



City Research Online

City St George's, University of London

Citation: Blake, D. & Hunt, A. (2016). Basis Risk and Pension Schemes: A Relative Modelling Approach (PI-1601). London, UK: Pensions Institute.

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/17351/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).



DISCUSSION PAPER PI-1601

Basis Risk and Pensions Schemes: A Relative Modelling Approach.

Andrew Hunt and David Blake.

January 2016

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University London
106 Bunhill Row
London EC1Y 8TZ
UNITED KINGDOM

<http://www.pensions-institute.org/>

Basis Risk and Pension Schemes: A Relative Modelling Approach

Andrew Hunt*

Pacific Life Re, London[†]

Corresponding author: a.o.d.hunt.00@cantab.net

David Blake

Pensions Institute, Cass Business School, City University London

January 2016

Abstract

For many pension schemes, a shortage of data limits their ability to use sophisticated stochastic mortality models to assess and manage their longevity risk. In this study, we develop a relative model for mortality, which compares the evolution of mortality rates in a sub-population with that observed in a larger reference population. We apply this relative approach to data from the CMI Self-Administered Pension Scheme study, using UK population data as a reference. We then use the relative approach to investigate the potential differences in the evolution of mortality rates between these two populations and find that, in many practical situations, basis risk is much less of a problem than is commonly believed.

JEL Classification: C33, C51, C52, G22, J26

Keywords: Mortality modelling, age/period/cohort models, relative

*We are grateful to Andrés Villegas for many useful discussions on this and related topics.

[†]This study was performed when Dr Hunt was a PhD student at Cass Business School, City University London, and therefore the views expressed within it are held in a personal capacity and do not represent the opinions of Pacific Life Re and should not be read to that effect.

models, longevity basis risk

1 Introduction

Longevity risk is increasingly recognised as a major risk in developed countries, as rising life expectancies place unanticipated strains on social security and healthcare systems (see Oppers et al. (2012)). As well as being of concern for governments, however, longevity risk also affects private organisations that have promised people an income for life, be this in the form of an insured annuity or an occupational pension. In the UK, this means that longevity risk affects the thousands of occupational pension schemes¹ established by companies during the 20th century to provide final salary pensions to their employees.

However, when it comes to managing the longevity risk in a pension scheme, actuaries face a critical problem: a shortage of mortality data for the scheme. A typical UK pension scheme has fewer than 1,000 members and may have reliable, computerised member records going back little more than a decade. This is insufficient for use with the sophisticated stochastic mortality models that have been developed in recent years to measure longevity risk in national populations, since these models require more data to estimate parameters robustly and longer time series to make projections into the future. While the insights gained from the study of national populations are useful for the study of longevity risk in pension schemes, actuaries are left with a nagging doubt: “What if my scheme is different from the national population?” The potential for divergence in mortality rates between the scheme and the national population is often called “basis risk”, and, anecdotally, is often given as a key reason holding back the use of standardised financial instruments (based on national data) to manage longevity risk in pension schemes.

¹In this paper, we refer to “pension schemes” which administer the provision of defined benefits to members. We draw a semantic distinction between a “pension scheme” and a “pension plan”, which we would use as a more general term for any defined benefit or defined contribution pension arrangement provided on either a group or an individual basis.

The actuarial profession in the UK initiated the Self-Administered Pension Scheme study in 2002 in an attempt to overcome these issues with data. The study pools data from almost all large occupational pension schemes in the UK, allowing insights about how typical pension schemes differ from the national population to be established.

In this paper, we use the data collected by the Self-Administered Pension Scheme study and develop a “relative” model for mortality in order to compare the evolution of mortality rates in UK occupational pension schemes directly with that observed in the national population. Such a relative model has the advantages of parsimony and robustness, important properties when dealing with the smaller datasets available for pension schemes. We then use this relative model to investigate the phenomenon of basis risk between pension schemes and the UK population, as well as the potential of using this approach on even smaller populations comparable with the size of an individual scheme. In doing so, we bring into question the potential importance of basis risk in small populations and find that in most contexts it is likely to be substantially outweighed by other risks in a pension scheme. This is investigated further in Hunt and Blake (2015e).

The outline of this paper is as follows. Section 2 describes the Self-Administered Pension Schemes (SAPS) study and how the population observed by it differs structurally from the national UK population. Section 3 discusses the “relative” modelling framework we will use to compare the mortality experience of these populations. Section 4 then applies this framework to data from the SAPS study, tests the models produced and considers the impact of parameter uncertainty on these conclusions. Section 5 uses the relative model to project mortality rates for the sub-population in the context of assessing the basis risk between it and the national population. Section 6 then assesses the feasibility of using the relative model for smaller populations which have sizes more comparable to those of actual UK pension schemes. Section 7 discusses some of the broader conclusions on the importance of basis risk we draw from this study, whilst Section 8 summarises our findings.

2 The Self-Administered Pension Scheme study

The Institute of Actuaries in England & Wales and the Faculty of Actuaries in Scotland initiated the SAPS study in 2002 to investigate the mortality experience of pensioner members of occupational pension schemes in the UK. Data from the SAPS study has been analysed by the Continuous Mortality Investigation (CMI) to produce the graduated mortality tables² in use by the majority of pension schemes in the UK for funding and accounting purposes.³ The CMI has also analysed the SAPS data in terms of the evolution of mortality during the study period⁴ and the differences in experience for schemes whose employers are in different industries.⁵

UK pension schemes with more than 500 pensioner members are asked to submit mortality experience data to the SAPS study after each triennial funding valuation. The CMI provides summaries of the aggregate of this data to members of the study, categorised across a number of different variables, at regular intervals.⁶ We have been provided with this data in a more complete form, comprising exposures to risk and death counts (unweighted by the amount of pension in payment) for individual ages and years for all men and women in the SAPS study between 2000 and 2011 by the CMI. A summary of the data used in this paper is given in Appendix A.

Since it is sampling from a distinct subset of the national population, the dataset collected by the SAPS study is atypical of the UK population data for a number of reasons:

- The dataset is the mortality experience of members of occupational, defined-benefit pension schemes. Typically, this will exclude the unemployed, the self-employed, those employed in the informal sector or those working for newer companies (which typically do not offer defined-benefit pensions).
- The dataset is the mortality experience of members of reasonably large

²The S1 tables in Continuous Mortality Investigation (2008) and the S2 tables in Continuous Mortality Investigation (2014a).

³The Pensions Regulator (2013a) and Sithole et al. (2012).

⁴See Continuous Mortality Investigation (2011).

⁵See Continuous Mortality Investigation (2012).

⁶See Continuous Mortality Investigation (2014c) for example.

pension schemes. According to The Pensions Regulator (2013b), only around 20% of UK pension schemes have more than 1,000 member in total, and therefore even fewer pensioner members. This means that employees of large, mature companies are likely to be over-represented in the SAPS study.

- The dataset is the mortality experience of pension schemes subject to triennial funding valuations. This means that it excludes most public sector employees, who are members of unfunded state pension schemes.
- The dataset is likely to have some individuals in receipt of pensions from multiple sources, for instance, because of employment at two or more different companies, and who will therefore be represented multiple times.
- The dataset will include members of UK pension schemes who emigrate and possibly die overseas, and who therefore would not be included in the UK national population mortality data.

These factors explain why the experience of the SAPS mortality study is believed to be a better proxy for the mortality experience of individual UK pension schemes (even those not included in the SAPS study). The mortality tables graduated from the SAPS data are therefore often used for pension scheme accounting and funding purposes, as opposed to tables graduated from national population data or the experience of individuals buying annuities directly from life insurers. However, they also mean that the future evolution of mortality rates for SAPS members may be different from that of the national population (although they may well be similar in other respects).

Unfortunately, the SAPS dataset poses a number of difficulties for use with the more sophisticated mortality modelling and projection techniques which have been developed in recent years. These include:

- relatively small exposures to risk compared with the national population (at most around 1.5 million members under observation in a single year), leading to greater parameter uncertainty especially in complex models;
- the short length of the study, with only twelve years of data in the sample for analysing the trends present; and

- the method of data collection - schemes submit data in respect of a three-year period at a lag of up to 18 months after the period ends - leads to a distinctive pattern of exposures shown in the data in Appendix A, with only partial data having been submitted to date for the last five years in the study.

For these reasons, it is still advisable to use national mortality data, with its larger exposures and longer period of availability, to produce projections of mortality rates. The SAPS data can then be used to quantify the ways that members of UK pension schemes are likely to differ from this baseline. We do this by means of a “relative” mortality model, which we now describe.

3 Relative mortality modelling

A “relative” mortality model for two populations is one that does not model mortality rates in a smaller population directly, but instead models the relative difference between those rates and those found in a larger, reference population. That is, it models the behaviour of the relative mortality rates, $R_{x,t}$, given by

$$R_{x,t} = f \left(\frac{\mu_{x,t}^{(S)}}{\mu_{x,t}^{(R)}} \right) \quad (1)$$

where $\mu_{x,t}^{(p)}$ are the mortality rates in the small population, S , and reference population, R . Typically, mortality rates in the reference population are modelled and projected independently of $R_{x,t}$.

A number of different models of this form have been proposed in order to analyse mortality for various different populations. Those which have explicitly adopted a relative modelling approach include the models of Jarner and Kryger (2011), which used a series of basis functions across age to model $R_{x,t}$ for Denmark compared to the wider EU and assume it mean reverts deterministically in future, and Villegas and Haberman (2014), which investigated the mortality of different socio-economic groups within the UK relative to the national average. However, a good many other multi-population mortality models which have been proposed, such as those of Carter and Lee (1992), Li and Lee (2005), Delwarde et al. (2006), Dowd et al. (2011), Cairns

et al. (2011b), Russolillo et al. (2011) and Wan and Bertschi (2015), can be rewritten as relative mortality models although this was not necessarily commented on by the authors. See Villegas and Haberman (2014) for a useful summary of many of these models and the similarities between them.

The advantage of a relative modelling approach is that it allows us to use a far simpler model for the relative mortality rates, $R_{x,t}$, than would be used for the reference population. This is desirable as we typically have insufficient data for the smaller population to estimate more complex models robustly, but would like to use a sophisticated model for the reference population in order to produce more accurate projections of mortality rates. In addition, there is no requirement that the data for the small population covers the same range of ages and years as that for the larger population.

3.1 The reference model

For the reference population, we choose to use the “general procedure” (GP) of Hunt and Blake (2014) in order to construct a model sufficient to capture all the significant information present in the national population data. This selects an appropriate model within the class of age/period/cohort (APC) models⁷ of the form

$$\ln \left(\mu_{x,t}^{(R)} \right) = \alpha_x^{(R)} + \sum_{i=1}^N f^{(R,i)}(x; \theta^{(R,i)}) \kappa_t^{(R,i)} + \gamma_{t-x}^{(R)} \quad (2)$$

where

- age, x , is in the range $[1, X]$, period, t , is in the range $[1, T]$ and hence that year of birth, y , is in the range $[1 - X, T - 1]$;
- $\alpha_x^{(R)}$ is a static function of age;
- $\kappa_t^{(R,i)}$ are period functions governing the evolution of mortality with time;
- $f^{(R,i)}(x; \theta^{(R,i)})$ are parametric age functions (in the sense of having a specific functional form selected a priori) modulating the impact of

⁷See Hunt and Blake (2015d) for a description of this class of models.

the period function dynamics over the age range, potentially with free parameters $\theta^{(R,i)}$,⁸ and

- $\gamma_y^{(R)}$ is a cohort function describing mortality effects which depend upon a cohort's year of birth and follow that cohort through life as it ages.

The GP selects the number of age/period terms, N , and the form of the age functions $f^{(R,i)}(x)$ in order to construct mortality models which give a close but parsimonious fit to the data. This way, we aim to extract as much information as possible from the national population dataset and have specific terms within the model corresponding to the different age/period or cohort features of interest.

3.2 The relative model

To analyse the data from the SAPS study, we propose using a model of the form

$$R_{x,t} = \ln \left(\frac{\mu_{x,t}^{(S)}}{\mu_{x,t}^{(R)}} \right) = \alpha_x^{(\Delta)} + \sum_{i=1}^N \Lambda^{(i)} f^{(R,i)}(x) \kappa_t^{(R,i)} + \Lambda^{(\gamma)} \gamma_{t-x}^{(R)} + \nu X_{t-x} \quad (3)$$

Apart from the νX_y term, this is an APC model of the same form as that used to model the reference population, i.e., with the same age/period terms and cohort parameters. However, these are modulated by factors, $\Lambda^{(j)}$ where $j \in \{1, \dots, N, \gamma\}$. The νX_{t-x} term, where X_y is a set of deterministic functions of year of birth and ν the corresponding regression coefficients, has been added to the APC structure in order to ensure that the model is identifiable under invariant transformations of the cohort parameters, as discussed in Appendix B.

