

Small Population Bias and Sampling Effects in Stochastic Mortality Modelling

Liang Chen, Andrew J.G. Cairns and Torsten Kleinow

Actuarial Research Centre,

Department of Actuarial Mathematics and Statistics and the Maxwell Institute for Mathematical Sciences, School of Mathematical and Computer Sciences, Heriot-Watt University, EH14 4AS, Edinburgh, U.K.

December 10, 2015

Abstract

We propose the use of parametric bootstrap methods to investigate the finite sample distribution of the maximum likelihood estimator for the parameter vector of a stochastic mortality model. Particular emphasis is placed on the effect that the size of the underlying population has on the distribution of the MLE in finite samples, and on the dependency structure of the resulting estimator: that is, the dependencies between estimators for the age, period and cohort effects in our model. In addition, we study the distribution of a likelihood ratio test statistic where we test a null hypothesis about the true parameters in our model. Finally, we apply the LRT to the cohort effects estimated from observed mortality rates for females in England and Wales and males in Scotland.

KEYWORDS — Small population, age effect, period effect, cohort effect, bootstrap, parameter uncertainty, systematic parameter difference, likelihood ratio test, power of test

1 Introduction

Stochastic mortality models are widely used as risk management tools in the insurance and pensions industry with the main application being the generation of plausible scenarios for future mortality rates. Many stochastic mortality models have been introduced in the last few decades. When new models have been developed the objective was mostly to improve the goodness of fit of the model to mortality data observed in relatively large populations: the Lee-Carter model and its refinements (e.g. Lee and Carter (1992); Renshaw and Haberman (2003); and Booth et al. (2006)) have been developed to provide a good fit to the mortality rates observed in the United States, England and Wales and the population of UK male assured lives; while the Cairns-Blake-Dowd (Cairns et al. (2006)) model was introduced for modelling the England and Wales males population.

In contrast, actuaries will often face the problem of modelling the mortality experience of much smaller populations, for example, the members of a pension scheme. Empirical

research has found that mortality rates of smaller populations exhibit significantly more variability compared to the observed rates in larger populations. Furthermore, models that fit large countries well, might not be appropriate for smaller populations, for example, Booth et al. (2006) showed that the Lee-Carter model provides a rather poor fit to the mortality experience of smaller populations. A related issue is that empirical data from smaller populations might only be available for a relatively short period, which makes mortality projections rather uncertain. As a result, a number of recent papers have aimed to develop models specifically for smaller populations: for example, the Saint Model of Jarner and Kryger (2011).

A common assumption for many of the proposed models is that the observed numbers of deaths are realisations of random variables with a Poisson distribution given the underlying mortality rates. The estimation of parameters of any such model is therefore based on samples from a Poisson distribution, and, as always in statistics, parameter uncertainty is related to the sample size. Furthermore, many results about the distribution of estimators and corresponding confidence intervals rely on the Maximum Likelihood theorem and large sample sizes.

The increased uncertainty about estimated parameters for small populations results in high levels of uncertainty about projected mortality rates. As a consequence future realised mortality rates will not only diverge from projected rates due to future sampling variation caused by the Poisson distribution, but might also diverge from projections since the projections themselves are uncertain.

In the actuarial literature, simulation techniques have been proposed for dealing with uncertain parameters and projected mortality rates. For example, Liu and Braun (2011) investigated mortality uncertainty by applying a block bootstrap method on the Lee-Carter model, and Brouhns et al. (2005) proposed Poisson bootstrap methods for mortality forecasting.

However, to the best of our knowledge, bootstrap methods have not been applied in a systematic way to investigate the impact of the size of a population on parameter and projection uncertainty. This is the focus of our research in this paper. We firstly apply Poisson parametric bootstrap methods to investigate how the variation of parameter estimates and projections is affected by the size of a population. The specific mortality model that we consider is a second generation CBD model with added cohort effect: see section 2 for details. We vary the size of the population by assigning weights to a chosen benchmark population, e.g. England and Wales males. In simulation studies we find that the size of the population has a significant effect on the variation of parameter estimates and projections.

By assigning weight to the benchmark population, we assume that the constructed small populations have the same underlying mortality rates as the benchmark population. In such a situation the uncertainty about projected mortality rates will be reduced if the information from the "true" benchmark parameter estimates can be used for fitting smaller populations. This raises the question of how we can test for systematic differences between the parameters driving mortality rates in a small population and a given null hypothesis about those parameters, where the null hypothesis might have been obtained from a model fitted to a much larger population. If no significant differences can be found then it seems reasonable to use mortality projections from the large population to generate scenarios for the small population. We therefore investigate the properties of a likelihood ratio (LR) test for all or some of the estimated parameters, and, in particular, consider the distribution of the test statistic based on the bootstrap simulations. This

allows us to investigate the power of the LR test and the effect of varying population sizes on the rejection rates. We find that the population size has a strong effect on the probability of a type II error.

In an empirical study, we apply the LR-test to test a null hypothesis about the cohort effect estimated for mortality rates of females in England and Wales and males in Scotland. The null hypothesis for both populations is the cohort effect estimated for males in England and Wales, and we find for both populations a significant divergence of the estimated cohort effect from that in the null hypothesis.

The remainder of the paper is organised as follows. Section 2 introduces the model, assumptions and the notations we apply. Section 3 discusses the process of simulation and investigates the distribution of the maximum likelihood estimates, the correlation between the estimates and how these will be affected by changing the population size. In Section 4, we investigate the effect of the population size on forecasting by projecting the parameters as well as the mortality rates. Section 5 introduces a likelihood ratio test for testing systematic deviations of the true parameters from a given null hypothesis. The power of the likelihood ratio test is also analysed and we then investigate how significant the impact of shifting and scaling parameters is on the fitted mortality rates and corresponding annuity prices in Section 6. Finally, Sections 7 and 8 include the LRT for testing a null hypothesis about the cohort effect only, and an empirical example for this test is provided. Section 9 provides our final conclusions.

