Modelling Mortality for Pension Schemes*

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

November 2016

Abstract

For many pension schemes, a shortage of data limits their ability to use sophisticated stochastic mortality models to assess and manage their exposure to longevity risk. In this study, we develop a mortality model designed for such pension schemes, which compares the evolution of mortality rates in a sub-population with that observed in a larger reference population. We apply this approach to data from the CMI Self-Administered Pension Scheme study, using UK population data as a reference. We then use the 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

*An extended version of this paper (Hunt and Blake (2016a)) is available on the Pensions Institute website (http://www.pensions-institute.org/workingpapers/wp1601.pdf), which contains additional results for female data and more technical details on the models used.

†We are grateful to Andrew Cairns and Pietro Millossovich for their helpful review on an earlier draft of this paper, and 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.

is much less of a problem than is commonly believed.
JEL Classification: C33, C51, C52, G22, J26

Keywords: Mortality modelling, age/period/cohort 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 schemes[1] established by companies 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

[1]In 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.

pension schemes.

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 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. This model has the advantages of parsimony and robustness, important properties when dealing with the smaller datasets available for pension schemes. We then use this 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 (2016b).

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 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 and model risk on these conclusions. Section 5 uses the 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 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 tables[2] 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 period[4] 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

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

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

[4]See Continuous Mortality Investigation (2011).

[5]See Continuous Mortality Investigation (2015).

[6]See 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 The proposed model

A number of different models have been proposed in order to analyse mortality for various different populations. Many of these, however, have assumed that the different populations are of comparable size (e.g., different countries in Hunt and Blake (2015c)) or the smaller populations are at least of sufficient size to be able to estimate a large number of parameters for them (for instance, see Dowd et al. (2011) and Villegas and Haberman (2014)). However, with UK pension schemes, the lack of available data means that we require a far simpler approach.

In order to achieve this, we fit a sophisticated model to the larger "reference" population (typically the national population), where we have sufficient data to be able to robustly estimate a larger number of age, period and cohort terms. For the sub-population, we then use the more limited data we have to estimate the difference in the level of mortality observed across different ages and a series of scaling factors, which scale the period and cohort parameters from the reference population for the sub-population. This approach is considerably more parsimonious than fitting a separate model for the sub-population, with a corresponding increase in the robustness of our parameter estimates. 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) models[7] of the form

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

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 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.

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

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

3.2 The sub-population model

To analyse the data for the sub-population, we use a model of the form

$$\ln\left(\mu_{x,t}^{(S)}\right) = \alpha_x^{(R)} + \alpha_x^{(\Delta)} + \sum_{i=1}^{N} \lambda^{(i)} f^{(R,i)}(x)\kappa_t^{(R,i)} + \lambda^{(\gamma)}\gamma_{t-x}^{(R)} + \nu X_{t-x} \qquad (2)$$

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 scaling factors, $\lambda^{(j)}$ where $j \in \{1, \ldots, N, \gamma\}$. 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. Details of these issues are given in an online appendix.

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

1. $\lambda^{(j)} = 0$: the sub-population has no dependence on the j^{th} age/period or cohort term; and
2. $\lambda^{(j)} = 1$: there is no difference between the reference and sub-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 structure[9] 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.

When fitting the 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 model to very small datasets comparable to the size of individual pension schemes, as done in Section 6.

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

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

4 Applying the model to SAPS data

4.1 The reference models for UK data

Our first task is to construct suitable mortality models for men in the national UK population.[11] 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,[12] plus a cohort term. All of these terms are shown in Figure 1.[13]

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

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

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 mostly standard and have been used to impose our desired demographic significance on the parameters. Details of these constraints and further technical information about the models used are given in an online appendix. However, they are arbitrary, in the sense that they do not affect the fit to historical data, and so care has been taken to ensure that our choice does not affect our conclusions.

[11] We have performed a similar analysis for women and find broadly comparable results. However, for reasons of space, these results are not presented here.

[12] Demographic 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.

[13] In Figure 1c, one of the most notable features of the cohort parameters is the presence of large outliers in 1919/20 and 1946/47. We believe, based on the analysis of Richards (2008)and Cairns et al. (2015), that these are not genuine cohort effects, hence we use indicator variables to remove the impact of outliers from the cohort parameters.