The choice of structure in Equation 3 is also motivated by the fact that we can write the mortality rates for the sub-population as

$$\ln \left(\mu_{x,t}^{(S)} \right) = \alpha_x^{(S)} + \sum_{i=1}^N \lambda^{(i)} f^{(R,i)}(x) \kappa_t^{(R,i)} + \lambda^{(\gamma)} \gamma_{t-x}^{(R)} + \nu X_{t-x} \quad (4)$$

⁸For simplicity, the dependence of the age functions on $\theta^{(R,i)}$ is suppressed in notation used in this paper, although it has been allowed for when fitting the model to data.

where $\alpha_x^{(S)} = \alpha^{(R)} + \alpha_x^{(\Delta)}$ and $\lambda^{(j)} = 1 + \Lambda^{(j)}$. We are therefore able to interpret $\alpha_x^{(\Delta)}$ as the difference in the level of mortality between the two populations, whilst the $\lambda^{(j)}$ correspond to the “sensitivity” of the small population to the j^{th} factor in the reference population. This form is often more convenient to use in practice, since it refers solely to mortality rates in the sub-population. However, it is important to bear in mind that it is still a relative model in structure. Further, in the form of Equation 4, it is possible to see the model as similar in spirit to that proposed by Russolillo et al. (2011), as discussed in Section 3.3.

It should be noted that there are two special cases for these sensitivities:

1. $\lambda^{(j)} = 0$ (i.e., $\Lambda^{(j)} = -1$): the small population has no dependence on the j^{th} age/period or cohort term; and
2. $\lambda^{(j)} = 1$ (i.e., $\Lambda^{(j)} = 0$): there is no difference between the reference and small populations with respect to the j^{th} factor.

In order to obtain a more parsimonious model, it may also be desirable to simplify the non-parametric structure⁹ for $\alpha_x^{(\Delta)}$ by constraining it to be of a specific parametric form, for example, a linear combination of a set of pre-defined basis functions. However, we must take care when doing so in order that the relative model is robust to changes in the identifiability constraints for the reference model, as discussed in Appendix B.

When fitting the relative model to data, we have a strong preference for parsimony due to the low volume of data for the sub-population. We therefore adopt a “specific-to-general” modelling approach: first testing a highly restricted form of the model with a parametric form for $\alpha_x^{(\Delta)}$ and $\lambda^{(j)} = \{0, 1\}$ and then relaxing these restrictions sequentially. The final model is chosen to maximise the Bayes Information Criteria (BIC),¹⁰ which penalises excessive parameterisation. This procedure is performed algorithmically, and is especially important when we apply the relative model to very small datasets comparable to the size of individual pension schemes, as done in Section 6.

⁹Defined in Hunt and Blake (2015d) as being fitted without any a priori structure or functional form.

¹⁰Defined as $\max(\text{Log-likelihood}) - 0.5 \times \text{No. free parameters} \times \ln(\text{No. data points})$.

3.3 Comparison with “three-way Lee-Carter”

It was noted above that many alternative multi-population mortality models have been proposed in the literature, including many which were explicitly designed as relative mortality models and others which can be re-written in relative form. For a summary and comparisons of some of these models, see Li and Hardy (2011) and Villegas and Haberman (2014).

Of these, the model which bears closest resemblance to the model outlined in Section 3.2 is the “three-way Lee-Carter” model of Russolillo et al. (2011). This extends the classic model of Lee and Carter (1992) into a third “dimension” of population, beyond the original two dimensions of age and period. They achieve this by including an extra covariate in the Lee-Carter predictor structure to represent the different populations, p , i.e.,

$$\ln \left(\mu_{x,t}^{(p)} \right) = \alpha_x^{(p)} + \lambda^{(p)} \beta_x \kappa_t \quad (5)$$

The parameters are fitted using multi-dimensional principal components techniques. Villegas and Haberman (2014) pointed out that an additional identifiability constraint is required to obtain a unique set of parameters, which they choose to be $\sum_p \lambda^{(p)} = N_p$, the number of populations. In a two-population setting, this can be re-written as a relative model, with

$$R_{x,t} = \alpha_x^{(\Delta)} + (\lambda^{(S)} - \lambda^{(R)}) \beta_x \kappa_t$$

and $\alpha_x^{(\Delta)}$ defined in the same fashion as in Equation 3.

We can, therefore, see that the relative model of Section 3 can be thought of as a “three-way” extension for multiple populations of the underlying model constructed by the general procedure for a single population, namely

$$\ln \left(\mu_{x,t}^{(p)} \right) = \alpha_x^{(p)} + \sum_i \lambda^{(p,i)} f^{(i)}(x) \kappa_t^{(i)} + \lambda^{(p,\gamma)} \gamma_{t-x}$$

We then introduce the νX_y term in order to ensure that the model does not depend upon the arbitrary identifiability constraints imposed in the reference model, as discussed in Appendix B. In our relative model, however, we set $\lambda^{(R,j)} = 1 \ \forall j$, as opposed to the $\lambda^{(R,j)} + \lambda^{(S,j)} = 1 \ \forall j$ constraint in Villegas and Haberman (2014). Our identifiability constraint implicitly establishes a

hierarchy between the populations, with population S subordinate to population R . Setting $\lambda^{(R,i)} = 1$ motivates the two-stage fitting process, with the age/period and cohort terms being fitted using data for the reference population alone. In our context, as the two populations are of very different sizes, this is both reasonable and unlikely to make a material difference to the fitted parameters. However, it means that the fitted parameters for the sub-population are conditional on those found for the reference model. It is, therefore, important that tests of the model include full allowance for parameter uncertainty in both populations.

As with the model of Russolillo et al. (2011) and the analysis of Villegas and Haberman (2014), it is also possible to apply our model to multiple sub-populations, such as those from different pension schemes. In this case, separate scaling factors would be required for each scheme. For multiple schemes, the hierarchical structure of the model is an advantage, since each scheme can be considered separately once the reference population has been estimated.


4 Applying the relative model to SAPS data

4.1 The reference models for UK data


Our first task is to construct suitable mortality models for men and women in the national UK population. To do this, we apply the GP to data from the Human Mortality Database (2014) for the period 1950 to 2011 and for ages 50 to 100. The GP produces a model with four age/period terms, described in Table 1,¹¹ plus cohort terms for both men and women in the UK. All of these terms are shown in Figures 1 and 2. Further details of the age functions used in this model and tests of the goodness of fit to data are given in Appendix C.

In Figures 1c and 2c, the most notable features of the cohort parameters for both men and women are the presence of large outliers in 1919/20 and 1946/47. We believe, based on the analysis of Richards (2008), that these


¹¹Demographic significance, as used in Table 1, is defined in Hunt and Blake (2015d) as the interpretation of the components of a mortality model in terms of the underlying biological, medical or socio-economic causes of changes in mortality rates which generate them.



(a) Age functions




(b) Period functions




(c) Cohort function


Figure 1: Age, period and cohort functions in the reference model for men in the UK



(a) Age functions



(b) Period functions



(c) Cohort function

Figure 2: Age, period and cohort functions in the reference model for women in the UK

Term	Description	Men	Women	
		Demographic Significance	Description	Demographic Significance
$f^{(R,1)}(x)\kappa_t^{(R,1)}$	Constant age function	Level of mortality curve	Constant age function	General level of mortality
$f^{(R,2)}(x)\kappa_t^{(R,2)}$	Linear age function	Slope of mortality curve	Linear age function	Slope of mortality curve
$f^{(R,3)}(x)\kappa_t^{(R,3)}$	Parabolic age function	Mid age-range mortality	Parabolic age function	Mid age-range mortality
$f^{(R,4)}(x)\kappa_t^{(R,4)}$	Parabolic age function	Younger age mortality	Parabolic age function	Younger age mortality

Table 1: Terms in the reference models constructed using the general procedure

are not genuine cohort effects, but are merely data artefacts arising from the surge of births following the large-scale demobilisations after the First and Second World Wars, which biases the calculation of the exposures to risk in the UK population data for those years. We do not expect to find similar outliers in the SAPS data as this is based on aggregating individual scheme-member data rather than population level estimates.¹² One method to solve this would be to adjust the UK population exposures data as proposed in Cairns et al. (2015). However, for simplicity, we choose to retain the original data and employ indicator variables to remove the impact of outliers from the relevant cohort parameters. These adjusted cohort parameters are then used in the analysis which follows.¹³

¹²This is borne out by using simple APC models fitted to the SAPS data, which show cohort parameters without these outliers.

¹³It is interesting to note that these outliers may impact the effectiveness of hedging strategies which use securities indexed to national population data, as the index will continue to show a large (but fictitious) effect for specific cohorts which will not be observed in the specific population being hedged. It is therefore important that any indices use national population data which has been adjusted to remove these data artefacts, possibly using the approach of Cairns et al. (2015).

As discussed in Hunt and Blake (2015a,b), many mortality models are not fully identified. To uniquely specify the parameters, we impose identifiability constraints. These constraints are arbitrary, in the sense that they do not affect the fit to data. However, they can be used to impose our desired demographic significance on the parameters.

Models generated by the GP impose the following standard identifiability constraints

$$\sum_{x=50}^{100} |f^{(R,i)}(x)| = 1 \quad \forall i, \quad R = \{\text{UKm}, \text{UKf}\} \quad (6)$$

on the age functions to ensure that they have a consistent normalisation scheme. This enables us to compare the magnitudes of the period functions both with each other and between populations and gauge their relative importance.¹⁴

In order to assist the visual comparison between the UK and SAPS data (the latter of which only spans ages 60 to 90 and years 2000 to 2011), we impose the following constraint on the period functions

$$\sum_{t=2000}^{2011} \kappa_t^{(R,i)} = 0 \quad \forall i, \quad R = \{\text{UKm}, \text{UKf}\} \quad (7)$$

This means that the period functions represent deviations from an “average” level of mortality in the period covered by the SAPS data, rather than over the whole period of the UK data.

The results of Hunt and Blake (2015b) also indicate that we need to impose constraints on the levels, linear and quadratic trends present in the cohort parameters. To identify their levels, we impose the following constraints on the cohort parameters for each of the reference populations

$$\sum_{y=1910}^{1951} n_y^{(S)} \gamma_y^{(R)} = 0, \quad R = \{\text{UKm}, \text{UKf}\} \quad (8)$$

$$S = \{\text{SAPSm}, \text{SAPsf}\}$$

¹⁴For both women and men, the second and third age/period terms use age functions which are “self-normalising” in the sense of Hunt and Blake (2015a).

where $n_y^{(S)}$ is the number of observations of each cohort in the SAPS data. As with the period functions, this means that the cohort parameters should be centred around zero over the range of the SAPS data, not the full range of the data covered for the UK population. To constrain the linear and quadratic trends in the cohort parameters, we impose

$$\sum_{y=1850}^{1961} n_y^{(R)} \gamma_y^{(R)} (y - \bar{y}) = 0, \quad RR = \{\text{UKm, UKf}\} \quad (9)$$

$$\sum_{y=1850}^{1961} n_y^{(R)} \gamma_y^{(R)} ((y - \bar{y})^2 - \sigma_y) = 0 \quad (10)$$

where $n_y^{(R)}$ is the number of observations of each cohort in the UK national data, $\bar{y} = 0.5(X + T - 1)$ and $\sigma_y = \frac{1}{X+T-1} \sum_y (y - \bar{y})^2$.

The justification for these constraints is that they allow us to remove linear and quadratic trends in the cohort parameters. This makes them conform better to the demographic significance for cohort parameters described in Hunt and Blake (2015d), namely that the cohort parameters should not have any long-term systematic trends. We impose this over the whole range of the UK data, which is considerably longer than the range covered by the SAPS data, since there appear to be short-term trends (lasting for a few decades) which are then reversed out over a longer time horizon. However, this means that over the shorter range of years of birth covered by the SAPS data, the cohort parameters appear to have strong trends.

It is important to note, however, that our demographic significance for the parameters is highly subjective and our choice of constraints is arbitrary. We have therefore taken appropriate steps in Appendix B to ensure that our choice of identifiability constraints does not affect either the mortality rates fitted by the relative model or our overall conclusions.

4.2 The relative models for the SAPS data

We now estimate the relative model using these reference age, period and cohort terms for the full SAPS dataset. As discussed in Section 3, we do this in stages using a specific-to-general procedure. We start with the simplest

and most restricted model, i.e., where $\alpha_x^{(\Delta)}$ is restricted to take a parametric form and we restrict the scaling factors $\lambda^{(j)}$ to be equal to zero. This model is referred to as Model 1 in Tables 2 and 3 below.

We then allow these restrictions to be relaxed sequentially. This means that, in turn, we estimate the relative model with all possible combinations of constraints, where $\alpha_x^{(\Delta)}$ is either parametric or non-parametric and where $\lambda^{(j)}$ can be restricted to be equal to zero, unity or allowed to vary freely. This gives us $486 (= 2 \times 3^5)$ different combinations of constraints for the two alternative structures for $\alpha_x^{(\Delta)}$ and three alternatives for each of the five different scaling factors, $\lambda^{(j)}$. For each of these different models, the goodness of fit to the data is calculated, as measured by the BIC. The model which gives the best fit to data (i.e., the highest BIC) is then selected as the preferred model, referred to as Model 8 in Tables 2 and 3, for the dataset. This process is illustrated in Figure 3.