2 The Model

We denote by $D(t, x)$ the number of deaths during calendar year $t = t_1, \dots, t_{n_y}$ at age $x = x_1, \dots, x_{n_a}$ and by $E(t, x)$ the corresponding central exposure to risk.

We will fit the following Poisson model to the observed death data, see Cairns et al. (2009):

$$D(t, x)|\theta \sim \text{Pois}(m(\theta, t, x)E(t, x)) \quad (1)$$

$$m(\theta, t, x) = -\log(1 - q(\theta, t, x)) \quad (2)$$

$$\text{logit } q(\theta, x, t) = \kappa_t^{(1)} + \kappa_t^{(2)}(x - \bar{x}) + \kappa_t^{(3)}((x - \bar{x})^2 - \hat{\sigma}_x^2) + \gamma_{t-x}^{(4)} \quad (3)$$

where the parameter vector θ is given by

$$\theta = (\kappa_t^{(1)}, \kappa_t^{(2)}, \kappa_t^{(3)}, \gamma_c^{(4)})$$

with the following interpretations:

- $\kappa_t^{(i)}$ is a period effect in year $t = t_1, \dots, t_{n_y}$ for each $i = 1, 2, 3$,
- $\gamma_c^{(4)}$ is the cohort effect for the cohort born in year $c = t - x$,
- \bar{x} is the mean of the age range we use for our analysis, and
- $\hat{\sigma}_x^2$ is the mean of $(x - \bar{x})^2$.

It is well known that the parameters in model (3) are not identifiable without imposing constraints on their values. We follow Cairns et al. (2009) and apply the following constraints on θ :

$$\sum_{c \in C} \gamma_c^{(4)} = 0, \quad \sum_{c \in C} c \gamma_c^{(4)} = 0, \quad \sum_{c \in C} c^2 \gamma_c^{(4)} = 0 \quad (4)$$

where $C = t_1 - x_{na}, \dots, t_{ny} - x_1$ is the set of all years of birth in a given dataset.

To estimate the parameters in (3) we apply maximum likelihood estimation. The log-likelihood function in our model is

$$l(\theta; D, E) = \sum_{t,x} D(t, x) \log[E(t, x)m(\theta, t, x)] - E(t, x)m(\theta, t, x) - \log[D(t, x)!] \quad (5)$$

where $m(\theta, t, x)$ is given by (2) and (3).

As mentioned earlier, in this paper we are concerned with the consequences of small exposures, or population sizes, on the distribution of the maximum likelihood estimator (MLE) $\hat{\theta}$ of θ . To study the distribution of the MLE $\hat{\theta}$ we will simulate death data $D(t, x)$ from the model in (1)–(3) using a given parameter vector θ_0 and different exposure sizes.

To ensure that our results are relevant for typical values of θ we first fit our model to death and exposure data observed in England and Wales during the years 1961 to 2011 for males aged 50 to 89. We then fix θ_0 to be equal to the estimated parameter vector $\hat{\theta}^{\text{EW}}$ for this data. Note that this is only an example for the true parameter vector θ_0 and our analysis can be applied to other choices of θ_0 . Mortality data for England and Wales are obtained from the Human Mortality Database.¹ Note that we do not exclude short cohorts from the estimation (different, therefore, from Cairns et al. (2009)).

The different exposure sizes used to simulate data in the remainder of this paper will be relative to the exposure $E_0(t, x)$ for a benchmark population. For reasons of practical relevance and consistency with our choice of θ_0 the benchmark population is the male population in England and Wales unless stated otherwise.

3 Distribution of MLE in Finite Samples

For any given parameter vector θ_0 and benchmark exposure $E_0(t, x)$ we define the small-sample exposure as

$$E^w(t, x) = wE_0(t, x)$$

for a constant $w \leq 1$.

We then simulate N scenarios for the death counts $D^w(t, x)$ using the model in (1)–(3) with $\theta = \theta_0$. Through our simulation we obtain N independent scenarios $D_j^w(t, x)$ for the death counts with

$$D_j^w(t, x) \sim \text{Pois}(m(\theta_0, t, x)wE_0(t, x)) \quad \text{for all } j = 1, \dots, N. \quad (6)$$

3.1 MLE

To obtain MLEs for $\hat{\theta}_j^w$ for each simulated scenario j and each w we maximise the log-likelihood function $l(\theta; D_j^w, E^w)$ as given in (5) subject to the constraints in (4), that is,

$$\hat{\theta}_j^w = \arg \max_{\theta} l(\theta; D_j^w, E^w) \quad (7)$$

Classical sampling theory tells us that

$$\sqrt{w} \left(\hat{\theta}_j^w - \theta_0 \right) \xrightarrow{\text{Dist}} N(0, H) \quad \text{as } w \rightarrow \infty$$

¹Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de (data downloaded on [date]).

for some positive semi definite matrix H (see Appendix A for further discussion).

Therefore, we would expect that, even in a finite sample, the co-variance of the distribution of $\hat{\theta}_j^w$ is approximately to $w^{-1}H$ and the correlations between different components of $\hat{\theta}_j^w$ are approximately independent of the relative population size w . Using the simulated sample $\hat{\theta}_1^w, \dots, \hat{\theta}_N^w$ we can investigate the finite-sample covariance and correlation matrices of $\hat{\theta}^w$. In Figure 1 we plot a graphical representation of the correlation matrices of $\hat{\theta}_j^w$ that we obtain for two values of w .