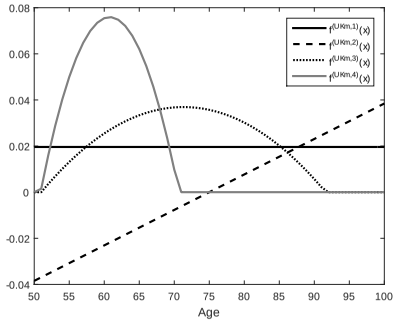

(a) Age functions

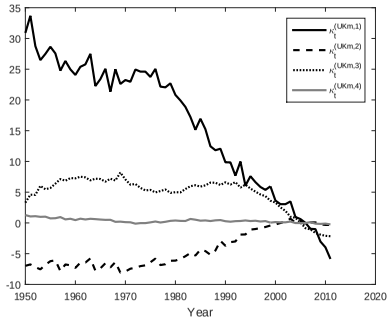

(b) Period functions

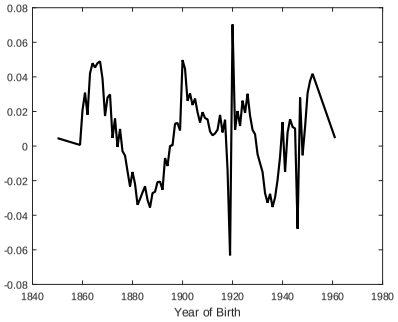

(c) Cohort function

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

4.2 The model for the SAPS data

We now estimate the 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 Table 2 below.

We then allow these restrictions to be relaxed sequentially. This means that, in turn, we estimate the model for the sub-population 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 Table 2, for the dataset. This process is illustrated in Figure 2.

Several of the models tested, with representative combinations of restrictions, are shown in Table 2 for the SAPS data.[14] 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 $\lambda^{(j)}$.

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

[14] In Table 2, "NP" stands for non-parametric while "P" stands for parametric.