Several of the models tested, with representative combinations of restrictions, are shown in Tables 2 and 3 for the male and female SAPS data.¹⁵ These have been chosen to illustrate the impact of relaxing various restrictions, for instance, comparing Models 1 and 2 illustrates the impact on the goodness of fit of using a non-parametric as opposed to a parametric structure for $\alpha_x^{(\Delta)}$, whilst comparing Models 3 and 4 illustrates the impact of introducing the set of cohort parameters from the reference population. The preferred model which maximises the fit to data is shown as Model 8. However, it is important to note that the fitting procedure tests all 486 possible combinations for the structure of $\alpha_x^{(\Delta)}$ and any combination of restrictions on $\lambda^{(j)}$.

For both men and women, the preferred model selects a parametric simplification for the difference in the level of mortality, $\alpha_x^{(\Delta)}$. This substantially reduces the number of free parameters in the preferred model, leading to greater parsimony. This is also borne out by comparing models which differ by the form of $\alpha_x^{(\Delta)}$, but have similar restrictions placed on the scaling factors, $\lambda^{(j)}$, e.g., Models 1 and 2, or Models 4 and 5 in Tables 2 and 3. In some respects, this supports the traditional actuarial practice of adjusting mortality rates for a pension scheme by taking a mortality table from a refer-

¹⁵In Tables 2 and 3, “NP” stands for non-parametric while “P” stands for parametric.




Figure 3: Flow chart illustrating the procedure for fitting and selecting the relative model

Model No.	1	2	3	4	5	6	7	8
$\alpha^{(\Delta)}$	P	NP	P	P	NP	P	NP	P
$\lambda^{(1)}$	0	0	1	1	1	1.36	1.37	1
$\lambda^{(2)}$	0	0	1	1	1	0.34	0.28	1
$\lambda^{(3)}$	0	0	1	1	1	1.12	1.18	1
$\lambda^{(4)}$	0	0	1	1	1	1.29	0.59	1
$\lambda^{(\gamma)}$	0	0	0	1	1	1.00	0.51	1
Log-likelihood $\times 10^3$	-2.04	-1.93	-1.98	-1.93	-1.86	-1.92	-1.85	-1.93
Free parameters	5	32	5	5	32	10	37	5
BIC $\times 10^3$	-2.06	-2.03	-1.99	-1.94	-1.95	-1.95	-1.96	-1.94

Table 2: Representative sets of restrictions for the relative model using male SAPS data

ence population (in this case, the full UK population) and making relatively simple adjustments to it. We also see from Figures 5a and 5b that $\alpha_x^{(\Delta)}$ is generally negative across all ages. This indicates that the SAPS population has generally lower levels of mortality rates than the national population, which is consistent with the results of Continuous Mortality Investigation (2011).

In the case of the male data, the procedure selects a model where all the $\lambda^{(i)}$ for the age/period and cohort terms are restricted to be equal to unity. This means that the model finds no difference between the evolution of mortality rates for men in the SAPS data and the national population.¹⁶ For the female SAPS data, $\lambda^{(1)}$, $\lambda^{(3)}$ and $\lambda^{(\gamma)}$ are allowed to vary freely, i.e., without any restrictions placed upon them at the estimation stage. This can be compared with Model 6 in Table 3, where these parameters are allowed to vary freely, but happen to take values relatively close to unity. This is developed further in Section 4.3.1.

We also see that the scaling factors for the period functions for women are greater than unity when their estimation is not restricted. Since mortality rates for women are generally falling in the UK, this implies that the rate of improvement in longevity is slightly faster for female members of occupa-

¹⁶In the terminology of Section 5, we say the model finds that there is level basis, but no trend basis.

Model No.	1	2	3	4	5	6	7	8
$\alpha^{(\Delta)}$	P	NP	P	P	NP	P	NP	P
$\lambda^{(1)}$	0	0	1	1	1	1.52	1.38	1.54
$\lambda^{(2)}$	0	0	1	1	1	1.07	1.06	1
$\lambda^{(3)}$	0	0	1	1	1	1.94	1.54	1.91
$\lambda^{(4)}$	0	0	1	1	1	1.22	0.68	1
$\lambda^{(\gamma)}$	0	0	0	1	1	0.83	0.76	0.78
Log-likelihood $\times 10^3$	-1.94	-1.91	-1.84	-1.81	-1.78	-1.79	-1.77	-1.79
Free parameters	5	32	5	5	32	10	37	8
BIC $\times 10^3$	-1.95	-2.00	-1.86	-1.82	-1.88	-1.82	-1.87	-1.81

Table 3: Representative sets of restrictions for the relative model using female SAPS data

tional pension scheme than for the national population. This contrasts with the findings of Continuous Mortality Investigation (2011), which found that the falls in standardised mortality ratios for the SAPS populations broadly mirrored the falls observed in the wider UK population. However, since the standardised mortality ratio is an aggregate measure of mortality, which takes account of the level of mortality rates, it is likely that the difference between our results and those of Continuous Mortality Investigation (2011) are not significant.

Finally, we note that the BICs of many of the models with different restrictions are very similar, meaning that there is not much to choose between them. This is developed further in Section 4.3.2. It may therefore be justifiable to select simpler models than suggested by looking just at goodness of fit, on the grounds that they may be more robust to parameter uncertainty or easier to project into the future, as done in Section 5. This will be even more important when we investigate smaller, pension scheme-sized datasets, as in Section 6.

4.3 Parameter uncertainty and model risk

We next consider the robustness of the preferred model selected, i.e., Model 8. We do this in two stages, by considering the different sources of uncertainty

outlined in Cairns (2000). First, we consider only parameter uncertainty, i.e., the uncertainty in the free parameters of the preferred model, on the assumption that the restrictions placed on the parameters in Model 8 are correctly specified. Second, we allow for model risk by allowing the procedure to select different models using the sequential procedure discussed above.

For both stages, we use a procedure based on the residual bootstrapping method of Koissi et al. (2006) to generate new pseudo-data. This resamples from the fitted residuals to generate new simulated death counts to which the model is refitted, allowing the uncertainty in the parameters to be measured. We do this first to allow for parameter uncertainty in the reference model. It is important to allow for parameter uncertainty in the reference model due to the hierarchical structure of the relative model, i.e., that the parameters for the reference model are implicitly assumed to be known when the relative model is fitted. Therefore, uncertainty in the parameters of the reference model can be magnified when we come to investigate the uncertainty in the parameters of the relative model.

The next step is to bootstrap new pseudo-data for the sub-population. When using a residual bootstrapping procedure, it is important that the fitted residuals being used contain as little structure as possible, so that as little information as possible in the original data is lost when these residuals are randomly resampled. This will be the case for models which provide a close fit to the data, i.e., a high maximum likelihood. Therefore, in our residual bootstrapping procedure we use the expected mortality rates and fitted residuals from Model 7, since this model has the highest log-likelihood in Tables 2 and 3 above. However, since Model 7 is outperformed by a number of other models when the goodness of fit is penalised for the number of parameters (i.e., it has lower BIC than other models), we do not specifically consider it further.

4.3.1 Parameter uncertainty

For the first stage, we consider only parameter uncertainty. To do this, we fit the relative model to 1,000 sets of pseudo death counts, generated by the Koissi et al. (2006) residual bootstrapping procedure. For each of these datasets, however, we do not test which set of restrictions give the best fit




Figure 4: Flow chart illustrating the procedure for fitting and selecting the relative model allowing for parameter uncertainty

to the data. Instead, we impose the same set of restrictions as were used for Model 8 in Tables 2 and 3. We fit the relative model with the restrictions in Model 6 (which allows all scaling factors to freely vary) used as a comparator. This process is illustrated in Figure 4.

Figure 5 shows the impact of parameter uncertainty on the level parameters for the SAPS population by showing the 95% fan chart. We see that the width of the confidence intervals for $\alpha_x^{(\Delta)}$ is approximately 0.025 for men and 0.05 for women across most ages. Because we are using a log-link function, these can be converted into confidence for the fitted mortality rates.¹⁷ This uncertainty limits the extent to which we can tell that there are systematic differences between populations. Uncertainty in the level of mortality in the

¹⁷To illustrate, if the model gave fitted mortality rates for men in the SAPS population of 2.00% at some age, this value has an approximate confidence interval of $[2.00\% - 0.025 \times 2.00\%, 2.00\% + 0.025 \times 2.00\%] = [1.95\%, 2.05\%]$.




Figure 5: 95% fan charts showing the level of parameter uncertainty in $\alpha_x^{(\Delta)}$

sub-population can have important implications for the cashflows from pension schemes, as we discuss in Hunt and Blake (2015e).

The dashed lines in Figure 5 show the parameter-certain estimates of $\alpha_x^{(\Delta)}$, which lie close to the centre of the confidence intervals given by relative models.¹⁸ This indicates that our method for estimating parameter uncertainty does not significantly bias the results, which is an important check of its suitability.

Table 4 shows the 95% confidence intervals for the scaling factors for men and women. The first thing to note from these results is that the scaling factors are subject to substantial parameter uncertainty. As the relative model is very parsimonious and contains relatively few free parameters, this should caution us against considering more sophisticated models for the SAPS population. For instance, we are unlikely to have sufficient data to robustly estimate separate period functions for the SAPS data compared with the reference population, which was done in Villegas and Haberman (2014).

From Table 4, we can easily apply a simple but important check of our

¹⁸There appears to be a systematic difference in the linear trend in $\alpha_x^{(\Delta)}$ for women, although this is of less concern, since such trends are not identified by the model and so are entirely dependent on the identifiability constraints in the model. See Appendix B for a discussion of this subject.

	Men		Women	
	Model 6	Model 8	Model 6	Model 8
$\lambda^{(1)}$	[1.18,1.49]	1	[1.26,1.67]	[1.27,1.69]
$\lambda^{(2)}$	[-0.37,1.20]	1	[0.50,1.58]	1
$\lambda^{(3)}$	[0.98,1.25]	1	[1.45,2.05]	[1.44,2.05]
$\lambda^{(4)}$	[-0.21,2.20]	1	[0.60,1.43]	1
$\lambda^{(\gamma)}$	[0.76,1.17]	1	[0.60,1.02]	[0.60,1.02]

Table 4: 95% confidence intervals for scaling factors in Model 6 and Model 8 fitted to male and female SAPS data

modelling approach by using an alternative method for determining suitable restrictions of the relative model such as a “general-to-specific” approach described in Campos et al. (2005). This would fit an unrestricted model (i.e., Model 6) to the data, observe the confidence intervals for each parameter and use these to determine which restrictions to apply. To illustrate, if the confidence interval for $\lambda^{(j)}$ included unity, the general-to-specific approach would impose $\lambda^{(j)} = 1$ on the grounds of statistical significance. From Table 4, we see that the confidence intervals for $\lambda^{(j)}, j \neq 1$, for men and $\lambda^{(2)}$ and $\lambda^{(4)}$ for women contain unity. Therefore, with the exception of $\lambda^{(1)}$ for men, the general-to-specific approach would arrive at the same set of restrictions for the preferred model as our approach, which is based solely on considering the goodness of fit of the relative model with different sets of restrictions.

We also see, by comparing the confidence intervals for the unrestricted parameters in Model 8 with their counterparts from Model 6, that imposing the preferred set of restrictions does not significantly affect the estimation of the other parameters in the model. This, again, acts as a useful check to ensure that the procedure we have used to select the preferred set of restrictions does not remove statistically significant parameters from the relative model, and gives us confidence that our approach merely removes unnecessary parameters and so leads to a more parsimonious model.

Inspection of the boxplots of the bootstrapped parameters from Model 6, shown in Figure 6, indicates that the confidence intervals appear roughly symmetric around their midpoints. However, on closer inspection, $\lambda^{(1)}$ shows substantial skewness. Investigating this further, Jarque-Bera tests on the


bootstrapped rejects the assumption of normality for $\lambda^{(1)}$ for both sexes and for $\lambda^{(3)}$ for women at the 5% level. This indicates that we cannot reliably use asymptotic methods based on the information matrix (similar to those used in Brouhns et al. (2002)) to allow for parameter uncertainty, since these methods assume that the parameters will be normally distributed. This justifies the use of residual bootstrapping procedures, such as the one proposed here, in order to properly investigate parameter uncertainty in these models.

4.3.2 Model risk


The second stage of testing the robustness of the model is to fit the relative model to the bootstrapped data without specifying the form of the preferred model. Instead, we allow the procedure to select a potentially different preferred model in each simulation. This allows for “model risk”, in the sense of Cairns (2000), i.e., the risk that the model selected is not an accurate representation of the true processes generating the data. This process is illustrated in Figure 7, and is conceptually similar to the approach developed in Yang et al. (2015). However, we are still selecting a preferred model from a relatively limited set of comparators, and so the procedure does not fully capture the potential for model risk.