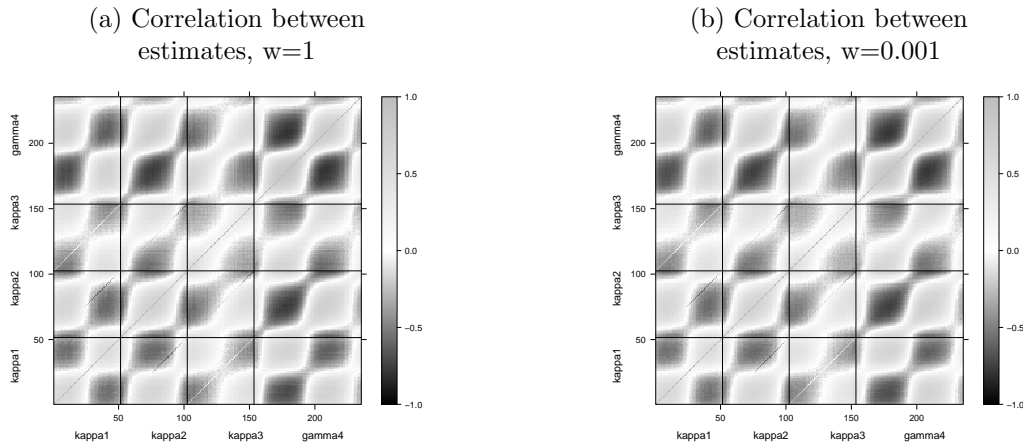


Figure 1: The empirical correlation matrix of the simulated parameter estimates $\hat{\theta}^w$ for different values of the population size $w = 1$ and $w = 0.001$. The grid lines at 51.5, 102.5 and 153.5 are used to visually separate the parameters $\kappa_t^{(1)}$, $\kappa_t^{(2)}$, $\kappa_t^{(3)}$, $\gamma_c^{(4)}$ from each other in both dimensions. For instance, the bottom left rectangle contains the correlations for $\hat{\kappa}^{w,(1)}$ for the 51 years from 1960 to 2011.

We conclude from Figure 1 that there are no significant differences between the empirical correlation matrices obtained from different population sizes, as predicted. However, individual components of $\hat{\theta}^w$ are not independent from each other as we would expect given the model in (1)–(3).

To investigate the finite-sample distribution of the MLE $\hat{\theta}^w$ further we plot the empirical mean together with 90% confidence intervals for each of the components of $\hat{\theta}^w$ in Figure 2.

We find for all population sizes considered that the empirical means of the simulated estimates fluctuate around the true parameter values θ_0 (solid line), which indicates that the MLE is approximately unbiased for all considered population sizes. However, the standard deviation of the estimator depends strongly on the size of the population, increasing significantly as the exposures get smaller as can be seen from the width of the confidence intervals.

The graphs on the right hand side of figure 2 show that the level of fluctuation increases approximately by a factor \sqrt{n} if the population size is reduced by a factor $1/n$, which is consistent with the asymptotic covariance matrix being proportional to $1/w$. It also suggests that the variance is generally stable for all the period effects over years, which is not the case for the cohort effect with wave shaped pattern. We notice that the standard deviation of $\gamma_{c,w}^{(4)}$ widens out considerably at both ends, reflecting the reducing number of observations that we have of the younger and older cohorts.