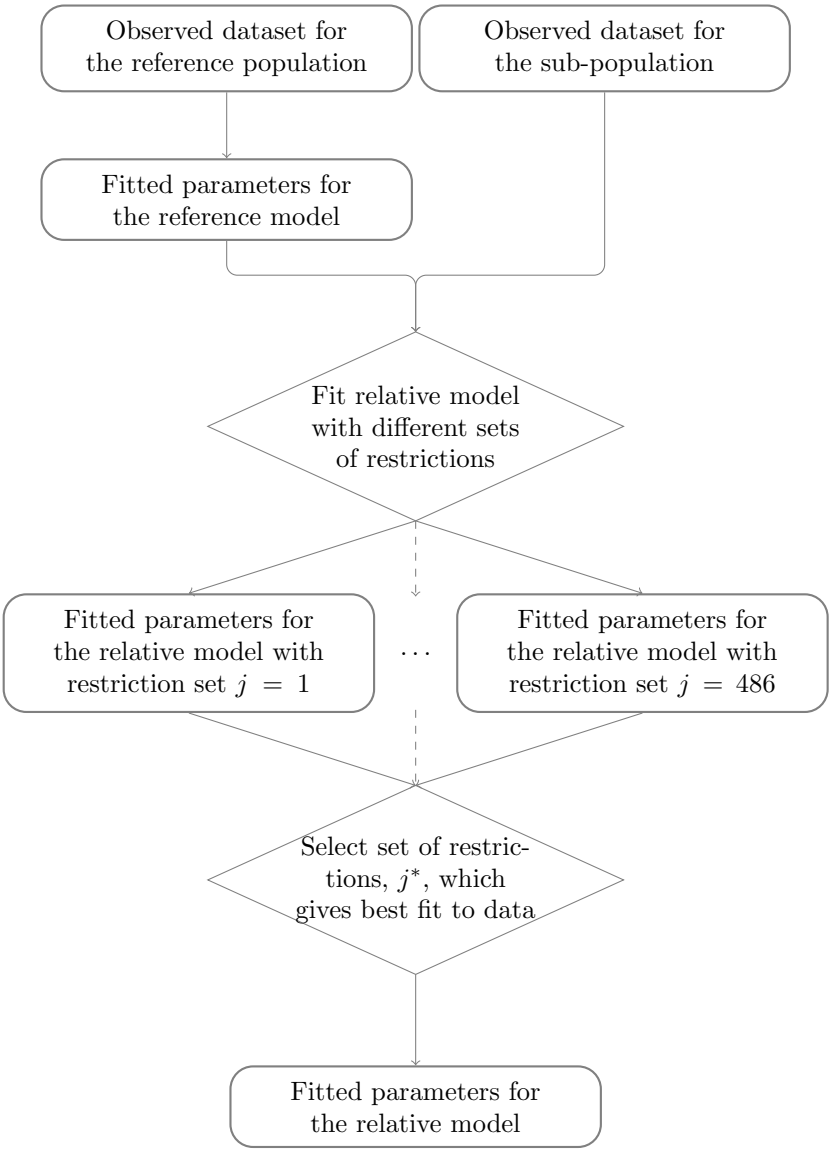

Figure 2: Flow chart illustrating the procedure for fitting and selecting the relative model

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 model using male SAPS data

similar restrictions placed on the scaling factors, $\lambda^{(j)}$, e.g., Models 1 and 2, or Models 4 and 5 in Table 2. In some respects, this supports the traditional actuarial practice of adjusting mortality rates for a pension scheme by taking a mortality table from a reference population (in this case, the full UK population) and making relatively simple adjustments to it. We also see from Figure 3 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).

We note that 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 is the same model as Model 4 in Table 2 however, the structure of Model 4 was selected a prioiri, whilst that for Model 8 was selected after an exhaustive search of all possible model structures. This means that the model finds no difference between the evolution of mortality rates for men in the SAPS data and the national population.[15] This is developed further in Section 4.3.1.

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 and so model risk (the risk of using an inappropriate model) may be

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

a potential issue. 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.

Both parameter uncertainty and model risk could be incorporated using a Bayesian approach with model averaging. However, doing so would use Bayesian Markov Chain Monte Carlo techniques, which are unfamiliar to most pension schemes and their advisors. Therefore, whilst Bayesian approaches have much to commend them, we will adopt a more familiar frequentist approach in addressing these issues.

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 model, i.e., that the parameters for the reference model are implicitly assumed to be known when the model for the sub-population 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 model.

The next step is to bootstrap new pseudo-data for the sub-population.

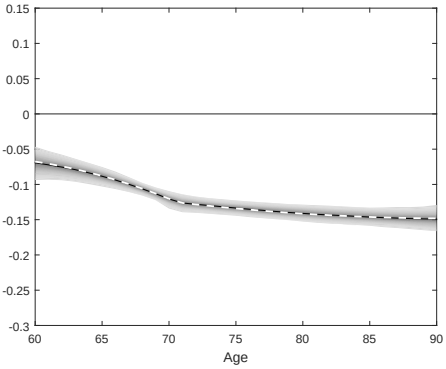

Figure 3: 95% fan chart showing the level of parameter uncertainty in $\alpha_x^{(\Delta)}$

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 Table 2. 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 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 to the data. Instead, we impose the same set of restrictions as were used for Model 8 in Table 2.

Figure 3 shows the impact of parameter uncertainty on the level parameters for the SAPS population by showing the 95% fan chart. We therefore conclude that the differences in the level of mortality between the national and SAPS populations are statistically significant across the entire age range.

However, substantial statistical uncertainty exists in the differences in the level of mortality, which can have important implications for the cashflows from pension schemes, as we discuss in Hunt and Blake (2016b).

Because the preferred model restricts the scaling factors for the reference period and cohort functions to be equal to unity, no parameter uncertainty is allowed for in their estimation.[16] Hence, allowing for parameter uncertainty alone will significantly understate the potential uncertainty in the approach and so we also need to consider model risk.

4.3.2 Model risk

The second stage of testing the robustness of the model is to fit the model to the bootstrapped data for the sub-population 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 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 only 36% of the bootstrapped datasets. Next, Table 3 shows the frequency of observing the various restrictions on the scaling factors in the preferred model. We note that the most common models chosen tend to restrict the scaling factors to equal unity in the same way as preferred in Model 8 in Table 2. The exception to this is for $\lambda^{(\gamma)}$, where an unrestricted value is preferred in the majority of cases.