Looking first at the preferred form of $\alpha_x^{(\Delta)}$, we find that, from 1,000 bootstrapped datasets, the preferred model restricts $\alpha_x^{(\Delta)}$ to have a parametric form in 36% of the datasets for men and 100% of the datasets for women. The modelling approach, therefore, overwhelmingly prefers imposing a parametric structure for $\alpha_x^{(\Delta)}$ for women over allowing this to vary freely, even when allowing for model risk. For men, in contrast, the majority of bootstrapped simulations favour a non-parametric $\alpha_x^{(\Delta)}$. This may relate to the fact that, for men, the preferred model tends to have the scaling factors restricted to unity, allowing the model to devote more degrees of freedom to optimising the fit across ages by using a non-parametric $\alpha_x^{(\Delta)}$.

Table 5 shows the frequency of observing the various restrictions on the scaling factors in the preferred model, based on the same 1,000 bootstrapped datasets. We note that the most likely form that these restrictions take is the one preferred for Model 8 in Tables 2 and 3. The exceptions to this are $\lambda^{(\gamma)}$ for men, where an unrestricted value is preferred, and $\lambda^{(3)}$ for women, where a restriction equal unity is preferred in the case where model risk is allowed



(a) Men



(b) Women

Figure 6: Boxplots of the bootstrapped parameters from Model 6




Figure 7: Flow chart illustrating the procedure for fitting and selecting the relative model allowing for parameter uncertainty and model risk

Restriction placed on:		$\lambda^{(j)} = 0$	$\lambda^{(j)} = 1$	$\lambda^{(j)}$ unrestricted
Men:	$\lambda^{(1)}$	0%	70%	30%
	$\lambda^{(2)}$	47%	53%	0%
	$\lambda^{(3)}$	0%	97%	3%
	$\lambda^{(4)}$	44%	55%	1%
	$\lambda^{(\gamma)}$	1%	37%	62%
Women:	$\lambda^{(1)}$	0%	45%	55%
	$\lambda^{(2)}$	2%	96%	2%
	$\lambda^{(3)}$	0%	71%	29%
	$\lambda^{(4)}$	70%	30%	0%
	$\lambda^{(\gamma)}$	0%	28%	72%

Table 5: Frequency with which different restrictions are placed upon the scaling factors in the preferred relative model, based on 1,000 bootstrapped datasets

for. We are unsure why this should be the case. However, we note that it is inevitable that some information in the original data will be lost due to the random resampling of the fitted residuals in the Koissi et al. (2006) approach. Therefore, it is likely that the preferred model for bootstrapped data will be simpler and have more restrictions placed upon it, as fewer parameters will be required to capture the reduced level of information in the bootstrapped data compared with the original data.

In summary, we find that there is substantial model risk for both sexes, and no one set of restrictions out of the available options is universally selected. This will be important when we project the model in Section 5. It should also, again, caution us against using overly complicated models for the SAPS populations, since there is substantial uncertainty not only in any parameter estimates found but also in the fundamental form of the model.

5 Basis risk and projecting mortality for the SAPS population

In Section 4, the relative model was applied to historical data for the SAPS population. Given projections of the reference population, we can also use

the relative model to map these into projections for the sub-population.

Many pension schemes are concerned that the mortality experience of the scheme in question will be substantially different to that of the national population. This is often and informally referred to as “basis risk”. This is important when assessing hedging strategies (for instance, in Li and Hardy (2011), Coughlan et al. (2011) and Cairns et al. (2013)) using financial instruments based on national mortality rates. More fundamentally, it is an important question when funding a pension scheme, since most standard projections for future mortality rates are based on analysing national populations (for instance, the CMI Mortality Projection Model in Continuous Mortality Investigation (2009) that is widely used in the UK).

Intuitively, basis risk can arise because of a difference in levels of mortality rates (e.g., the specific population exhibiting systematically higher or lower mortality rates than the reference population as a result of characteristics such as socio-economic status which will change only slowly) and a difference in trends in mortality rates (i.e., mortality rates evolving differently in the sub-population, for instance, due to preferential access to new medications) between the two populations. In order to be more precise in our analysis, we define the following:

- the basis: the difference in mortality rates between two populations;
- the level basis: the difference in the level of mortality rates across ages between two populations at a defined point in time;
- the trend basis: the difference in the evolution of mortality rates between two populations;
- level basis risk: the risk arising due to uncertainty in the level basis;
- trend basis risk: the risk arising due to uncertainty in the trend basis in future; and
- basis risk: the aggregate of level basis risk and trend basis risk.

To clarify these definitions, if we know that mortality rates are different in two populations, but we also know how they are different, then there is a basis, but no basis risk. For example, if we knew that population A had

mortality rates that were 5% higher across all ages than population B, but these improved 1% p.a. faster, then we could still construct portfolios using securities linked to mortality in population B to hedge mortality in population A perfectly. Hence, we would say that there is a basis, but no basis risk. Basis risk arises because we cannot measure the differences in level and trend across different populations perfectly, e.g., we might believe the level basis is 5% across all ages but this is subject to error (i.e., level basis risk) and the true value could lie between 4% and 6%. This distinction is not allowed for in most models of “basis risk” (for instance, Li and Hardy (2011) and Haberman et al. (2014)), but we believe our definitions allow for a clearer understanding and attribution of basis risk.

Similarly, we draw a distinction between differences (and uncertainty in the differences) in the level of mortality between two populations and the rates of change between them. This distinction is widely made in practice, where it is common to consider the base table and improvements in mortality rates separately when selecting mortality assumptions.

Level differences can be measured relatively easily using traditional actuarial methods which are well within the capabilities of modern scheme actuaries. Hence, level basis risk is not often a primary concern, albeit we believe that it may be understated in many situations (see Hunt and Blake (2015e)). In contrast, the difference in the evolution of mortality rate between populations is more difficult to measure reliably and, consequently, trend basis risk is of greater concern to many scheme actuaries.

In terms of the relative model of Equation 3, level basis can be thought of as relating to $\alpha_x^{(\Delta)}$ and trend basis to $\lambda^{(j)}$. Therefore, we note that if parameter uncertainty and model risk are not allowed for, our proposed approach will not allow for basis risk in the sub-population, since we have no uncertainty in the mortality rates in the sub-population, conditional on knowing mortality rates in the reference population. Parameter uncertainty alone is sufficient to introduce level basis risk, since this allows for uncertainty in $\alpha_x^{(\Delta)}$. However, in many models, $\lambda^{(j)}$ are restricted and hence there will still be no uncertainty in the trend basis in the reference population, allowing for parameter uncertainty alone. Hence, it is only appropriate to talk about “basis risk” in conjunction with our model when both parameter uncertainty and model risk are allowed for when making projections.

This trade-off is common to many multi-population mortality models designed to measure basis risk. More complicated models can allow for a more sophisticated analysis and quantification of basis risk than simpler models, but are more difficult to estimate and less robust when fitted to small datasets. Our approach has been specifically designed for situations where there is relatively little data over a short period range to make best use of sparse data. However, we acknowledge that this makes it less effective at modelling basis risk than other models. We discuss this trade-off further in Section 7.

In order to evaluate the potential impact of basis risk between the UK and SAPS populations, we first need to project mortality rates for the national population. However, it is important that our projections of mortality rates are “well-identified” in the sense of Hunt and Blake (2015a,b) in that they do not depend upon our chosen identifiability constraints. To project the reference population, we therefore adopt the techniques of Hunt and Blake (2015b) and use random walks with drift

$$\boldsymbol{\kappa}_t^{(R)} = \boldsymbol{\mu}^{(R)} \begin{pmatrix} 1 \\ t \end{pmatrix} + \boldsymbol{\kappa}_{t-1}^{(R)} + \boldsymbol{\epsilon}_t^{(R)} \quad (11)$$

where $\boldsymbol{\kappa}_t^{(R)} = \left(\kappa_t^{(R,1)}, \dots, \kappa_t^{(R,N)} \right)^\top$, $\boldsymbol{\mu}^{(R)}$ is a matrix of drift coefficients with respect to the period “trends”, $(1, t)^\top$, and $\boldsymbol{\epsilon}_t^{(R)}$ are normally distributed, contemporaneously correlated innovations. For the cohort parameters, we make projections using an AR(1) around “well-identified” drifts

$$\gamma_y^{(R)} - \beta^{(R)} X_y = \rho^{(R)} (\gamma_{y-1}^{(R)} - \beta^{(R)} X_{y-1}) + \varepsilon_y \quad (12)$$

where $\beta^{(R)}$ is a matrix of drift coefficients with respect to the cohort “trends”, X_y .¹⁹ The deterministic functions in X_y are chosen to ensure that the projections are “well-identified”, i.e., that the projected mortality rates for the

¹⁹We have used the same notation for the trends, X_y , in Equation 12 as was used for the additional functions of year of birth in the relative model in Equation 3. However, the reader should be aware of the slight difference in definition between these two contexts, namely that in Equation 12, $X_y = (1, (y - \bar{y}), ((y - \bar{y})^2 - \sigma_y))^\top$, whilst in Equation 3, $X_y = ((y - \bar{y}), ((y - \bar{y})^2 - \sigma_y))^\top$, i.e., X_y did not possess a constant.

reference population do not depend upon the identifiability constraints used when fitting the model. To achieve this in the context of the reference models developed in Section 4.1 and Appendix C, we have

$$X_y = (1, (y - \bar{y}), ((y - \bar{y})^2 - \sigma_y))^T$$

$$\beta^{(R)} = \left(\beta_0^{(R)}, \beta_1^{(R)}, \beta_2^{(R)} \right) \quad R = \{\text{UKm}, \text{UKf}\}$$

Any dependence between mortality rates for men and women is not relevant to the following discussion, where only the relationships between mortality rates in the reference and sub-populations for the same sex are investigated. Therefore, in these projections, we do not take into account any dependence between male and female mortality rates in the reference population, and consequently project these populations independently. A more complete analysis of the mortality and longevity risks in pension schemes, such as in Hunt and Blake (2015e), would need to allow for dependence between sexes in the reference population. For techniques which could allow for dependence between these populations, see Hunt and Blake (2015c) and the references therein.

To illustrate the basis between the SAPS and UK populations, we consider annuity values at age 65 (calculated using a real discount rate of 1% p.a.). We perform 1,000 Monte Carlo simulations using the time series processes above to give projected mortality rates in the national population, which are then used to generate projected mortality rates in the SAPS population using the relative mortality models for men and women separately. Basis risk is accounted for using the relative model in three stages:

1. First, we allow only for the impact of the random innovations, $\epsilon_t^{(R)}$ and $\varepsilon_y^{(R)}$, on projected mortality rates, i.e., we allow for process risk in the terminology of Cairns (2000). We do this by using Equations 11 and 12 to project stochastically the period and cohort parameters found for the reference population in Section 4.1, and then using the preferred relative model estimated in Section 4.2 and shown as Model 8 in Tables 2 and 3. Note that this approach only allows for the basis between the two populations, and not for basis risk as defined above. Using this technique, we find correlations between annuity values in the UK and SAPS populations of over 99% for both men women.

2. Second, we allow for parameter uncertainty in both populations. To do this, we use the approach illustrated in Figure 4 to generate new parameters for both the reference and the sub-populations. The time series processes in Equations 11 and 12 are then re-estimated for the bootstrapped period and cohort parameters for the reference model, and mortality rates for the reference and sub-populations projected from these. By allowing for uncertainty in the parameters governing the basis between the two populations, this approach allows for level basis risk for both populations and trend basis risk for the female SAPS projections. Trend basis risk, however, is still not allowed for in the male SAPS projections since the $\lambda^{(j)}$ for this population are restricted to be equal to unity, and so are not subject to parameter uncertainty.

When allowing for parameter uncertainty, we find correlations between annuity values in the UK and SAPS populations of around 98% for both men and women. This indicates that parameter uncertainty alone has not added significantly to the basis risk between the two populations. This is surprising, given the results of Section 4.3.1 as shown in Figures 5 and 6, which showed relatively high levels of uncertainty in the levels and scaling parameters. However, this may indicate that the basis risk arising from different rates of change in mortality in different populations may not be particularly significant, as discussed in Section 7.

3. Finally, we allow for model risk in the selection of the preferred model for the sub-population. We do this using the same procedure as illustrated in Figure 7 to generate new parameters for the reference population and a new preferred model for the sub-population. The time series processes in Equations 11 and 12 are then re-estimated for the bootstrapped period and cohort parameters for the reference model, and mortality rates for the reference and sub-populations projected from these. Because all parameters in the model are subject to uncertainty using this method (i.e., even the restrictions that were previously found are reassessed), this approach allows for both level and trend basis risk in both populations. Using this procedure, we observe correlations between annuity values in the UK and SAPS populations of 85% for men and 93% for women. Thus it is the potential for model mis-specification which adds most significantly to the basis risk for both populations.




