as to make the subsequent correction of underlying data errors nearly impossible (possibly using a "shadow payroll" for annuity payments, leaving significant data risks with the purchaser of the longevity protection), reinsurers and single counterparties may exhibit a much greater degree of tolerance in relation to such problems, with more scope to make adjustments for historical data errors.

5.3 Longevity Swaps

Longevity swaps are one means of transferring longevity risk without transferring the corresponding asset risks on the annuity or pension portfolio. While the name "swap" has connotations with derivative and capital market instruments, longevity swaps can also be structured as insurance or reinsurance contracts.

In short, a longevity swap involves an agreement where:

- One party agrees to pay the other party a fixed series of cash flows (the "fixed" part of the swap, or the "premium" on an insurance contract); while
- The other party agrees to pay the actual (uncertain) cash flows relating to annuity or pension benefits (the "floating" part of the swap, or the "claims" on an insurance contract).

In this way, a holder of longevity risk can transform the uncertain future annuity or pension benefit cash flows into a fixed series of payments. The "fixed" or "premium" payments are set based on mortality assumptions of the organisation offering the longevity protection with some risk margin.

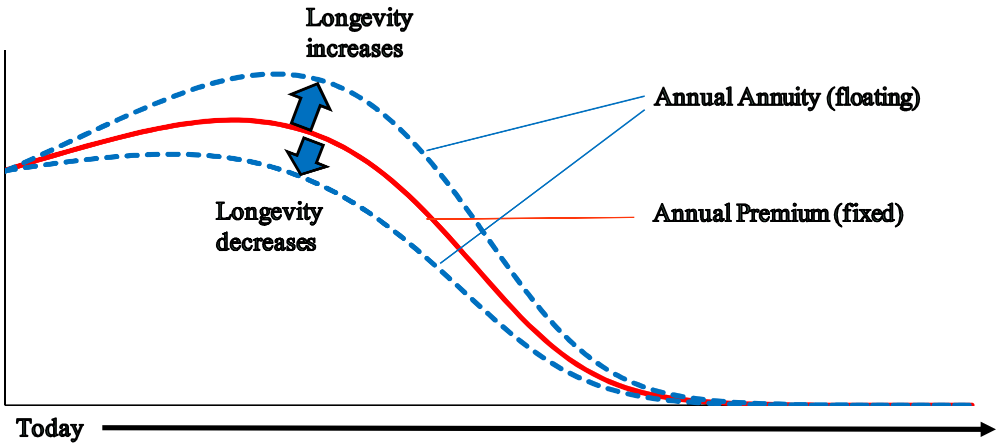

Figure 33: Operation of a Longevity Swap

5.4 Pricing Margins for Longevity Risk

Leaving aside the determination of a "best estimate" for future mortality and allowance for expenses, the margins that writers of longevity risk require for pricing purposes depend on a number of considerations:

- *Novelty of Risk*
  For any "new" type of capital market instrument, it has often been observed that investors demand a "novelty premium" over and above what might be required based purely on risk-neutral pricing theory. This is a phenomenon that has often been observed in the early days of insurance-linked securities looking to transfer insurance risk to the capital markets.

- *Scarcity of Longevity Risk*
  Normal considerations in terms of supply and demand also play a significant role in the pricing of longevity risks. In general, the total amount of market capacity that exists for organisations wishing to take on more longevity risk is thought to be far less than the amount of longevity risk currently held within insurance companies, pension funds and government entities. This probably also

contributes to market pricing that is higher than what a pure risk-neutral pricing theory would suggest.

- *Liquidity*
  Due to the long effective duration of longevity protection solutions and the very slow release of utilised longevity capacity, market pricing also reflects a liquidity premium.

- *Parametric Triggers*
  Because of the simplicity that comes with standardisation, longevity hedging solutions based on parametric triggers would generally be expected to be cheaper than indemnity solutions which require significantly more complex analysis.