Table 3 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 Table 2. The exception to this is for $\lambda^{(\gamma)}$,

[16] This is in contrast to when the model is applied to female data, where some of the scaling factors are allowed to vary freely and hence are subject to parameter uncertainty.

	$\lambda^{(j)}=0$	$\lambda^{(j)}=1$	$\lambda^{(j)}$ unrestricted
$\lambda^{(1)}$	0%	70%	30%
$\lambda^{(2)}$	47%	53%	0%
$\lambda^{(3)}$	0%	97%	3%
$\lambda^{(4)}$	44%	55%	1%
$\lambda^{(\gamma)}$	1%	37%	62%

Table 3: Frequency with which different restrictions are placed upon the scaling factors in the preferred model, based on 1,000 bootstrapped datasets

where an unrestricted value is preferred.

In summary, we find that there is substantial model risk, 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 model was applied to historical data for the SAPS population. Given projections of the reference population, we can also use the 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 mortality 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, known differences in mortality rates between populations form the mortality basis, not the 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 this constitutes the basis between the populations. In this case, we could still construct portfolios using securities linked to mortality in population B to hedge mortality in population A perfectly. 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 (2016b)). 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 model of Equation 2, 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 $\alpha_x^{(\Delta)}$, as shown in Figure 3. However, in our preferred model for the sub-population, the $\lambda^{(j)}$ are restricted to unity and hence there will still be no uncertainty in the trend basis in the reference population, when allowing for parameter uncertainty alone. Hence, it is only appropriate to talk about "basis risk in conjunction with our preferred model for the SAPS data 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

$$\boldsymbol{\kappa}_t^{(R)} = \mu^{(R)} \begin{pmatrix} 1 \\ t \end{pmatrix} + \boldsymbol{\kappa}_{t-1}^{(R)} + \boldsymbol{\epsilon}_t^{(R)} \tag{3}$$

where $\boldsymbol{\kappa}_t^{(R)} = \begin{pmatrix} \kappa_t^{(R,1)}, & \dots & \kappa_t^{(R,N)} \end{pmatrix}^\top$, $\mu^{(R)}$ is a matrix of drift coefficients with respect to the period "trends", $\begin{pmatrix} 1, & t \end{pmatrix}^\top$, 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

$$\gamma_y^{(R)} - \beta^{(R)} \begin{pmatrix} 1 \\ (y-\bar{y}) \\ ((y-\bar{y})^2 - \sigma_y) \end{pmatrix} = \rho^{(R)} \left[\gamma_{y-1}^{(R)} - \beta^{(R)} \begin{pmatrix} 1 \\ (y-1-\bar{y}) \\ ((y-1-\bar{y})^2 - \sigma_y) \end{pmatrix} \right] + \varepsilon_y \tag{4}$$

where $\beta^{(R)}$ is a matrix of drift coefficients with respect to the cohort "trends", $\begin{pmatrix} 1, & (y-\bar{y}), & ((y-\bar{y})^2 - \sigma_y) \end{pmatrix}^\top$. These deterministic functions are chosen to ensure that the projections are "well-identified", i.e., that the projected mortality rates for the reference population do not depend upon the identifiability constraints used when fitting the model.

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 (2016b), 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 by using the approach developed here and allowing fully for both parameter uncertainty and model risk. 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%. Analysing the impact of process, parameter and model risk separately, we find that it is the potential for model mis-specification which adds most significantly to the basis risk for both populations.

Figure 4 shows scatter plots of annuity values calculated using mortality rates in the UK and SAPS populations. First, we note that the systematic longevity risk (indicated by the range of values the annuity value can take) 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,[17] indicating that basis risk may be considerably less important than is widely believed believed. Of the remaining 10% of the uncertainty attributable to basis risk, we find that approximately 8% is attributable to model risk and only 2% to parameter uncertainty.[18] Whilst it is true that the preferred model for the male SAPS data may understate the impact of parameter uncertainty in a more general case (since all the $\lambda^{(j)}$ are set to unity), this still indicates that the large majority of basis risk arises from fundamental uncertainty over the correct model to use. This is discussed further in Section 7.

Second, Figure 4 shows that, the points tend to cluster depending on the

[17] As measured as the proportion of the observed variance explained by a regression of the SAPS annuity value on the national population annuity value.