Figure 8: Projected annuity values for the UK and SAPS populations from 1,000 Monte Carlo simulations

Note that this analysis looks only at annuity values (i.e., the expected present value of payments to an individual) and so does not consider the idiosyncratic risk that would also be present in the benefits payable from a pension scheme. This was investigated in Donnelly (2014), Aro (2014) and, in particular, in Hunt and Blake (2015e) where we find this is likely to be substantial for even relatively large pension schemes.

Figure 8 shows scatter plots of annuity values calculated using mortality rates in the UK and SAPS populations for men and women in the third, most general case (i.e., incorporating process risk, parameter uncertainty and model risk). First, we note that, for both sexes, the systematic longevity risk (indicated by the range of values the annuity value can take, e.g., 18 to 24 in the case of men) is far greater than the basis risk. Indeed, the systematic longevity risk accounts for around 90% of the uncertainty in an annuity value for the SAPS population using ordinary linear regression, indicating that basis risk may be considerably less important than is commonly believed. This is discussed further in Section 7.

Second, Figure 8 shows that, for both sexes, the points tend to cluster depending on the preferred set of restrictions found. Studies which do not allow for potential model risk will, therefore, only observe one of these clusters and hence understate the true potential for basis risk.

However, it is important to note that even when model risk is allowed for, there is limited trend basis risk between the two populations. This is because the same processes, i.e., $\kappa_t^{(R)}$ and $\gamma_y^{(R)}$, control the evolution of mortality in both populations, albeit scaled by factors, $\lambda^{(j)}$, in the sub-population which are uncertain. This is in contrast with other studies, such as Hunt and Blake (2015c), which have allowed for different time series processes in each population. This helps explain why the correlations we find are somewhat higher than those found in other studies of basis risk, such as Cairns et al. (2013). However, we note that most of these studies used sub-populations which were considerably larger and covered a longer period of time than the SAPS population. Consequently, there is a trade-off. On the one hand, we might wish to use more complicated models that might give a more accurate assessment of basis risk, but which require larger volumes of data to estimate robustly and, therefore, might involve using data for a larger sub-population which is less relevant for the mortality experience of a specific pension scheme (for instance, the CMI Assured Lives dataset). On the other hand, we might prefer to use simpler models, which can be robustly estimated from smaller datasets that are likely to be more relevant to the specific scheme experience, but give a less accurate assessment of basis risk. The impact of this trade-off is discussed in Section 7.

Finally, the importance of model risk and parameter uncertainty will tend to increase if we consider populations smaller than the SAPS population, as we do in Section 6. We would therefore expect to see correlations of a similar size to those found in other studies for population sizes that are more typical of UK pension schemes, due to the greater parameter uncertainty and model risk, even without allowing for different period and cohort processes in the two populations. In addition, the cashflows experienced by a pension scheme will also have (potentially substantial) idiosyncratic risk due to the relatively low number of lives under observation. This suggests that, for most pension-scheme-sized populations, it is impossible to distinguish between the trend basis risk arising from different processes in each population and the basis risk arising from a model such as ours where the two processes are the same, but we include parameter and model uncertainty. This is discussed further in Section 7 and Hunt and Blake (2015e).

6 Applying the relative model to small populations

While the SAPS population is small compared with the national UK population, it does have annual exposures to risk of over one million lives each for men and women, and so still represents a population larger than almost all occupational pension schemes (with the exception of some state schemes). However, the methods developed in this paper can be applied to significantly smaller populations, such as those more comparable with the size of large occupational pension schemes.

As discussed in Section 4.2, the relative model applied to the SAPS population exhibited a strong preference for parsimony. However, parameter uncertainty and model risk were still important considerations, even with a relatively simple model and the full SAPS data. It is therefore exceedingly likely that in even smaller populations, these considerations will dominate what we can and cannot realistically say about the evolution of mortality of a small sub-population such as that associated with an individual pension scheme.

We investigate the effect of population size on the ability of the relative model to measure mortality differences with the national population by randomly generating scheme-sized exposures to risk and death counts (denoted by lower-case s) based on the SAPS data. We adopt the following procedure to generate pseudo-data for a scheme with N lives (considering each sex separately):

1. We first rescale SAPS exposures, $E_{x,t}^{(S)}$, to give a proxy for smaller pension schemes with approximately N members. We could, in principle, do this very simply by setting

$$E_{x,t}^{(s)} = E_{x,t}^{(S)} \times \frac{N}{\sum_{\xi} E_{\xi,t}^{(S)}}$$

This would give a scheme with a constant total exposure to risk (N) over each year, but the same pattern of exposures to risk as the SAPS population across different ages. However, this simple approach does not capture the pattern of exposures across years seen in the actual

SAPS data, due to the partial submission of scheme data in the first and last few years of the SAPS datasets (discussed in Section 2, see also Figure 13a). This means that, were we to artificially generate a scheme of the same size as the SAPS population, we would not recover the observed SAPS exposures and so would obtain inconsistent results. Since we will apply this procedure to generate pseudo-scheme data for schemes of widely varying sizes, up to and in excess of the full SAPS data, it is essential that our results are consistent with the results we found in previous sections. Consequently, we amend the scaling factors so that

$$E_{x,t}^{(s)} = E_{x,t}^{(S)} \times \frac{5N}{\sum_{\xi} \sum_{\tau=2004}^{2008} E_{\xi,\tau}^{(S)}}$$

This modifies the denominator to reflect the average exposure to risk in the SAPS data in years 2004-2008, for which almost all relevant pension schemes have submitted data to the SAPS study. This approach replicates the full SAPS data when we generate a scheme of the same size as the SAPS population (including the pattern of relatively low exposures to risk for the first and last years, along with the pattern of exposures at different ages found in the SAPS data).

2. We then generate random death counts for the scheme by modelling them as Poisson random variables. To do this, we use the exposures to risk generated using both the procedure above and the crude mortality rates observed in the SAPS dataset,

$$D_{x,t}^{(s)} \sim Po \left(\frac{D_{x,t}^{(S)}}{E_{x,t}^{(S)}} E_{x,t}^{(s)} \right)$$

We fit the relative model to this pseudo-scheme data, testing all 486 sets of possible restrictions on the parameters to determine the preferred model using the same procedure described in Section 4.3.2. This procedure is illustrated in Figure 9. Such an approach is conceptually similar to the “semi-parametric” bootstrapping technique in Brouhns et al. (2005), except we rescale the exposures in order to simulate the range of different scheme sizes present in the UK.




Figure 9: Flow chart illustrating the procedure for generating data and fitting the relative model to scheme-sized populations, allowing for parameter uncertainty and model risk

To gain a better understanding of the impact of the size of the population on the complexity of the preferred model, we apply this procedure for scheme sizes at regular intervals in the range $N \in (10^2, 10^6)$ and for 1,000 sets of random death counts at each scheme size. This range of population sizes covers almost the entire range of pension scheme sizes in the UK, and the fitting of multiple models allows for potential model risk in the selection of the preferred model. The results of this procedure for men and women are shown in Figures 10 and 11.

First, let us consider the results shown in Figures 10a and 11a. These figures show that the probability of the procedure preferring a parametric restriction for $\alpha^{(\Delta)}$ is almost unity for schemes with up to around half a million male members and approximately one million female members. This indicates an overwhelming preference for parametric restrictions for $\alpha_x^{(\Delta)}$ in all but the very largest schemes with memberships far in excess of all but the largest state schemes in the UK. The implication of this is that making simple adjustments to a standard mortality table will be sufficient to capture the difference in levels in mortality for almost all UK schemes, with little or no need to graduate a bespoke table (even if the data is available).

Looking at the scaling factors for the age/period and cohort terms, we see that, typically, the smallest schemes (fewer than 1,000 members of each sex) are indifferent between restricting $\lambda^{(j)}$ to be equal to zero or unity. For instance, Figure 10b shows that the procedure imposes the restriction $\lambda^{(1)} = 0$ and $\lambda^{(1)} = 1$ for men in approximately 50% of the simulations for small schemes, with $\lambda^{(1)}$ being estimated without restrictions in almost no cases. This pattern is repeated for the other scaling factors shown in Figures 10 and 11. Since the restrictions $\lambda^{(j)} = 0$ and $\lambda^{(j)} = 1$ give models with the same number of free parameters, the choice between them depends entirely on the log-likelihood found when fitting the model. However, the difference between $\lambda^{(j)} = 0$ and $\lambda^{(j)} = 1$ is the difference between a model which allows mortality rates to change with time and a static model of mortality ($\lambda^{(j)} = 0 \forall j$). We therefore find that, in very small schemes it is almost impossible to say whether or not mortality rates are changing, let alone whether the rate of change differs from the national population.

Looking at Figure 10b again, we see that for larger schemes, with around 10,000 to 100,000 members, the relative model has a clear preference for set-




Figure 10: Restrictions placed on the relative model for different volumes of male data







(a) $\alpha_x^{(\Delta)}$ (b) $\kappa_t^{(1)}$ (c) $\kappa_t^{(2)}$ (d) $\kappa_t^{(3)}$ (e) $\kappa_t^{(4)}$ (f) γ_y

Figure 11: Restrictions placed on the relative model for different volumes of female data

ting $\lambda^{(1)} = 1$ for men, which is preferred in almost all simulations for schemes with around 200,000 members. This pattern also holds for $\lambda^{(3)}$ for men, and $\lambda^{(1)}$, $\lambda^{(3)}$ and (to a lesser extent) $\lambda^{(4)}$ for women. The cohort scaling factors also exhibit similar behaviour, although there is a greater preference for restricting $\lambda^{(\gamma)}$ to equal zero for small schemes and it is only for schemes with around 500,000 men or women that there is a significant preference for setting $\lambda^{(\gamma)} = 1$. For the other scaling factors ($\lambda^{(2)}$ for both sexes and $\lambda^{(4)}$ for men), the model is broadly indifferent between imposing $\lambda^{(j)} = 0$ and $\lambda^{(j)} = 1$ for all sizes and schemes.

The implication of this is that, although there is sufficient evidence to suggest mortality is improving in these larger schemes (unlike the smaller schemes discussed above), there is not enough data to quantify any differences in this improvement between the scheme and the national population. This supports the use of projection methods based on the national population for the majority of pension schemes in the UK. It also makes it unlikely that we can detect trend basis risk between the scheme and the national population for schemes with fewer than 100,000 members of each sex. It also shows that there is insufficient evidence to justify the inclusion of a set of cohort parameters for all but the largest pension schemes, a result which agrees with the findings of Haberman et al. (2014).

The preference for a freely varying $\lambda^{(1)}$ for some scaling factors for men in schemes with around two million members in Figure 10b illustrates that it is only in the very largest schemes, with over one million members of each sex, do we find that there is sufficient data to estimate unrestricted $\lambda^{(j)}$. Therefore, it is only for these very large schemes that we can quantify any difference in the evolution of mortality rates between a pension scheme and the national population. However, the results of Section 5 indicate that, even when trend basis is allowed for, the impact on annuity values is likely to be quite limited, especially when considered in the context of the other mortality and longevity risks in the scheme. This is investigated further in Hunt and Blake (2015e).

In summary, we find that, for datasets that are the same size as a typical UK pension scheme, there is insufficient data to make more than a few simple adjustments to reflect the level basis. For most practical circumstances, we would therefore be unable to quantify any trend basis in a pension scheme,

and it is most convenient to assume that the changes in mortality in the scheme are equal to those in the national population. Therefore, for small schemes, we find that the basis risk is determined solely by the uncertainty in estimating the level basis, rather than the trend basis. This is examined further in Hunt and Blake (2015e). Given that trend basis risk is often given as a key concern for why pension schemes are reluctant to use index based hedging instruments to manage their longevity risk and, instead, prefer bespoke arrangements, we believe that much of this trepidation is misplaced, as we now discuss.

7 Discussion: Basis risk in pension schemes

There has been a lot of work regarding the quantification of basis risk between different populations, most notably in Plat (2009b), Salhi and Loisel (2009), Li and Hardy (2011), Coughlan et al. (2011), Cairns et al. (2013), Li et al. (2015) and Haberman et al. (2014). The analysis of this risk has also motivated many of the multi-population mortality models that have recently been proposed, such as those of Dowd et al. (2011), Cairns et al. (2011a), Zhou et al. (2014), Villegas and Haberman (2014) and Hunt and Blake (2015c). However, much of this work to date is not directly relevant to the situation faced by many UK pension schemes when assessing and trying to manage their longevity risk.