- *Diversification Benefits*
  Part of the reason for the scarcity of longevity risk capacity relates to the way in which many organisations already hold significant amounts of longevity risk. Most prominently, many governments with social security and old-age pension obligations already hold much more longevity risk than would be desirable.

  However, for some investors, longevity may be viewed as a desirable risk because it is largely uncorrelated with other asset classes (eg equities, bonds, etc). At another level, there are some classes of investors for whom longevity risk would be negatively correlated with existing portfolios, potentially allowing them to reduce the risk margins that they would demand for taking on longevity risk.

  Key among those investor classes for whom longevity risk is negatively correlated with existing portfolios would be holders of significant amounts of mortality risk, including extreme mortality tail risks like pandemic risk and war risk. Most internationally diversified life reinsurers would fall into this category and this is one of the main reasons for the key interest of reinsurers in the market for longevity risk transfer. Because of the diversification benefits between lifetime annuities and pension business (carrying long-term longevity risks) with large in force portfolios of mortality business (especially long-term contracts with guaranteed premium rates), combined with synergies related to the natural fit between mortality and longevity research, large life insurers / reinsurers could regard themselves as a "natural home" for longevity risk.

  The following chart shows the level of correlation between Australian male mortality at ages 40-44 (representing a portfolio of mortality risks) and at ages 70-74 (representing a portfolio of annuitants). From the chart, it is evident that correlations between mortality rates at younger and older ages have fluctuated over time, partly due to the relatively small size of the Australian population, but have generally been positive, reflecting the "period effects" observed in mortality improvement fluctuations. There have also been numerous periods during which mortality improvements in the younger and older segments of the population have diverged, with zero or even negative correlations observed, reflecting the impact of cohort and other effects. Examples would include mortality improvements associated with road safety, which would affect younger lives much more than older lives, or mortality improvements associated with stroke, which would have the opposite effect.

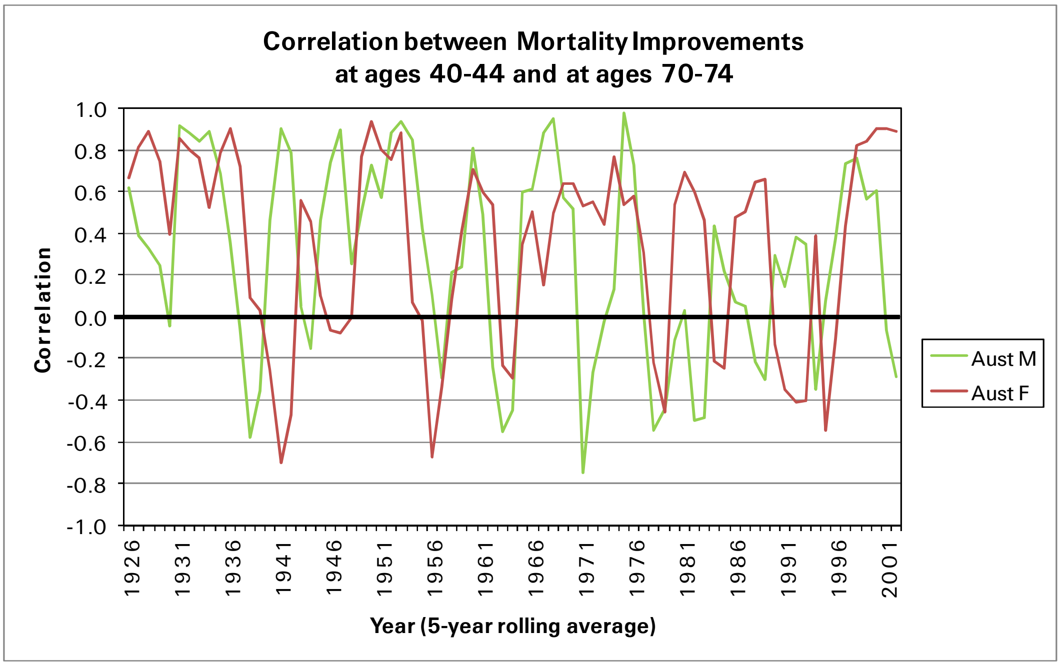

Figure 34: Australia Male – Correlation between Mortality Improvements at 40-44 and 70-74