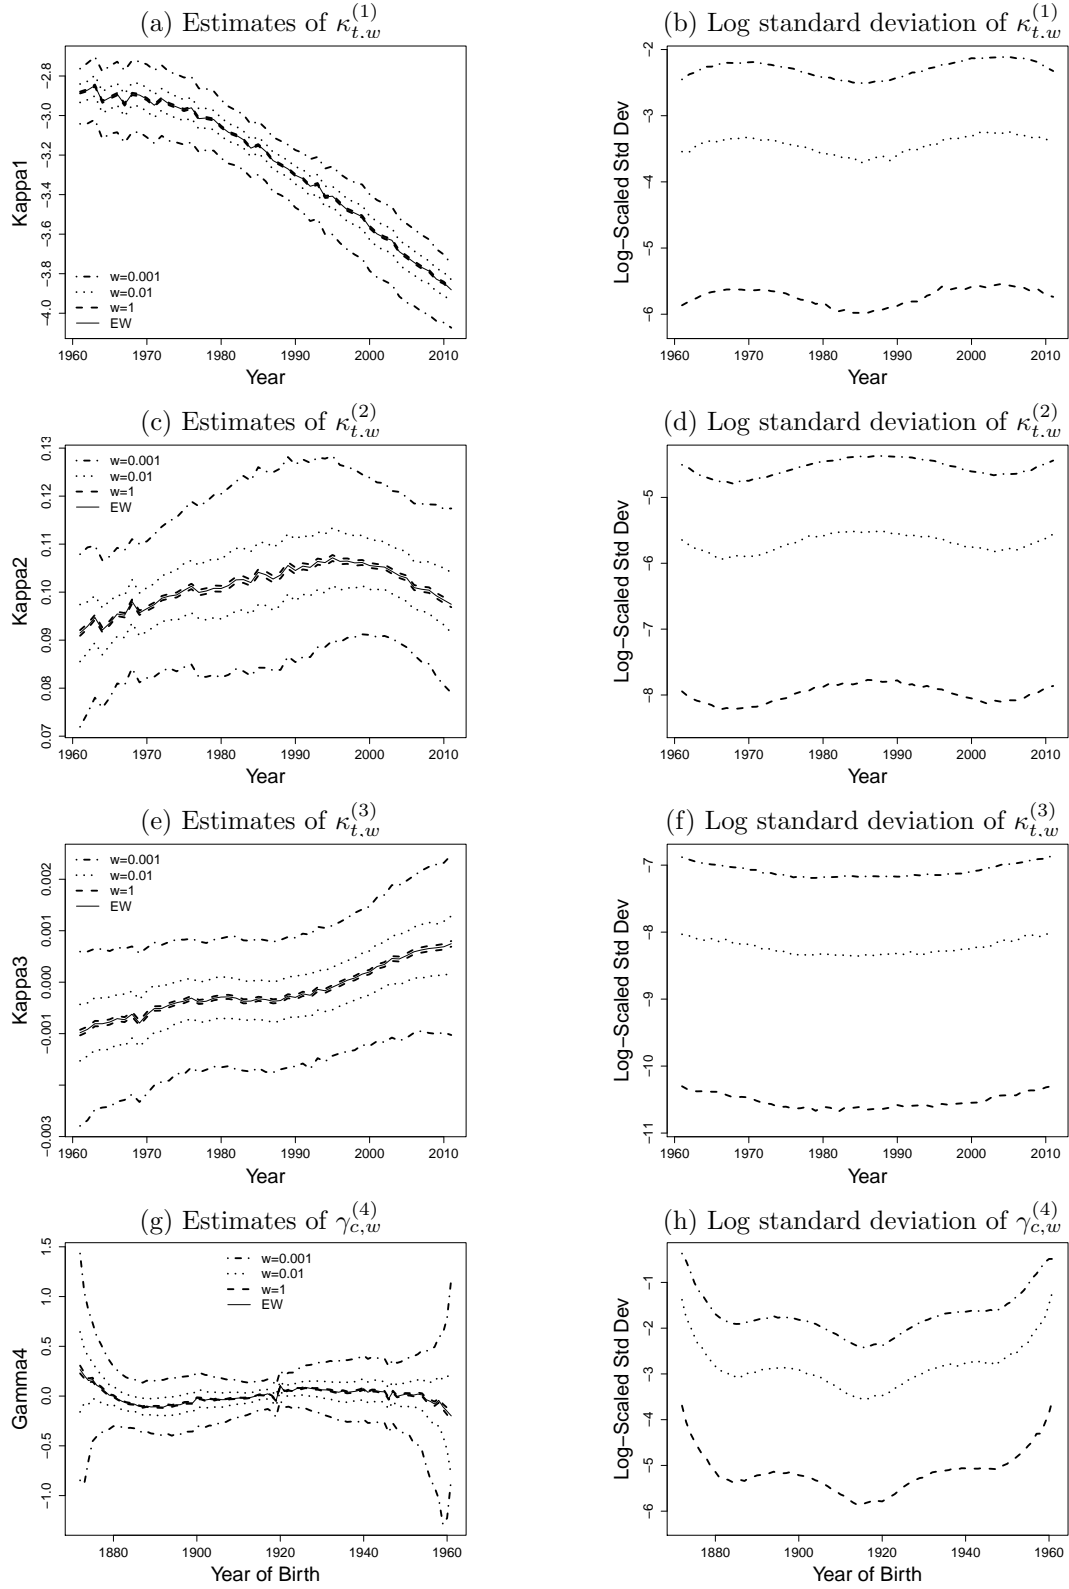


Figure 2: The distribution of MLEs: The mean and confidence interval of the MLEs of $\kappa_{t,w}^{(1)}$, $\kappa_{t,w}^{(2)}$, $\kappa_{t,w}^{(3)}$, $\gamma_{c,w}^{(4)}$, with respect to year t and year of birth c respectively, of populations with $w = 1$ (dashed line), 0.01 (dotted line), 0.001 (dot dashed line), together with the parameter estimates for the England and Wales population (solid line). Note: The upper bound of the CI in the left column is the 95% quantile of the distribution and the lower bound is the 5% quantile

4 Mortality Projections

While fitting the model in (1)–(3) to observed mortality data only requires the estimation of the period effects $\kappa_t = (\kappa_t^{(1)}, \kappa_t^{(2)}, \kappa_t^{(3)})'$ and the cohort effect $\gamma_c^{(4)}$, projecting mortality rates into the future requires a model for values of κ_t for $t > t_{n_y}$ where t_{n_y} is the last year for which mortality data are available. Similarly, future values of the cohort effect $\gamma^{(4)}$ are also required.

The most common approach to obtain future values of κ and $\gamma^{(4)}$ is to consider these parameter vectors as observed trajectories of stochastic processes and fit a parametric time series model to each trajectory. In the following we will fit a three-dimensional random walk to κ and a stationary AR(1) model to $\gamma^{(4)}$, as in Cairns et al. (2009). We will then discuss the estimation of the parameters of those models based on the values of θ_0 and $\hat{\theta}_j^w$ for different values of w . This will allow us to investigate the impact of the relative population size w on the estimators for the parameters of the κ and $\gamma^{(4)}$ processes.

For the estimation of those parameters and the projections of the period effects and the cohort effect we will consider two approaches. Firstly, we will use a frequentist approach to obtain point estimates of the process parameters ignoring any uncertainty about those estimates. In our further analysis we will follow a Bayesian approach to incorporate parameter uncertainty into our mortality projections.

4.1 Projecting Period Effects

As mentioned above, we model the period effects κ as a three-dimensional normal random walk.

$$\Delta\kappa_t = \mu + L\epsilon_t \quad (8)$$

where $\Delta\kappa_t = \kappa_t - \kappa_{t-1}$ and the $\epsilon_t = (\epsilon_t^{(1)}, \epsilon_t^{(2)}, \epsilon_t^{(3)})'$ are independent random vectors with a multivariate standard normal distribution. The parameter vector μ is the 3×1 drift vector of the random walk and L is the 3×3 Cholesky decomposition of the covariance matrix $V = LL'$.

4.1.1 Point Estimators

Having generated N scenarios for the number of deaths according to (6) and having estimated the parameter vector $\hat{\theta}_j^w$ in each scenario as in (7), we can now apply the random walk model to the period effects in θ_0 and $\hat{\theta}_j^w$ for every generated scenario j . We then apply the usual (i.e. maximum likelihood) point estimators $\hat{\mu}_j^w$ and \hat{V}_j^w for each simulated scenario of $D_j^w(t, x)$. The estimators for the three components of $\hat{\mu}_j^w$ (scenario j) are

$$\hat{\mu}_j^w(i) = \frac{1}{n_y - 1} \sum_{t=t_1}^{t_{n_y}} (\hat{\kappa}_{t,j}^{(i),w} - \hat{\kappa}_{t-1,j}^{(i),w}); \quad i = 1, 2, 3 \quad (9)$$

and the entries of the estimated 3×3 covariance matrix \hat{V}_j^w are

$$\hat{V}_j^w(i, k) = \frac{1}{n_y - 1} \sum_{t=t_1}^{t_{n_y}} \left[\left(\Delta\hat{\kappa}_{t,j}^{(i),w} - \hat{\mu}_j^w(i) \right) \left(\Delta\hat{\kappa}_{t,j}^{(k),w} - \hat{\mu}_j^w(k) \right) \right]; \quad i, k = 1, 2, 3. \quad (10)$$

The corresponding estimators for μ and V for the true trajectory θ_0 are defined similarly.