Partly, this is because the populations being considered in these studies are far larger in terms of the size of the exposures to risk than that of a typical (or, indeed, even a very large) UK pension scheme. This enables the authors of these studies to adopt a “general-to-specific” approach when analysing trend basis risk: first mortality models are fitted separately to the different populations under investigation and then any dependence between the period or cohort parameters is analysed. This approach is exemplified by the study of Li et al. (2015), which statistically determined whether or not to simplify a model by using the same sets of parameters for different populations (which is a very specific form of dependence). Such an approach therefore starts from the assumption that mortality rates will have different patterns of evolution in different populations, and then looks for evidence of similarities.

Such an approach is entirely reasonable when looking at large populations where there is sufficient data to estimate sophisticated mortality models in each population under investigation. However, this is not the situation in which most pension schemes find themselves. Instead, with relatively little data, it is necessary for them to adopt a “specific-to-general” approach, such as that underlying the relative model proposed in this paper. As there is insufficient data to estimate many sub-population-specific parameters robustly, a specific-to-general methodology starts from the assumption that mortality rates in the sub-population evolve in the same fashion as those in the reference population and then looks for evidence of differences between the two. This approach naturally leads to more parsimonious models, which are therefore likely to be more robust. However, it is less likely to overturn the null hypothesis of no trend basis risk, especially when parameter uncertainty and model risk are included in any analysis. This is the trade-off between the ability to model basis risk fully and the simplicity and robustness of the model for small datasets discussed in Section 5.

Our findings suggest that large volumes of data (in terms of both the size of the exposures to risk and the period range of the data) are required to overturn the null hypothesis of no trend basis risk, especially when parameter uncertainty and model risk are included in the analysis. For the full SAPS dataset, the simple relative model we have proposed achieves relatively good and parsimonious fits to the data for both men and women, as shown in Section 4. Furthermore, for the smaller datasets more typical of UK pension schemes, even simpler models which fix the scaling factors in the model are preferred, as shown in Section 6. This is consistent with the results of Haberman et al. (2014), which found that it is only possible to quantify basis risk for very large schemes.

In addition, in order to estimate the more complicated multivariate time series processes used in many of the general-to-specific models we need longer periods of data than a typical pension scheme has. For instance, to estimate the cointegration-based models of Salhi and Loisel (2009) and Hunt and Blake (2015c) requires several decades of mortality data, which is usually far in excess of what a pension scheme will have itself. Similarly, Haberman et al. (2014) found that eight years or more of data is required for the quantification of basis risk, even for very large pension schemes. Specific-to-general models, however, do not require such long data ranges, as they start from

the assumption that information about the reference population can be used to fill in gaps in the data if required.

However, Section 5 shows that projections from the relative model have many of the features we would expect from models which use more complicated time series processes, when appropriate allowance is made for parameter uncertainty and model risk, despite there being no genuine trend basis risk using the relative approach. This implies that it may be impossible to distinguish between genuine trend basis risk and the effects of parameter uncertainty and model risk in practice. Indeed, it is noticeable that few of the studies to date which have investigated basis risk allow for parameter uncertainty and model risk, and so the findings of these studies potentially wrongly attribute differences in historical improvements in mortality between different populations to basis risk and, thus, overstate its importance.

Finally, we note that the confusion between basis and basis risk, and the distinction between the level and trend bases, may cause issues with some models. For instance, many models proposed for “basis risk”, e.g., Jarner and Kryger (2011), are actually models of the basis according to our definition. Furthermore, models which allow for trend basis risk using different processes in each population often do not allow for level basis risk by ignoring parameter uncertainty, e.g., Zhou et al. (2014), and so may understate its importance in smaller populations. We therefore believe that it is important to make these distinctions to ensure that all users of multi-population mortality models are able to communicate effectively about the advantages and disadvantages of the different approaches.

We find that for most UK pension schemes, the existence or not of trend basis between the scheme and the UK population is of little practical relevance. The scheme will never have sufficient information to be able to say with confidence that the improvements in mortality it experiences are significantly different from that in the reference population, as any such differences will be overwhelmed by the other sources of risk and uncertainty present in the scheme.

This is not to dispute that trend risk can exist between different countries or amongst highly distinct sub-populations of a reference population. Indeed, there are good reasons to suggest that it does and that there is sufficient data

to estimate it reliably using a general-to-specific approach as in previous studies. For instance, many studies (for instance in Li and Hardy (2011) and Hunt and Blake (2015c)) investigate differences between the evolution of mortality rates in different countries. However, populations in different countries may have different diets, lifestyles and access to healthcare, and so would be expected to have different patterns of evolution in mortality rates. Other studies, such as in Villegas and Haberman (2014) consider the differences in the evolution of mortality rates between highly selective sub-populations of a country (for instance, based on deprivation). The sub-populations in these studies have, therefore, been constructed in such a fashion as to maximise the likelihood of observing different patterns in the evolution of mortality rates.²⁰

Nor do we argue that the evolution of mortality rates in a pension scheme *is* the same as in the reference population. It may be true that for very large schemes, we may have sufficient data to be able to detect trend basis (even when allowing for parameter uncertainty and model risk) if there is quite a large difference in the evolution of mortality rates between the two populations.

However, a pension scheme, whose only membership requirement was employment with a particular company, would be expected to be more similar to the national population or differ only due to persistent selection effects which affect the level of mortality rates (i.e., level basis) but not how mortality rates evolve with time (i.e., trend basis). In order to have sufficient data to reject the assumption that the evolution of mortality rates in the pension scheme is the same as in the national population, the scheme must be very large (such as being the pension scheme for a large and long-established national company) and so entry to such schemes is likely to be relatively unselective. Therefore, these schemes are more likely to represent a fair cross section of the UK population. Consequently, the circumstances where we

²⁰As well as being a highly selected sub-population of the UK population, the data for CMI Assured Lives has also varied considerably in the socio-economic makeup of the relevant population over the period of the data due to changes in the UK annuity market. Since this dataset was used in Cairns et al. (2011a), Dowd et al. (2011) and Cairns et al. (2013), it is therefore unclear whether any difference in the evolution of mortality detected by these studies is the result of genuine trend basis risk or simply a result of the changing composition of the dataset.

have enough data to quantify trend basis (for example, the pension scheme of a large, national employer) are also the circumstances when trend basis is least likely to be important. In most practical situations, we will never have sufficient data to quantify any trend basis and therefore an assumption of no difference between the evolution of mortality rates in the national population and the pension scheme is both practical and parsimonious. However, by assuming no trend basis a priori also implies assuming no trend basis risk.

The practical implications of these results are important for the development of any market in longevity hedging. Since trend basis risk is unlikely to be important enough to be statistically significant, it is also unlikely to be financially significant. If longevity risk is felt to be important, hedging can be achieved by use of standardised instruments based on projected changes in mortality rates in a reference population, making adjustments to reflect the level of mortality observed in the pension scheme. Concerns that the trend basis risk will make such hedges ineffective, such as those raised against the EIB longevity bond (see Blake et al. (2006)), should be regarded as secondary compared with the other risks a pension scheme faces, such as idiosyncratic mortality risk. Bespoke products, such as longevity swaps tailored to the characteristics of the pension scheme, should be regarded primarily as vehicles for hedging and transferring these other risks, rather than any trend basis risk for the scheme, and their cost effectiveness judged accordingly, as discussed in Hunt and Blake (2015e).

8 Conclusions

In conclusion, in this study we present a relative model for mortality in a sub-population, which models the mortality rates observed in a small population relative to those observed in a larger reference population. Such a model has the advantages of being more parsimonious compared with the approach of fitting separate mortality models for both populations, which has been adopted in many multi-population mortality studies, and so is better suited to situations where there is little data for the sub-population.

We then apply the relative model to investigate the mortality rates observed in the SAPS study of UK pension schemes. We find that this simple

model is sufficient to achieve a good and parsimonious fit to the available data and reasonable projections of mortality rates. Specifically, we find that, in aggregate, members of UK occupational pension schemes generally experience lower levels of mortality rates than the national population, which are also improving at a faster rate than those in the national population. However, we find relatively high levels of uncertainty in estimating the parameters even in this simple model and that the data is insufficient to uniformly prefer one model over any other. Furthermore, when we apply the relative modelling approach to sub-populations which are smaller than the SAPS population, and closer in size to those of typical UK pension schemes, we find that the modelling approach prefers very simple, highly restricted models, which do not allow for any difference in the evolution of mortality between the reference and sub-populations.

In order to analyse how mortality rates differ between populations more completely, we introduce a new set of definitions for basis risk. These definitions seek to distinguish between differences in the level of mortality rates between populations and differences in their rates of change, and also to restrict discussion of basis risk to a discussion of uncertainty in these differences. We feel that this allows for a more complete discussion of what different models can, and cannot, say about basis risk.

These considerations lead us to the belief that a full analysis of trend basis risk is not possible with the datasets realistically available for most pension schemes. This is because such an analysis would require more sophisticated models than the relative model proposed, with separate processes operating in each population. We find that, in pension-scheme-sized datasets, we will never have sufficient information to determine whether there is any difference in the evolution of mortality rates in the sub-population compared with the reference population when the other risks present are properly accounted for. Therefore, we believe that an assumption of no difference in the evolution of mortality rates between the two populations is practical and parsimonious. Consequently, we conclude that concerns regarding trend basis risk in the development of the market for longevity hedging and risk management tools for pension schemes are misplaced.

A Summary of SAPS data

We are indebted to the CMI for kindly providing death counts and exposures, weighted by individual lives, for the SAPS population for the period 2000 to 2011 and ages 60 to 90. These relate to all pensioners in the surveyed pension schemes, and so include people receiving benefits after retiring at normal retirement age, those who retired early or in ill-health, and those in receipt of spousal benefits. It is likely that some of these sub-populations will have different mortality characteristics, especially those retiring in ill-health. However, such cases represent a relatively small proportion of the SAPS data and are unlikely to materially impact our results.

Large pension schemes in the UK submit their mortality experience to the SAPS study following completion of a triennial funding valuation. Therefore, each submission is in respect of data with a considerable time delay, e.g., data submitted on 30 June 2013 may result from a funding valuation with an effective date of 31 December 2011 (due to the time taken to perform the valuation) and cover the period 1 January 2009 to 31 December 2011. Consequently, the last few years of the SAPS data only reflects a partial submission to date of the mortality experience of the schemes which will, ultimately, submit data to the study. However, we have no reason to believe that the schemes that have submitted to date are an unrepresentative subsample of the SAPS population, and so do not believe this biases our results.

Similarly, there are fewer submissions for the earliest years of the SAPS data. Unlike the most recent years, the missing data for this period will never be received by the CMI. Therefore, we only have data we consider complete for roughly the period 2004 to 2008.²¹

Figures 12 and 13 summarise the patterns of deaths and exposures for

²¹However, we note that Continuous Mortality Investigation (2014b) and Continuous Mortality Investigation (2014c) have been published subsequently to us obtaining the data used in this study from the CMI. These working papers included new data in respect of the SAPS study for 2012 and 2013, respectively, along with revisions to the data for years prior to 2012 caused by new pension schemes submitting data to the study. In the interests of avoiding errors caused by merging multiple sources of data, we have not combined this new data with that provided previously by the CMI and, therefore, it has not been included in this study. However, we have investigated the impact the new data would have on our findings if it were included, and are satisfied that it would not affect our results materially.




Figure 12: Exposures to risk and death counts in the SAPS dataset by age

men and women across age and time.

B Identifiability in the relative model

In Hunt and Blake (2015a) and Hunt and Blake (2015b), we discussed the identifiability issues in age/period and age/period/cohort mortality models, respectively. In particular, we find that almost all APC mortality models possess “invariant” transformations, i.e., transformations of the parameters of the model which leave the fitted mortality rates unchanged. In order to find a unique set of parameters, we impose a set of identifiability constraints on them. Typically, these are chosen so that we can assign our desired interpretation of the demographic significance to the parameters in question. However, because this interpretation is subjective, it is important that our choice of identifiability constraints does not have any impact on any observable quantities. For instance, we discuss in Hunt and Blake (2015a,b) how to ensure that projected mortality rates are independent of the choice of identifiability constraints.

The relative model in Equation 3 does not possess any additional identifiability issues in and of itself, once the parameters from the reference population are known. However, due to the relative structure, transformations of the parameters in the reference population model will have knock-on effects for those in the relative model. It is important therefore that invariant trans-




Figure 13: Exposures to risk and death counts in the SAPS dataset by year

formations of the reference model are also invariant for the relative model, so that our choice of identifiability constraints for the reference population does not affect the suitability of the relative model. This requirement will determine both the nature of the set of deterministic functions of year of birth, X_y in Equation 3, and the nature of any parametric simplification imposed upon $\alpha_x^{(\Delta)}$, i.e., if $\alpha_x^{(\Delta)}$ is restricted to be a linear combination of a set of basis functions

$$\alpha_x^{(\Delta)} = \sum_{i=1}^n \alpha^{(i)} g^{(i)}(x)$$

then the nature of the basis functions, $g^{(i)}(x)$, will be determined by the identifiability issues present in the model. We, therefore, consider each of the different forms that the invariant transformations of the reference model can take in turn, in order to ensure that they will not affect the relative model.