[18] Found by comparing the projected annuity values when either only systematic longevity risk or systematic longevity risk and parameter uncertainty are considered.

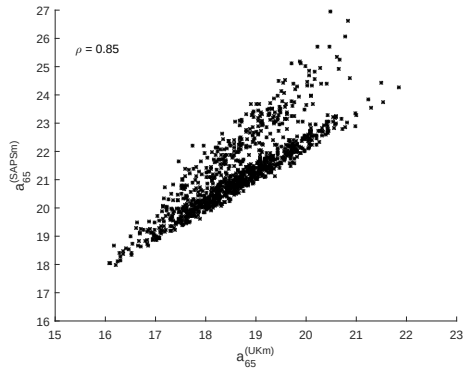

Figure 4: Projected annuity values for the UK and SAPS populations from 1,000 Monte Carlo simulations

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. In particular, we find that the second cluster of points observed in Figure 4 arises from the 30% of simulations where $\lambda^{(1)}$ is unrestricted (see Table 3. These tend to find values of $\lambda^{(1)} > 1$ and hence faster improvements in mortality rates in the SAPS population than the reference population, as shown in Figure 4.

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., $\boldsymbol{\kappa}_t^{(R)}$ 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. Although we have not performed a detailed investigation,[19] a good rule of thumb says that parameter uncertainty is inversely proportional to the size of the dataset in terms of lives. Hence, if we reduce the size of the population from over one million lives to around 100,000, we would expect the importance of parameter uncertainty to more than treble. Consequently, Wwe 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 (2016b).

6 Applying the 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.

[19] Section 6 looks only at in-sample model fitting rather than projection and basis risk.

As discussed in Section 4.2, the 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 subpopulation such as that associated with an individual pension scheme.

We investigate the effect of population size on the ability of the model to measure mortality differences with the national population by randomly generating scheme-sized exposures to risk and death counts based on the SAPS data. We do this first by scaling the exposures to risk from the SAPS data appropriately to proxy for pension schemes of the desired size. Then, we generate random death counts for the scheme by modelling them as Poisson random variables with the expected number of deaths found using the crude mortality rates observed in the SAPS dataset. We fit the 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.

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 are shown in Figure 5.

First, let us consider the results shown in Figure 5a for the level of mortality in the smaller pension schemes. 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. 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).