The following table provides further confirmation of positive correlations between mortality improvements at different ages, showing for selected countries, the proportion of the time when the correlation between mortality improvements at ages 40-44 and 70-74 has been greater than some specified values over time:

Country	Period	Gender	ρ > 0.00	ρ > 0.25	ρ > 0.50	ρ > 0.75
Australia	1924-2004	Male	71.4%	54.5%	37.7%	19.5%
		Female	72.7%	64.9%	50.6%	23.4%
United Kingdom	1925-2004	Male	72.4%	64.5%	42.1%	21.1%
		Female	90.8%	82.9%	65.8%	36.8%
United States	1936-2004	Male	78.5%	64.6%	43.1%	30.8%
		Female	80.0%	76.9%	70.8%	40.0%

Another way to consider the diversification benefits of mortality and annuity portfolios would be to consider the impact of mortality improvements on sample term and annuity policies.

- For a term policy, we have calculated the impact of mortality improvements as the ratio of the (present value of death claims using actual Australian population *cohort* mortality for a policy *issued in year y*) to the (present value of death claims using actual Australian population *period* mortality for a policy in *year y*)[3];

---

[3] Formulaically, this ratio can be represented as:

$$\frac{\sum_{t=0}^{T-1} v^{t+0.5} \frac{l_{x+t}^{y+t}}{l_x^y} q_{x+t}^{y+t}}{\sum_{t=0}^{T-1} v^{t+0.5} \frac{l_{x+t}^{y}}{l_x^y} q_{x+t}^{y}}$$

where q_x^y is the mortality rate for age x in year y,

l_x^y is the number of lives alive at age x in year y

$l_{x+t+1}^{y+t+1} = l_{x+t}^{y+t}.(1-q_{x+t}^{y+t}).(1-lapse)$ represents survival based on cohort mortality starting from year y

$l_{x+t+1}^{y} = l_{x+t}^{y}.(1-q_{x+t}^{y}).(1-lapse)$ represents survival based on period mortality in year y

v^t is $(1+i)^{-t}$

- For an annuity, we have calculated the impact of mortality improvements as the ratio of the (present value of annuity claims using actual Australian population *cohort* mortality for a policy *issued in year y*) to the (present value of annuity claims using actual Australian population *period* mortality for a policy *in year y*)[4].

For the purposes of these calculations, we considered 5 and 20 year term assurance policies sold to a male aged 40 and a 20 year life annuity sold to a male aged 65, discounting cash flows at a net interest rate of 3% pa, with lapses of 10% pa for the term assurance policies. For example, in calculating the impact of mortality improvements on a 5-year term assurance policy issued to a male aged 40 in the year 2000, we firstly calculated the value of a term assurance using the mortality of a 40 year old in 2000, a 41 year old in 2001, a 42 year old in 2002, a 43 year old in 2003 and a 44 year old in 2004, and then compared this to the value of a term assurance using mortality for ages 40 to 44 in 2000.

The chart below shows the impact of mortality improvements on term and annuity policies using Australian population mortality rates for selected ages and policy terms for issue years from 1923 to 2003. A value greater than 100% means that mortality improvements have increased the value of claims under the policy, while a value less than 100% means that mortality improvements have decreased the value of claims under the policy.