4.1.2 Bayesian Estimation – Parameter Uncertainty

As mentioned earlier we model uncertainty about the parameters μ and V by applying a Bayesian approach to estimation. We denote by p the density of the prior joint distribution of the two parameters. Assuming that we have no prior knowledge about the true values of μ and V , we use the Jeffreys prior density

$$p(\mu, V) \propto |V|^{-\frac{3}{2}},$$

where $|V|$ is the determinant of V (see for example, Gelman et al. (1995)). Using this prior distribution in each scenario j , the posterior distribution is given by the inverse Wishart distribution for V and a multivariate normal distribution for μ , that is,

$$\left(\tilde{V}_j^w\right)^{-1} |\Delta \hat{\kappa}_j^w \sim \text{Wishart}(n_y - 2, (n_y - 1)^{-1}(\hat{V}_j^w)^{-1}) \quad (11)$$

$$\tilde{\mu}_j^w | \tilde{V}_j^w, \Delta \hat{\kappa}_j^w \sim N(\hat{\mu}_j^w, (n_y - 1)^{-1} \tilde{V}_j^w) \quad (12)$$

where $\hat{\mu}_j^w$ and \hat{V}_j^w are the estimates obtained from $\hat{\theta}_j^w$ as defined in (9) and (10).

4.1.3 Empirical comparison

For our empirical study we simulate $N = 1000$ scenarios for different values of w and plot the empirical density of the point estimator $\hat{\mu}^w$ in (9) based on the sample $\hat{\mu}_1^w, \dots, \hat{\mu}_N^w$ on the left hand side of Figure 3. To incorporate parameter uncertainty we draw a further sample of size $M = 100$ from the posterior distribution of $\tilde{\mu}_j^w$ in (12) in each scenario $j = 1, \dots, N$. The empirical density of $\tilde{\mu}_j^w$ from these $N \times M$ realisations is shown on the right hand side of Figure 3.

By comparing the densities in the two columns of that figure we observe that the additional parameter uncertainty increases the variance of the empirical distributions of the drift estimators. This can be explained by investigating the source of uncertainty to the drift. The variation to the point estimator $\hat{\mu}^w(i)$ with no allowance for parameter uncertainty comes from the Poisson noise in the number of deaths from the bootstrap simulations, while the variance of the Bayesian estimator $\tilde{\mu}^w(i)$ with allowance for extra parameter uncertainty also includes the uncertainty (equation 12) from the posterior distribution given the Poisson noise.

We also find in Figure 3 that the size of a population affects the uncertainty about the drift vector μ . The variance of the empirical finite sample distribution of both estimators, $\hat{\mu}$ and $\tilde{\mu}$ decreases significantly when the population size increases, although the difference between $w = 1$ and $w = 0.01$ is rather small as is particularly obvious for the Bayesian estimator $\tilde{\mu}$.

However, for smaller values of w we find that the population size has a much more pronounced effect on the variance. For example, the range of likely values of $\tilde{\mu}^{0.001}$ is significantly wider than the range of values of $\tilde{\mu}^{0.1}$ and $\tilde{\mu}^1$ reflecting the uncertainty about μ^w that we have already observed in Figure 2 top left). The same argument applies to the point estimators $\hat{\mu}$.

To investigate parameter uncertainty further we calculate the standard deviations for the distributions of $\hat{\mu}$ and $\tilde{\mu}$ in Figure 3. Those standard deviations are shown in Table 1. We observe that the standard deviation of the point estimator $\hat{\mu}$ is increased

approximately by a factor $\sqrt{10}$ if the population size is reduced by a factor 10. The situation becomes more complicated when for the Bayesian estimator $\tilde{\mu}$ since the variance of the posterior distribution affects the finite sample variance of the estimator. There is no obvious proportional relationship between population size and variation, which suggests that the size of the population is not the only determinant of the variance of $\tilde{\mu}$.

		$i = 1$	$i = 2$	$i = 3$
Point estimator $\hat{\mu}^w(i)$	w=1	0.0000966	0.0000071	0.00000113
	w=0.1	0.0003050	0.0000217	0.00000343
	w=0.01	0.0009777	0.0000727	0.00001068
	w=0.001	0.0028787	0.0002206	0.00003387
Bayesian estimator $\tilde{\mu}^w(i)$	w=1	0.00369	0.000173	0.00000936
	w=0.1	0.00396	0.000222	0.0000162
	w=0.01	0.00620	0.000505	0.0000458
	w=0.001	0.01689	0.001566	0.0001478

Table 1: The finite sample standard deviation of $\hat{\mu}$ and $\tilde{\mu}$.

To investigate the impact of the relative population size w and the inclusion of parameter uncertainty on the empirical distribution of the estimated co-variance matrix V of the random walk in (8) we compare the empirical means of \hat{V} and \tilde{V} obtained for different values of w . The means of the estimated co-variance matrix \hat{V} are:

$$E[\hat{V}^1] = \begin{pmatrix} 6.82 \times 10^{-4} & 2.12 \times 10^{-5} & 5.42 \times 10^{-7} \\ 2.12 \times 10^{-5} & 1.41 \times 10^{-6} & 2.99 \times 10^{-8} \\ 5.42 \times 10^{-7} & 2.99 \times 10^{-8} & 4.30 \times 10^{-9} \end{pmatrix}$$

$$E[\hat{V}^{0.01}] = \begin{pmatrix} 18.7 \times 10^{-4} & -1.52 \times 10^{-5} & 4.61 \times 10^{-6} \\ -1.52 \times 10^{-5} & 1.25 \times 10^{-5} & -1.89 \times 10^{-7} \\ 4.61 \times 10^{-6} & -1.89 \times 10^{-7} & 0.99 \times 10^{-7} \end{pmatrix}$$