First, the scaling factors in the relative model do not depend upon the normalisation scheme of the age/period terms in the reference model. Normalisation schemes are imposed by using a transformation of the form

$$\{\hat{f}^{(R,i)}(x), \hat{\kappa}_t^{(R,i)}\} = \left\{ \frac{1}{a^{(i)}} f^{(R,i)}(x), a^{(i)} \kappa_t^{(R,I)} \right\}$$

and so it is obvious that $\Lambda^{(i)} \hat{f}^{(R,i)}(x) \hat{\kappa}_t^{(R,i)} = \Lambda^{(i)} f^{(R,i)}(x) \kappa_t^{(R,i)}$.

Second, we know from Hunt and Blake (2015a) that all APC models are invariant under the transformation

$$\{\hat{\alpha}_x^{(R)}, \hat{f}^{(R,i)}(x), \kappa_t^{(R,i)}, \hat{\gamma}_y^{(R)}\} = \{\alpha_x^{(R)} - a^{(i)} f^{(R,i)}(x), f^{(i)}(x), \kappa_t^{(R,i)} + a^{(i)}, \gamma_y^{(R)}\} \quad (13)$$

i.e., the model using the transformed parameter set gives exactly the same fitted mortality rates. This allows us to impose the “level” of the period functions, $\kappa_t^{(R,i)}$, via the identifiability constraints, such as imposing $\sum_t \kappa_t^{(R,i)} = 0$ or $\kappa_T^{(R,i)} = 0$. However, such a set of identifiability constraints is arbitrary, and so should not have any consequences for our relative modelling approach.

Accordingly, we require that our relative model in Equation 3 is also invariant if the transformed parameters are used for the reference population. In order to ensure this, we require that Equation 3 is invariant under the transformation

$$\hat{\alpha}_x^{(\Delta)} = \alpha_x^{(\Delta)} - a^{(i)} \Lambda^{(i)} f^{(R,i)}(x) \quad (14)$$

This transformation can be accommodated without $\alpha_x^{(\Delta)}$ fundamentally changing form if

1. $\alpha_x^{(\Delta)}$ is non-parametric, as in the original specification in Equation 3; or
2. if $\alpha_x^{(\Delta)}$ is restricted to be of parametric form, then $\alpha_x^{(\Delta)} = \sum_{i=1}^N \alpha^{(i)} f^{(i)}(x) + \sum_{i=N+1}^n \alpha^{(i)} g^{(i)}(x)$, i.e., the age functions in the reference model form a subset of the basis functions, $g^{(i)}(x)$.

As an example, consider the case where our model for the reference population is the “classic APC” model of Hobcraft et al. (1982)

$$\begin{aligned} \ln(\mu_{x,t}^{(R)}) &= \alpha_x^{(R)} + \kappa_t^{(R)} + \gamma_{t-x}^{(R)} \\ R_{x,t} &= \alpha_x^{(\Delta)} + \Lambda^{(1)} \kappa_t^{(R)} + \Lambda^{(\gamma)} \gamma_{t-x}^{(R)} + \nu X_{t-x} \end{aligned}$$

The classic APC model is invariant under the transformation

$$\{\hat{\alpha}_x^{(R)}, \hat{\kappa}_t^{(R)}, \hat{\gamma}_y^{(R)}\} = \{\alpha_x^{(R)} - a, \kappa_t^{(R)} + a, \gamma_y^{(R)}\}$$

Substituting the transformed parameters into the relative model gives

$$\begin{aligned}\hat{R}_{x,t} &= \hat{\alpha}_x^{(\Delta)} + \hat{\Lambda}^{(1)}\hat{\kappa}_t^{(R)} + \hat{\Lambda}^{(\gamma)}\hat{\gamma}_{t-x}^{(R)} + \hat{\nu}X_{t-x} \\ &= \hat{\alpha}_x^{(\Delta)} + \hat{\Lambda}^{(1)}(\kappa_t^{(R)} + a) + \hat{\Lambda}^{(\gamma)}\gamma_{t-x}^{(R)} + \hat{\nu}X_{t-x}\end{aligned}$$

In order to ensure $\hat{R}_{x,t} = R_{x,t}$, we must have $\hat{\Lambda}^{(1)} = \Lambda^{(1)}$, $\hat{\nu} = \nu$ and $\hat{\alpha}_x^{(\Delta)} = \alpha_x^{(\Delta)} - a\Lambda^{(1)}$. The requirement that $\hat{\alpha}_x^{(\Delta)}$ is of the same form as $\alpha_x^{(\Delta)}$ implies that any parametric simplification for $\alpha_x^{(\Delta)}$ must be of the form $\alpha_x^{(\Delta)} = \alpha^{(1)} + \sum_{j=2}^n \alpha^{(j)}g^{(j)}(x)$, i.e., it has a constant basis function, $g^{(1)}(x) = 1$, in order that the relative model does not change if the levels of the period functions are transformed.

Third, the values of $\Lambda^{(i)}$ depend upon the definition of the age functions in the reference model. “Equivalent” models for the reference population, which use different definitions for the age functions but give identical fitted mortality rates, will give different values of $\Lambda^{(i)}$. To see this, consider a reference model of the form²²

$$\begin{aligned}\ln\left(\mu_{x,t}^{(R)}\right) &= \alpha_x^{(R)} + \kappa_t^{(R,1)} + (x - \bar{x})\kappa_t^{(R,2)} + \gamma_{t-x}^{(R)} \\ R_{x,t} &= \alpha_x^{(\Delta)} + \Lambda^{(1)}\kappa_t^{(R,1)} + \Lambda^{(2)}(x - \bar{x})\kappa_t^{(R,2)} + \Lambda^{(\gamma)}\gamma_{t-x}^{(R)} + \nu X_{t-x}\end{aligned}$$

The model for the reference population is equivalent to a model of the form

$$\ln\left(\mu_{x,t}^{(R)}\right) = \alpha_x^{(R)} + \hat{\kappa}_t^{(R,1)} + x\hat{\kappa}_t^{(R,2)} + \gamma_{t-x}^{(R)}$$

with $\hat{\kappa}_t^{(R,1)} = \kappa_t^{(R,1)} - \bar{x}\kappa_t^{(R,2)}$ and $\hat{\kappa}_t^{(R,2)} = \kappa_t^{(R,2)}$. The corresponding relative model in this case is

$$\hat{R}_{x,t} = \hat{\alpha}_x^{(\Delta)} + \hat{\Lambda}^{(1)}\hat{\kappa}_t^{(R,1)} + \hat{\Lambda}^{(2)}x\hat{\kappa}_t^{(R,2)} + \Lambda^{(\gamma)}\gamma_{t-x}^{(R)} + \hat{\nu}X_{t-x}$$

However, in this situation, we find that we require $\hat{\Lambda}^{(2)}x = \Lambda^{(2)}(x - \bar{x}) + \Lambda^{(1)}\bar{x}$ to give the same fitted mortality rates for both reference models. In this case,

²²We call this model the “reduced Plat” model, since it was suggested in Plat (2009a) as being a reduced form of the model tested in that paper that might be more suitable for high ages. This model can also be thought of as an extension to model M6 in Cairns et al. (2009), with a static age function, or as an extension to the “CBDX” model discussed in Hunt and Blake (2015a) with a cohort term.

the relationship between the two is a function of age, x , which contradicts the assumption that the scaling factors are constants. Consequently, we find that the values of the scaling factors and the fit provided by the relative model will depend on the specifics of the age functions in the reference model and will differ between equivalent models.

Finally, identifiability under transformations of the cohort parameters is not as straightforward. From Hunt and Blake (2015b), we found that APC models may have unidentifiable trends which are allocated between the age/period and cohort terms by the identifiability constraints. Invariance of the mortality rates in the relative model to a different allocation of these trends in the reference model depends upon the deterministic regressors, X_y , we added to the relative model in Equation 3, and the form of any parametric simplification of $\alpha_x^{(\Delta)}$. This is illustrated by the following example.

Consider the example of the classic APC model for the reference population again. In addition to the transformation above, the classic APC model is also invariant under the following two transformations involving the cohort parameters

$$\begin{aligned}\{\hat{\alpha}_x^{(R)}, \hat{\kappa}_t^{(R)}, \hat{\gamma}_y^{(R)}\} &= \{\alpha_x^{(R)} - b, \kappa_t^{(R)}, \gamma_y^{(R)} + b\} \\ \{\hat{\alpha}_x^{(R)}, \hat{\kappa}_t^{(R)}, \hat{\gamma}_y^{(R)}\} &= \{\alpha_x^{(R)} + c(x - \bar{x}), \kappa_t^{(R)} - c(t - \bar{t}), \gamma_y^{(R)} + c(y - \bar{y})\}\end{aligned}$$

where $\bar{x}$ and $\bar{t}$ are defined in a similar fashion to $\bar{y}$ in Section 4.1. Invariance of the relative model under the first of these transformations requires $\hat{\Lambda}^{(\gamma)} = \Lambda^{(\gamma)}$ and $\hat{\alpha}_x^{(\Delta)} = \alpha_x^{(\Delta)} - b\Lambda^{(\gamma)}$, and therefore that any parametric restriction placed upon $\alpha_x^{(\Delta)}$ must have a constant basis function, $g^{(1)}(x) = 1$, as discussed above in respect of the level of $\kappa_t^{(R)}$.

However, substituting the transformed parameters from the second transformation in Equation 3, we find

$$\begin{aligned}\hat{R}_{x,t} &= \hat{\alpha}_x^{(\Delta)} + \hat{\Lambda}^{(1)}\hat{\kappa}_t^{(R)} + \hat{\Lambda}^{(\gamma)}\hat{\gamma}_{t-x}^{(R)} + \hat{\nu}X_{t-x} \\ &= \hat{\alpha}_x^{(\Delta)} + \hat{\Lambda}^{(1)}(\kappa_t^{(R)} - c(t - \bar{t})) + \hat{\Lambda}^{(\gamma)}(\gamma_{t-x}^{(R)} + c((t - \bar{t}) - (x - \bar{x}))) + \hat{\nu}X_{t-x}\end{aligned}$$

In order to have $\hat{R}_{x,t} = R_{x,t}$, we require

- $\hat{\Lambda}^{(j)} = \Lambda^{(j)}$, i.e., that our sensitivities do not change from one set of identifiability conditions to any other;

- $\hat{\nu}X_y = \nu X_y - c(\lambda^{(\gamma)} - \lambda^{(1)})(y - \bar{y})$, i.e., we can add terms linear in year of birth to the deterministic term without it fundamentally changing form, and therefore that our deterministic regressors contain a linear trend in year of birth; and
- $\hat{\alpha}_x^{(\Delta)} = \alpha_x^{(\Delta)} - c\lambda^{(1)}(x - \bar{x})$, i.e., we can add linear functions to any parametric form for $\alpha_x^{(\Delta)}$ without it fundamentally changing form, and therefore that it must be either non-parametric or have a linear function of age, $g^{(2)}(x) = x - \bar{x}$, amongst the basis functions used in any parametric restriction.

In addition to the identifiability issues discussed here, it is also important that any parametric simplification for $\alpha_x^{(\Delta)}$ consists of more than one, constant term. As discussed in Tuljapurkar and Edwards (2009), multiple terms in $\alpha_x^{(\Delta)}$ allow higher moments of the observable distribution of deaths in the sub-population (such as the variance of age at death) to be captured by the relative model, as well as the difference in life expectancy between the two populations. These higher moments are important in the allowance for idiosyncratic risk in the sub-population, which is likely to be important in many circumstances, such as those discussed in Hunt and Blake (2015e).

We also see from the analysis above that the form of our deterministic regressors, X_y , will depend upon the mortality model being used for the reference population. From Hunt and Blake (2015b), if the model for the reference population contains age functions which span the polynomials to order p , then there will be unidentified polynomial trends in the cohort parameters of order $p + 1$. We must therefore ensure that the deterministic regressors in Equation 3 span the polynomials to order $p + 1$ and that any parametric simplification for the age function, $\alpha_x^{(\Delta)}$, also contains a basis function of the form $g^{(i)}(x) = x^{p+1}$.

For the classic APC model, $p = 0$ and therefore we would require that the deterministic regressors and age function are, at least, of linear order. Similarly, for the reduced Plat model and the models constructed by the general procedure in Section 4.1 and Appendix C, $p = 1$ and we require that the deterministic regressors are at least of quadratic order.