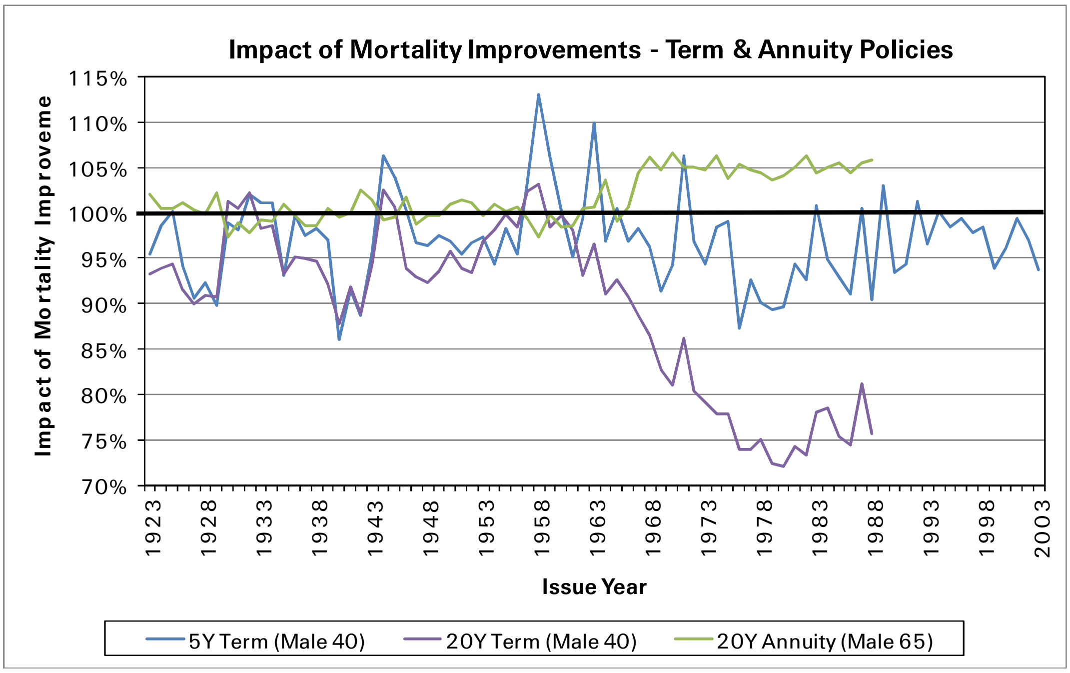

Figure 35: Australia Male – Impact of Mortality Improvements on Term & Annuity Policies

It is notable that the impact of mortality improvements has increased significantly since the mid-1960s although there are signs that rates of mortality improvement have slowed more recently. It should be noted that the impact of mortality improvements on annuities in the chart above is underestimated because a 20-year annuity is assumed (in order to for there to be sufficient mortality data to support the calculations for more recent years). Nevertheless, it is clear from historical Australian population mortality rates that there have been offsetting impacts for term and annuity policies as a result of mortality improvement.

---

[4] Formulaically, this ratio can be represented as: $\dfrac{\sum_{t=0}^{T-1} v^{t+0.5} \dfrac{l_{x+t+0.5}^{y+t}}{l_x^y}}{\sum_{t=0}^{T-1} v^{t+0.5} \dfrac{l_{x+t+0.5}^{y}}{l_x^y}}$

6. Conclusions

In this paper, we have considered the various uncertainties inherent in longevity risk, in respect of the current base mortality applicable to an annuitant or pensioner portfolio as well as the future mortality improvements that might be experienced by such a portfolio.

In relation to base mortality levels, we have seen Australian post retirement mortality improving rapidly in the past 35 years, more so for males than for females. The mortality levels vary substantially between different sub-groups within the overall population, with very large differences in mortality observed for lives in different socio-economic classes (approximated by benefit size, occupation, residence, etc) as well as material impacts on longevity resulting from self-selection, especially in markets where annuity purchase is a voluntary decision.