and the mean values of the Bayesian estimator \tilde{V} are

$$E[\tilde{V}^1] = \begin{pmatrix} 7.58 \times 10^{-4} & 2.37 \times 10^{-5} & 6.06 \times 10^{-7} \\ 2.37 \times 10^{-5} & 1.58 \times 10^{-6} & 3.35 \times 10^{-8} \\ 6.06 \times 10^{-7} & 3.35 \times 10^{-8} & 4.79 \times 10^{-9} \end{pmatrix}$$

$$E[\tilde{V}^{0.01}] = \begin{pmatrix} 20.90 \times 10^{-4} & -1.74 \times 10^{-5} & 5.11 \times 10^{-6} \\ -1.74 \times 10^{-5} & 1.39 \times 10^{-5} & -2.10 \times 10^{-7} \\ 5.11 \times 10^{-6} & -2.10 \times 10^{-7} & 1.11 \times 10^{-7} \end{pmatrix}.$$

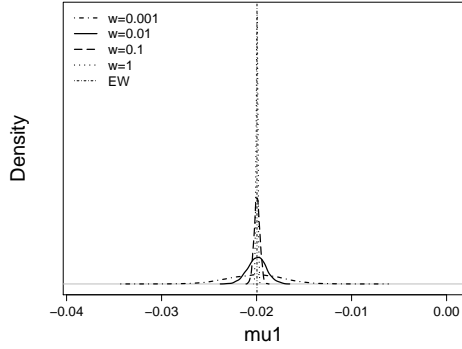
The corresponding estimated covariance matrices, V^{EW} , for England and Wales based on the single sample paths of κ_t and γ_c and the mean of Bayesian estimator are

$$\hat{V}^{EW} = \begin{pmatrix} 6.70 \times 10^{-4} & 2.16 \times 10^{-5} & 4.94 \times 10^{-7} \\ 2.16 \times 10^{-5} & 1.31 \times 10^{-6} & 3.18 \times 10^{-8} \\ 4.94 \times 10^{-7} & 3.18 \times 10^{-8} & 3.30 \times 10^{-9} \end{pmatrix}$$


$$E[\tilde{V}^{EW}] = \begin{pmatrix} 5.49 \times 10^{-4} & 1.80 \times 10^{-5} & 1.05 \times 10^{-7} \\ 1.80 \times 10^{-5} & 1.07 \times 10^{-6} & 2.19 \times 10^{-8} \\ 1.05 \times 10^{-7} & 2.19 \times 10^{-8} & 3.06 \times 10^{-9} \end{pmatrix}$$

Comparing the mean values of \hat{V} and \tilde{V} with the estimates obtained from the England and Wales data we find significant differences in the estimated covariances. In particular,


(a) Density of $\mu_w^{(1)}$ with No Allowance to Parameter Uncertainty



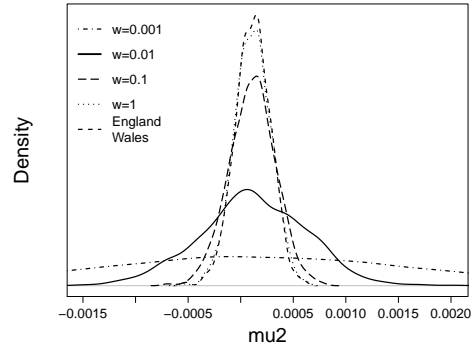
(b) Density of $\mu_w^{(1)}$ with Allowance to Parameter Uncertainty




(c) Density of $\mu_w^{(2)}$ with No Allowance to Parameter Uncertainty



(d) Density of $\mu_w^{(2)}$ with Allowance to Parameter Uncertainty



(e) Density of $\mu_w^{(3)}$ with No Allowance to Parameter Uncertainty



(f) Density of $\mu_w^{(3)}$ with Allowance to Parameter Uncertainty

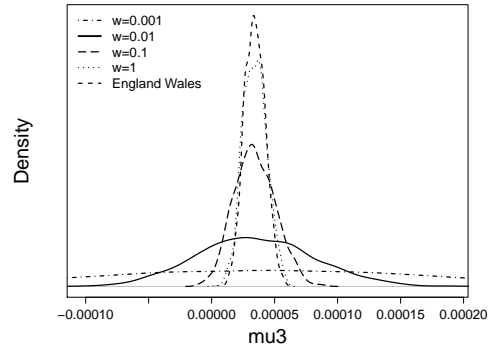


Figure 3: The impact of population size on the distribution of the random walk drift, from population of $w = 1$ (dotted line), $w = 0.1$ (long-dashed line), $w = 0.01$ (solid line), $w = 0.001$ (dot dashed line) and England and Wales (vertical line). The left column is the density of drift without allowance to the parameter uncertainty; the right column is the density of drift with allowance to the parameter uncertainty.

for smaller populations (e.g. $w=0.01$) sampling variation pushes up significantly estimates of the covariance matrix. In addition, sampling variation also widens the distribution of V around these mean values for smaller values of w . On the other hand, for a given value of w , the inclusion of full Bayesian parameter uncertainty moving from \hat{V} to \tilde{V} has rather less of an impact.

Finally, the projected parameters based on the Bayesian estimates $\tilde{\mu}$ and \tilde{V} are shown in Figure 6. As we expected, the prediction intervals reflecting the uncertainty about future values of the period effects are very wide for small populations. The plots also suggest that the mean of the co-variances are right biased from the England and Wales. The variance of projection for all the populations are much higher than the estimates, due to the additional normal randomness added in the forecasting model. However, the left column shows that there is no obvious proportional relationship between the population size and projection variance. By investigating the mean co-variance matrices, we find that the increase of $E[V_{3,3}^w]$ from $w = 0.01$ to $w = 1$ is of the highest among the three period effects, which suggests that the standard deviation of projection for $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ is not as sensitive as $\kappa_t^{(3)}$ to the change of population size.