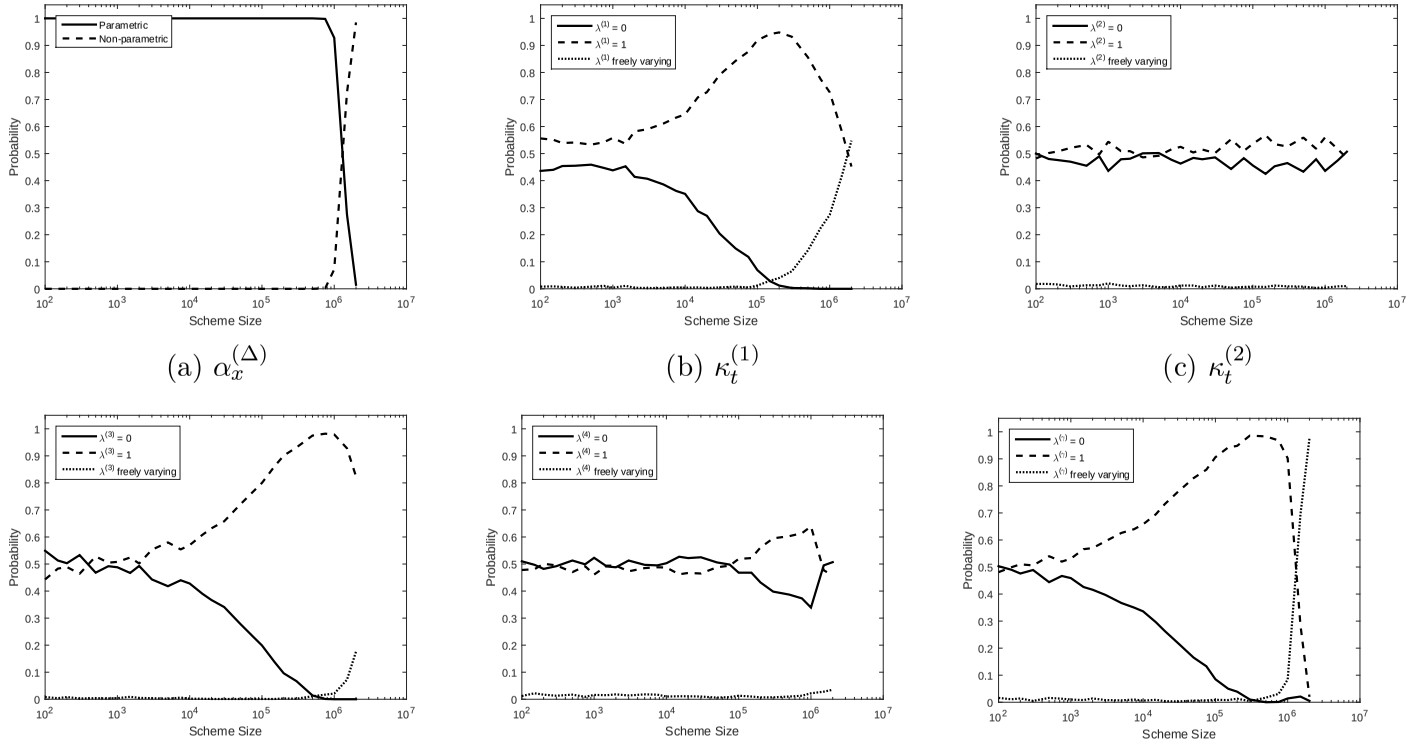

(a) $\alpha_x^{(\Delta)}$ (b) $\kappa_t^{(1)}$ (c) $\kappa_t^{(2)}$

(d) $\kappa_t^{(3)}$ (e) $\kappa_t^{(4)}$ (f) γ_y

Figure 5: Restrictions placed on the relative model for different volumes of male data

Looking at the scaling factors for the age/period and cohort terms, we see that, typically, the smallest schemes (fewer than 1,000 members) are indifferent between restricting $\lambda^{(j)}$ to be equal to zero or unity. For instance, Figure 5b shows that the procedure imposes the restriction $\lambda^{(1)} = 0$ and $\lambda^{(1)} = 1$ 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 Figure 5. 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 5b again, we see that for larger schemes, with around 10,000 to 100,000 members, the model has a clear preference for setting $\lambda^{(1)} = 1$, which is preferred in almost all simulations for schemes with around 200,000 members. This pattern also holds for $\lambda^{(3)}$ and $\lambda^{(\gamma)}$ (albeit for slightly larger schemes). For the other scaling factors, 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 any trend basis 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 scheme-specific 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 in schemes with around one million members in Figure 5b illustrates that it is only in

the very largest schemes 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 (2016b).

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 (2016b). 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.

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 (2009), 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. (2011), 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 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, especially when parameter uncertainty and model risk are included in any analysis. This is the trade-off between the ability to model the mortality basis 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, especially when parameter uncertainty and model risk are included in the analysis. For the full SAPS dataset, the simple 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 trend basis 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 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 proposed 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 the mortality 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 mortality basis according to our definition, since it does not allow for any uncertainty in the basis in future. 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 modelling 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 basis 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 as fashion as to maximise the likelihood of observing different patterns in the evolution of mortality rates.

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, we note that very large pension schemes or aggregated data sets from many pension schemes (such as the SAPS dataset) may change in composition over the period of the data. It will therefore be unclear if genuine trend basis is being detected, or merely a change in the composition of the underlying data.[20]

[20]In particular, we note that the data for CMI Assured Lives has varied considerably in

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 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.

The practical implications of these results are important for the development of any market in longevity hedging. It is commonly believed in industry that basis risk is sufficiently important as to prevent the feasibility of index-based hedging instruments for longevity risk. To date, the vast majority of longevity risk transfer has been performed on a bespoke basis, with only limited attempts (such as the Kortis bond discussed in Hunt and Blake (2015c)) to develop instruments linked to systematic risks alone. However, we do not believe this view is consistent with the correlations (and hence hedge effectivenesses) of around 85% in Section 5.