In relation to future mortality improvements, we have seen that differences in the chosen mortality model, the sample data or period used for fitting such models as well as statistical fluctuations can all have a very material impact on the valuation of annuity liabilities. Rather than advocating any particular model for projecting future mortality rates, we would instead urge actuaries to be vigilant in the use of such models, recognising that each model will have significant limitations when used for forecasting and using the models more as a means to quantify uncertainty, than as a means of determining a definitive best estimate of longevity liabilities.

For the consumer, awareness of the uncertainty surrounding their own potential longevity should increase the incentive to guarantee an income over their future lifetime. For the provider of such a financial product, that uncertainty should be a key element in forming product pricing and risk management. The level of uncertainty implied by the volatility that can result from using different, but similarly plausible models, as well as uncertainty related to the fitted parameters, underlying data and statistical volatility (before any allowance for diversification with other lines of business), tends to be at the upper end or even beyond the normal allowances built into Australia's solvency and capital adequacy standards.

With an improved understanding of the uncertainty in future mortality rates, a market for hedging and transfer of longevity risk has been developing in recent years. This is in spite of challenges posed by the long duration of the risks to be transferred, including the significantly delayed recognition of longevity experience deviations and the need for long-term creditworthy counterparties. A number of longevity swap contracts have already been executed (including in Australia) and many more are currently under discussion, with global life insurers and reinsurers, for whom longevity risk provides natural diversification against large mortality portfolios, acting as one of the few "natural" acquirers of pure longevity risk (without the corresponding asset risks). While the capital markets look set to develop further solutions for the hedging of longevity risk, it seems that a role will remain for intermediaries that are able provide more tailored solutions that are free of basis risk, warehousing portfolios of longevity risk with the option of ultimately transferring these risks to the capital markets.

References

[1] Australian Prudential Regulatory Authority (2007), *Prudential standard LPS 2.04 Solvency Standard and Prudential Standard LPS 3.04 Capital Adequacy Standard.*

[2] Booth H and Tickle L (2004), "Beyond Three Score Years and Ten: Prospects for Longevity in Australia", *People and Place*, vol.12, no.1, pp.15-27.

[3] Booth H and Tickle L (2008), "Mortality modelling and forecasting: A review of methods", Australian Demographic & Social Research Institute (ADSRI), Working Paper No.3, Australian National University.

[4] Coughlan G, *et al* (2007), *Life Metrics: A toolkit for measuring and managing longevity and mortality risks*, Version 1.1, March 13, JPMorgan Pensions Advisory Group.

[5] Cairns AJG, *et al* (2007), *A quantitative comparison of stochastic mortality models using data from England & Wales and the United States*, Pensions Institute Discussion Paper PI-0701, JPMorgan LifeMetrics.

[6] Cairns AJG, *et al* (2008a), *Mortality Density Forecasts: An Analysis of Six Stochastic Mortality Models*, Pensions Institute Discussion Paper PI-0801, JPMorgan LifeMetrics.

[7] Cairns AJG, *et al* (2008b), *Backtesting Stochastic Mortality Models: An Ex-Post Evaluation of Multi-Period-Ahead Density Forecasts*, Pensions Institute Discussion Paper PI-0803, JPMorgan LifeMetrics.

[8] Continuous Mortality Investigation (2002a), *Working Paper 1: An interim basis for adjusting the "92" Series mortality projections for cohort effects*, Institute and Faculty of Actuaries.

[9] Continuous Mortality Investigation (2002b), *Working Paper 2: Responses to the draft report entitled "A proposed interim basis for adjusting the "92" Series mortality projections for cohort effects"*, Institute and Faculty of Actuaries.

[10] Continuous Mortality Investigation (2004a), *Working Paper 3: Projecting future mortality: A discussion paper*, Institute and Faculty of Actuaries.

[11] Continuous Mortality Investigation (2004b), *Working Paper 11: Responses to Working Paper 3 entitled "Projecting future mortality: A discussion paper" and further commentary theron*, Institute and Faculty of Actuaries.