4.2 Projecting the Cohort Effect

As mentioned earlier we fit an AR(1) model to the cohort effect. We will not investigate how additional parameter uncertainty influences mortality projections, but will only use point estimates for the parameters in the AR(1) model. To be precise, our model is given by

$$\gamma_{c+1}^{(4)} = \alpha_0 + \alpha_\gamma (\gamma_c^{(4)} - \alpha_0) + \epsilon_{c+1} \quad (13)$$

Figure 2 shows that the variance of the estimated cohort effect is very large for the very early and very late years of birth, in particular, for $w = 0.01$ and 0.001 . This is a consequence of the very few observations available for those cohorts. We therefore remove the cohorts with six or less observations.

However, removing short cohorts could significantly influence the estimated values of the parameters. To investigate the effect of removing short cohorts in more detail we plot the empirical densities for the parameters in (13) based on the estimated parameters in each simulated scenario j for $w = 1$. We find that the distribution of $\hat{\alpha}_0$ is not significantly affected by removing cohorts which is also the case for the estimated variance of ϵ_c when more than 4 cohorts removed. In contrast, $\hat{\alpha}_\gamma$ is shifted to the left as more cohorts are removed. Further we notice that the variance of the estimators for all three parameters stays approximately unchanged regardless of how many cohorts are removed.

After having removed cohorts with six or less observations from the data, we fit the AR(1) model in (13) to the rest of the cohort effects. The resulting density of the parameter estimates of the model are shown in Figure 5. All of the parameter estimates and the standard deviation of error terms appear to be biased relative to the estimate for England and Wales, regardless of the size of population. However, we find that reducing the population size will greatly increase the mean bias as well as the uncertainty.

We now forecast the cohort effect from $\hat{\gamma}_{t_{ny}-x_1-6}^{(4),w} = \hat{\gamma}_{1955}^{(4),w}$ instead of $\hat{\gamma}_{1961}^{(4),w}$ and the result is shown in Figure 6d. The variation in the projected cohort effects for the years 1956 to 1961 now comes from the Poisson and Normal randomness, which is not as great as variation at the two tails of the estimates observed in Figure 2 where no cohorts have been removed. Within the sample, the confidence intervals are narrower for cohorts with

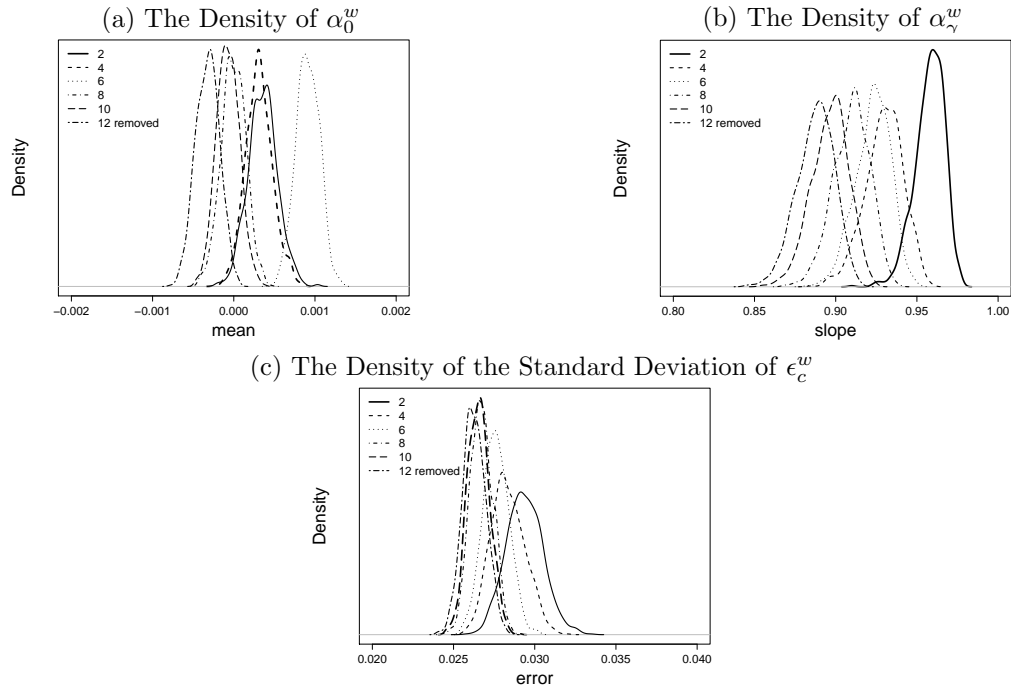


Figure 4: The effect of removing the cohort effects of short cohorts on the distribution of the parameter estimators of the $AR(1)$ model for $w = 1$. We investigate the distribution of parameter estimates when the first and last 1, 2, 3, 4, 5, 6 cohorts are removed

greater numbers of observed years (ranging from 7 to 40) and greater numbers of deaths since variance is reduced by having more number of observations.