However, we believe a lot of this is due to the confusion between basis (which can be solved by efficient structuring of the hedge) and basis risk and, furthermore, the confusion between level basis risk (which can be reduced through the methods discussed above) and trend basis risk (which

the socio-economic makeup of the relevant population over its lifetime due to changes in the UK annuity market. Since this dataset was used in Cairns et al. (2011), Dowd et al. (2011) and Cairns et al. (2013), it is 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.

is more of an "unknown unknown). 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 (2016b).

8 Conclusions

In conclusion, in this study we present a parsimonious 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 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 proposed 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 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 6 and 7 summarise the patterns of deaths and exposures for men and women across age and time.

References

Blake, D., Cairns, A. J., Dowd, K., 2006. Living with mortality: Longevity bonds and other mortality-linked securities. British Actuarial Journal 12 (1), 153–197.

[21]However, 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.

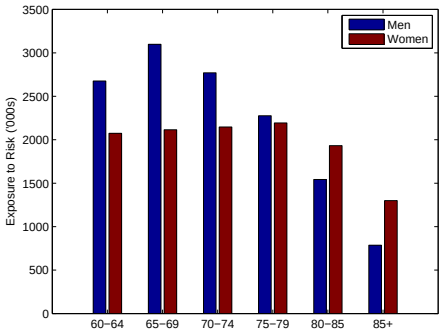

(a) Exposure to risk by age band

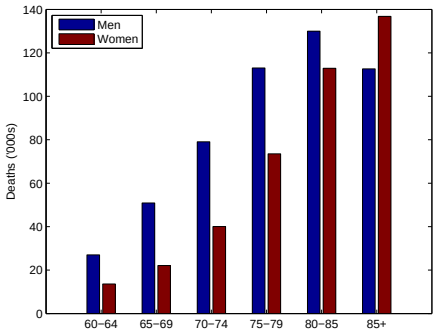

(b) Death count by age band

Figure 6: Exposures to risk and death counts in the SAPS dataset by age

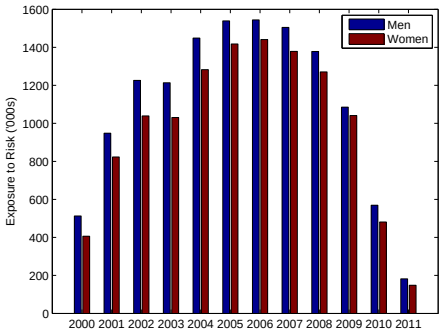

(a) Exposure to risk by year

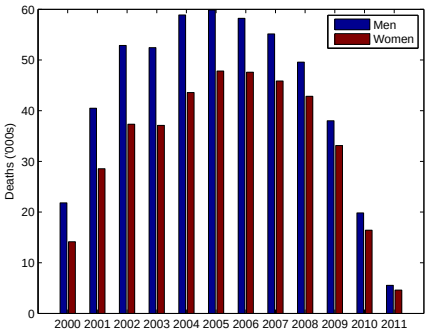

(b) Death count by year

Figure 7: Exposures to risk and death counts in the SAPS dataset by year

Cairns, A. J., 2000. A discussion of parameter and model uncertainty in insurance. Insurance: Mathematics and Economics 27 (3), 313–330.

Cairns, A. J., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., Khalaf-Allah, M., 2011. Mortality density forecasts: An analysis of six stochastic mortality models. Insurance: Mathematics and Economics 48 (3), 355–367.

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

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

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-papers-34-and-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-papers-38-and-39

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, 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

Continuous Mortality Investigation, 2015. Working Paper 86 - An investigation into the mortality experience by industry classification of pensioners of self-administered pension schemes for the period 2006 to 2013.

Coughlan, G. D., Khalaf-Allah, M., Ye, Y., Kumar, S., Cairns, A. J., 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.

Dowd, K., Cairns, A. J., 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.

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., 2016a. Basis risk and pension schemes: A relative modelling approach. Tech. rep., Pensions Institute PI 16-01, Cass Business School, City University London.

Hunt, A., Blake, D., 2016b. 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.

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.

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., 2009. 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.

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-benefits.aspx

The Pensions Regulator, 2013b. The Purple Book. Tech. rep.
URL http://www.thepensionsregulator.gov.uk/doc-library/research-analysis.aspx

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.

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.