[12] Continuous Mortality Investigation (2005), *Working Paper 15: Projecting future mortality: A discussion paper: Towards a proposal for stochastic methodology*, Institute and Faculty of Actuaries.

[13] Continuous Mortality Investigation (2007a), *Revised Working Paper 20: Stochastic projection methodologies: Further progress and P-Spline model features, example results and implications*, Institute and Faculty of Actuaries.

[14] Continuous Mortality Investigation (2007b), *Revised Working Paper 25: Stochastic projection methodologies: Lee-Carter model features, example results and implications*, Institute and Faculty of Actuaries.

[15] Continuous Mortality Investigation (2007c), *Working Paper 30: The CMI Library of Mortality Projections*, Institute and Faculty of Actuaries.

[16] Continuous Mortality Investigation (2009a), *Working Paper 35: The graduation of the CMI Self-Administered Pension Schemes 2000-2006 mortality experience,* Institute and Faculty of Actuaries

[17] Continuous Mortality Investigation (2009b), *Working Paper 37: Version 1.1 of the CMI Library of Mortality Projections*, Institute and Faculty of Actuaries.

[18] Continuous Mortality Investigation (2009c), *Working Paper 38: A Prototype Mortality Projections Model: Part One – An Outline of the Proposed Approach*, Institute and Faculty of Actuaries.

[19] Continuous Mortality Investigation (2009d), *Working Paper 39: A Prototype Mortality Projections Model: Part Two – Detailed Analysis*, Institute and Faculty of Actuaries.

[20] Continuous Mortality Investigation (2009e), *Working Paper 41: CMI Mortality Projections Model: Feedback on Consultation and Issue of 'CMI_2009'*, Institute and Faculty of Actuaries.

[21] Continuous Mortality Investigation Reports (2009), *Number 23 Graduations of the 1999-2002 Life Office Mortality Experiences* , Institute and Faculty of Actuaries.

[22] Human Mortality Database (www.mortality.org)

[23] Humble R and Wilson B (2008), *Drivers of Longevity Projections in Particular with Reference to Smoking*, Presented to the Staple Inn Actuarial Society, 22 April.

[24] The Individual Annuity Experience Committee (2009), *2000-04 Individual Payout Annuity Experience Report*, Society of Actuaries.

[25] Johnson B & Blackwell L (2007), *Review of methods for estimating life expectancy by social class using the ONS Longitudinal Study*, Office for National Statistics.

[26] JP Morgan LifeMetrics (http://www.jpmorgan.com/pages/jpmorgan/investbk/solutions/lifemetrics)

[27] The Life Risk Insurance Committee (2004), *Report on the Lump Sum Experience Investigation 1998-99*, Institute of Actuaries of Australia.

[28] Madrigal AM, *et al* (2009), *What longevity predictors should be allowed for when valuing pension scheme liabilities?*, presented to the Institute of Actuaries Sessional Meeting, September.

[29] Office for National Statistics News Release (2007), *Variations persist in life expectancy by social class*, Office for National Statistics.

[30] Ribeiro R and di Pietro V (2009), "Longevity risk and portfolio allocation", *JP Morgan Investment Strategies No.57*, 5 June.

[31] Richards S & Jones G (2004), *Financial Aspects of Longevity Risk*, presented to the Staple Inn Actuarial Society, 26 October.

[32] Stevenson M & Wilson A (2008), *Mortality of Public Sector Scheme Pensioners 2005-2007 Update*, 4th Financial Services Forum, Institute of Actuaries of Australia, 19-20 May.

[33] Willets R (1999), *Mortality in the Next Millenium*, Presented to the Staple Inn Actuarial Society, 7 December.

[34] Willets R, *et al* (2004), *Mortality in the 21st Century*, Presented to the Faculty of Actuaries , 15 March, Presented to the Institute of Actuaries, 26 April.

[35] Yiasoumi C (2010), "Longevity swaps: Fad or saviour?", *The Actuary*, 1 April.