In summary, the identifiability issues present in APC mortality models

and discussed in Hunt and Blake (2015a,b) have important consequences for the relative mortality modelling approach used in this study. Most importantly, we require an additional νX_y term in the model and must be careful when specifying any parametric simplification for $\alpha_x^{(\Delta)}$, in order to ensure that our results do not depend on the arbitrary identifiability constraints we impose on the reference model. In the context of the reference model used in this study, described in Section 4.1, this means that we need the term

$$\nu X_y = \nu_1(y - \bar{y}) + \nu_2((y - \bar{y})^2 - \sigma_y)$$

in Equation 3, and any parametric simplification of $\alpha_x^{(\Delta)}$ must be of the form

$$\begin{aligned} \alpha_x^{(\Delta)} &= (\alpha^{(1)}, \alpha^{(2)}, \alpha^{(3)}, \alpha^{(4)}, \alpha^{(5)}) \begin{pmatrix} f^{(1)}(x) \\ f^{(2)}(x) \\ f^{(3)}(x) \\ f^{(4)}(x) \\ ((x - \bar{x})^2 - \sigma_x) \end{pmatrix} \\ &= \sum_{i=1}^{N+1} \alpha^{(i)} \tilde{f}^{(R,i)}(x) \end{aligned}$$

where $f^{(i)}(x)$ are the parametric age functions in the reference model, described in Table 1.

C Models constructed by the “general procedure” for the UK

In Hunt and Blake (2014), a “general procedure” for constructing mortality models tailored to the specific features of individual datasets was proposed. In outline, this

- starts from a simple static mortality model with a non-parametric static age function;
- sequentially adds age/period terms to the model to detect and capture the age/period structure in the data:
 - structure is detected by adding a non-parametric age/period term which will identify the feature explaining the largest proportion of the remaining structure in the data;

- then this term is simplified into a parametric form which identifies the same feature more parsimoniously and with greater demographic significance;
- then the statistical significance and robustness of the term is tested;
- finally adds a cohort term once all age/period structure has been captured by the model;
- tests the standardised deviance residuals of the model for any remaining structure, independence, and normality.

This procedure was applied to data from the Human Mortality Database (2014) for men and women in the UK for ages 50 to 100 and years 1950 to 2011 in order to construct mortality models capable of capturing all the relevant information in the data and therefore allowing it to be projected appropriately.

A brief description of the terms in the models and their demographic significance is given in Table 1. A fuller list of the parametric age functions in the “toolkit” developed as part of the general procedure is given in the Appendix of Hunt and Blake (2014).

As discussed in Section 4, we also require additional identifiability constraints in order to obtain a unique set of parameters when fitting the model to data. These are given in Section 4 and have been chosen to aid comparability between the models for the reference population and the relative model in Equation 3.

When fitting the final models, we obtain the parameters shown in Figures 1 and 2. These models have BICs of -1.96×10^4 and -2.00×10^4 for men and women, respectively, with 407 free parameters for both populations.²³ We also test the standardised deviance residuals from fitting the model as part of the general procedure. The moments of the residuals and a Jarque-Bera test of their normality is given in Table 6. We can see that the residuals are close to normal and therefore pass the relevant Jarque-Bera tests for normality at

²³For comparison, the Lee and Carter (1992) model fitted to the same data obtains BICs of -2.71×10^4 and -2.69×10^4 with 161 free parameters for both populations.




Figure 14: Correlations for sequential years and ages of the residuals from fitting the model developed by the general procedure to data for the UK

the 5% level (p-values of 8.7% for both men and women), although they are slightly leptokurtic for both datasets. We also see from Figures 14a and 14b that there appears to be relatively little correlation structure over consecutive ages, although the residuals show significant autocorrelations during the early part of the data range, which diminishes towards the end of the period of the data.

	Residual mean	Standard deviation	Residual skewness	Residual kurtosis	Jarque-Bera statistic
Men	0.00	0.93	-0.05	3.15	4.84
Women	-0.00	0.93	0.02	3.19	4.82

Table 6: Moments of the residuals from fitting the model developed by the general procedure to data for the UK for men and women in the UK

The heat maps for the residuals shown in Figure 15 indicate that the residuals for both sexes in the UK have very little remaining structure in them. There is possibly some remaining structure around age 80 for both men and women, although this appears to be specific to only a few neighbouring years and therefore it is difficult to add an age/period term to capture this without overfitting the models




Figure 15: Heat maps of the residuals from fitting the model developed by the general procedure to data for the UK

References

- Aro, H., 2014. Systematic and nonsystematic mortality risk in pension portfolios. *North American Actuarial Journal* 18 (1), 59–67.
- Blake, D., Cairns, A. J. G., Dowd, K., 2006. Living with mortality: Longevity bonds and other mortality-linked securities. *British Actuarial Journal* 12 (1), 153–197.
- Brouhns, N., Denuit, M. M., Van Keilegom, I., 2005. Bootstrapping the Poisson log-bilinear model for mortality forecasting. *Scandinavian Actuarial Journal* 2005 (3), 212–224.
- Brouhns, N., Denuit, M. M., Vermunt, J., 2002. Measuring the longevity risk in mortality projections. *Bulletin of the Swiss Association of Actuaries* 2, 105–130.
- Cairns, A. J. G., 2000. A discussion of parameter and model uncertainty in insurance. *Insurance: Mathematics and Economics* 27 (3), 313–330.
- Cairns, A. J. G., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., Khalaf-Allah, M., 2011a. Mortality density forecasts: An analysis of six stochastic mortality models. *Insurance: Mathematics and Economics* 48 (3), 355–367.

Cairns, A. J. G., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., Ong, A., Balevich, I., 2009. A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. *North American Actuarial Journal* 13 (1), 1–35.

Cairns, A. J. G., Blake, D., Dowd, K., Coughlan, G. D., Khalaf-Allah, M., 2011b. Bayesian stochastic mortality modelling for two populations. *ASTIN Bulletin* 41 (1), 29–59.

Cairns, A. J. G., Blake, D., Dowd, K., Kessler, A., 2015. Phantoms never die: Living with unreliable mortality data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, Forthcoming.

Cairns, A. J. G., Dowd, K., Blake, D., Coughlan, G. D., 2013. Longevity hedge effectiveness: A decomposition. *Quantitative Finance* 14 (2), 217–235.

Campos, J., Ericsson, N. R., Hendry, D. F., 2005. General-to-specific modeling: An overview and selected bibliography. *Federal Reserve Discussion Papers*.

Carter, R., Lee, D., 1992. Modeling and forecasting US sex differentials. *International Journal of Forecasting* 8, 393–411.

Continuous Mortality Investigation, 2008. Working Paper 35 - The graduations of the CMI self-administered pension schemes 2000-2006 mortality experience: Final “S1” series of mortality tables.

URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-35>

Continuous Mortality Investigation, 2009. Working Paper 38 - A prototype mortality projection model: Part one - an outline of the proposed approach.

URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-38>

Continuous Mortality Investigation, 2011. Working Paper 53 - An initial investigation into rates of mortality improvement for pensioners of self-administered pension schemes.

URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-53>

Continuous Mortality Investigation, 2012. Working Paper 61 - An investigation into the mortality experience by industry classification of pensioners.

URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-61>

Continuous Mortality Investigation, 2014a. Working Paper 71 - Graduations of the CMI SAPS 2004-2011 mortality experience based on data collected by 30 June 2012: Final “S2” series of mortality tables.

URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-71>

Continuous Mortality Investigation, 2014b. Working Paper 73 - Analysis of the mortality experience of pensioners of self-administered pension schemes for the period 2005 to 2012.

URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-73>

Continuous Mortality Investigation, 2014c. Working Paper 76 - Analysis of the mortality experience of pensioners of self-administered pension schemes for the period 2006 to 2013.

URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-76>

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

Delwarde, A., Denuit, M. M., Guillén, M., Vidiella-i Anguera, A., 2006. Application of the Poisson log-bilinear projection model to the G5 mortality experience. *Belgian Actuarial Bulletin* 6 (1), 54–68.

Donnelly, C., 2014. Quantifying mortality risk in small defined-benefit pension schemes. *Scandinavian Actuarial Journal* (1), 41–57.

Dowd, K., Cairns, A. J. G., Blake, D., Coughlan, G. D., 2011. A gravity model of mortality rates for two related populations. *North American Actuarial Journal* 15 (2), 334–356.

Haberman, S., Kaishev, V. K., Millossovich, P., Villegas, A. M., Baxter, S. D., Gaches, A. T., Gunnlaugsson, S., Sison, M., 2014. Longevity basis risk: A methodology for assessing basis risk. Tech. rep., Cass Business School, City University London and Hymans Robertson LLP.

Hobcraft, J., Menken, J., Preston, S. H., 1982. Age, period and cohort effects in demography: A review. *Population Index* 48 (1), 4–43.

- Human Mortality Database, 2014. Human Mortality Database. Tech. rep., University of California, Berkeley and Max Planck Institute for Demographic Research.
URL www.mortality.org
- Hunt, A., Blake, D., 2014. A general procedure for constructing mortality models. *North American Actuarial Journal* 18 (1), 116–138.
- Hunt, A., Blake, D., 2015a. Identifiability in age/period mortality models. Tech. rep., Pensions Institute PI 15-08, Cass Business School, City University London.
- Hunt, A., Blake, D., 2015b. Identifiability in age/period/cohort mortality models. Tech. rep., Pensions Institute PI 15-09, Cass Business School, City University London.
- Hunt, A., Blake, D., 2015c. Modelling longevity bonds: Analysing the Swiss Re Kortis Bond. *Insurance: Mathematics and Economics* 63, 12–29.
- Hunt, A., Blake, D., 2015d. On the structure and classification of mortality models. Tech. rep., Pensions Institute PI 15-06, Cass Business School, City University London.
- Hunt, A., Blake, D., 2015e. Transferring mortality and longevity risks in pension schemes via bespoke longevity swaps. Work in Progress.
- Jarner, S. F., Kryger, E. M., 2011. Modelling mortality in small populations: The SAINT model. *ASTIN Bulletin* 41 (2), 377–418.
- Koissi, M., Shapiro, A., Hognas, G., 2006. Evaluating and extending the Lee-Carter model for mortality forecasting: Bootstrap confidence interval. *Insurance: Mathematics and Economics* 38 (1), 1–20.
- Lee, R. D., Carter, L. R., 1992. Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association* 87 (419), 659–671.
- Li, J. S.-H., Hardy, M. R., 2011. Measuring basis risk in longevity hedges. *North American Actuarial Journal* 15 (2), 177–200.
- Li, J. S.-H., Zhou, R., Hardy, M. R., 2015. A step-by-step guide to building two-population stochastic mortality models. *Insurance: Mathematics and Economics* 63, 121–134.

- Li, N., Lee, R. D., 2005. Coherent mortality forecasts for a group of populations: An extension of the Lee-Carter method. *Demography* 42 (3), 575–594.
- Oppers, S. E., Chikada, K., Eich, F., Imam, P., Kiff, J., Kisser, M., Soto, M., Kim, Y. S., 2012. The financial impact of longevity risk. Tech. rep., International Monetary Fund.
- Plat, R., 2009a. On stochastic mortality modeling. *Insurance: Mathematics and Economics* 45 (3), 393–404.
- Plat, R., 2009b. Stochastic portfolio specific mortality and the quantification of mortality basis risk. *Insurance: Mathematics and Economics* 45 (1), 123–132.
- Richards, S. J., 2008. Detecting year-of-birth mortality patterns with limited data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171 (1), 279–298.
- Russolillo, M., Giordano, G., Haberman, S., 2011. Extending the Lee-Carter model: A three-way decomposition. *Scandinavian Actuarial Journal* 2011 (2), 96–117.
- Salhi, Y., Loisel, S., 2009. Longevity basis risk modeling: A co-integration based approach. Tech. rep., University of Lyon.
- Sithole, T., Haberman, S., Verrall, R., 2012. Second international comparative study of mortality tables for pension fund retirees: A discussion paper. *British Actuarial Journal* 7 (3), 650–671.
- The Pensions Regulator, 2013a. Scheme funding. Tech. rep.
URL <http://www.thepensionsregulator.gov.uk/codes/code-funding-defined-benefi>
- The Pensions Regulator, 2013b. The Purple Book. Tech. rep.
URL <http://www.thepensionsregulator.gov.uk/doc-library/research-analysis.asp>
- Tuljapurkar, S., Edwards, R., 2009. Variance in death and its implications for modeling and forecasting mortality. NBER Working Papers.
- Villegas, A. M., Haberman, S., 2014. On the modeling and forecasting of socioeconomic mortality differentials: An application to deprivation and mortality in England. *North American Actuarial Journal* 18 (1), 168–193.

- Wan, C., Bertschi, L., 2015. Swiss coherent mortality model as a basis for developing longevity de-risking solutions for Swiss pension funds: A practical approach. *Insurance: Mathematics and Economics* 63, 66–75.
- Yang, B., Li, J., Balasooriya, U., 2015. Using bootstrapping to incorporate model error for risk-neutral pricing of longevity risk. *Insurance: Mathematics and Economics* 62, 16–27.
- Zhou, R., Wang, Y., Kaufhold, K., Li, J. S.-H., Tan, K. S., 2014. Modeling period effects in multi-population mortality models: Applications to Solvency II. *North American Actuarial Journal* 18 (1), 150–167.