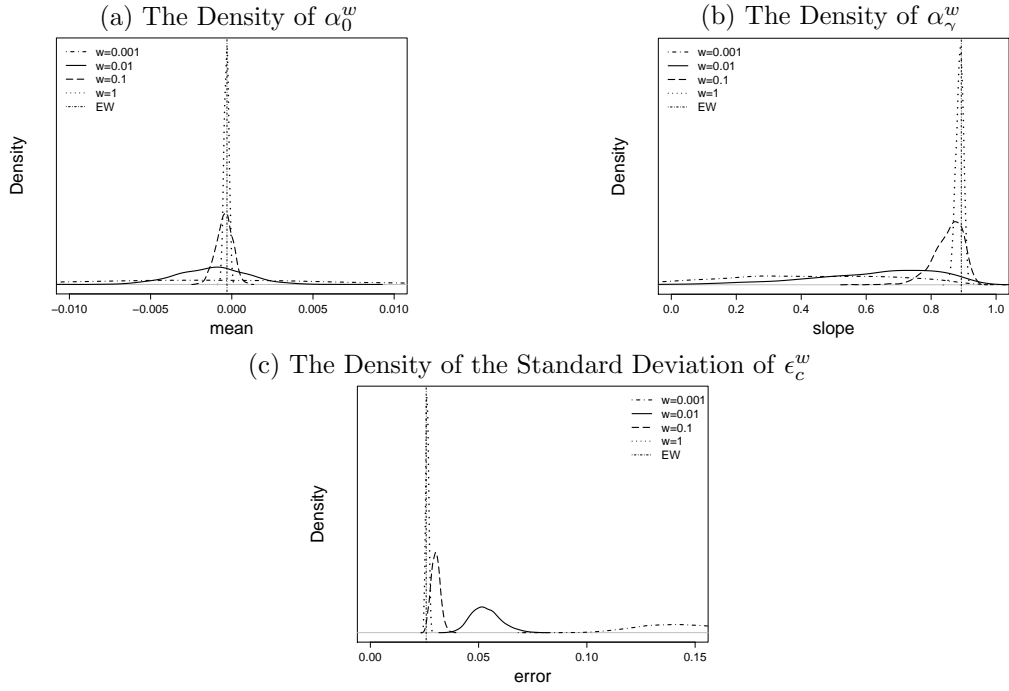


Figure 5: The comparison of the distribution of the parameter estimates of the $AR(1)$ model between the constructed populations $w = 1$ (dotted line), $w = 0.1$ (long dashed), $w = 0.01$ (solid), $w = 0.001$ (dot dashed) and England and Wales (vertical line)

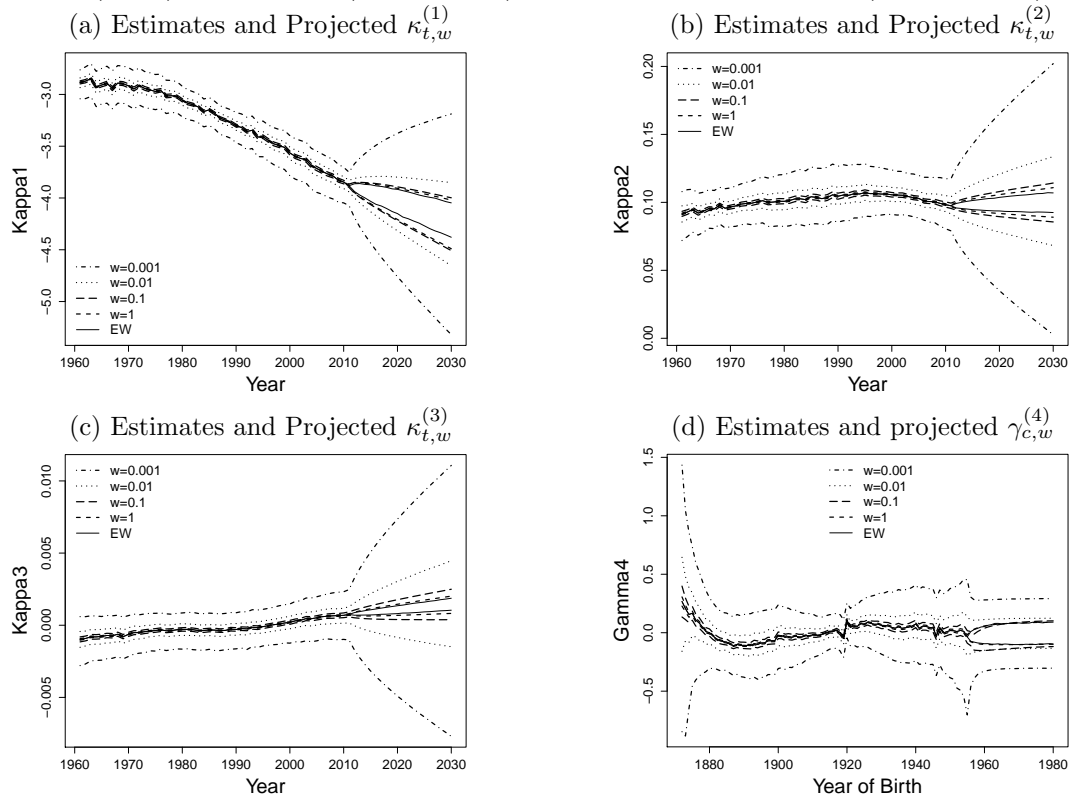


Figure 6: The comparison of twenty-year forward projection of $\kappa_{t,w}^{(1)}$, $\kappa_{t,w}^{(2)}$, $\kappa_{t,w}^{(3)}$, $\gamma_{c,w}^{(4)}$, of weight $w = 1, 0.1, 0.01, 0.001$ with England and Wales. Note: We forecast the cohort effect from the last sixth cohort instead of the very last one due to the cohort removal. The upper bound of the CI is the 95% quantile of the distribution and lower bound is the 5% quantile. Parameter uncertainty is allowed in the projection.

4.3 Projected Mortality Rates

Based on the projected period and cohort effects we can now turn to the projection of mortality rates using our model in (1)–(3). Figure 7 shows the twenty-year forward projections of mortality rates at ages 65 and 85. We compare the predicted rates with and without the allowance for parameter uncertainty for all the constructed populations with the projections based on the England and Wales data. Unsurprisingly, the uncertainty about future mortality rates increases as the forecast horizon increases. The other two factors which significantly influence the projection uncertainty are age and population size.

It seems that the projection variation for large populations is more effected by the age than the variation of projected rates for smaller populations. The variance of the two larger populations ($w = 1, 0.1$) and England and Wales increases as the age increases from 65 to 85. At age 65 (and also at age 85), these three populations have prediction intervals which are of similar width. However, the much wider prediction intervals for the two smaller populations seem to be less affected by age.

Reducing the population size results in greater uncertainty about mortality forecasts for both ages. For example, the uncertainty is much greater for the smaller populations ($w = 0.01, 0.001$) at both ages 65 and 85. This means that there is considerable uncertainty about future mortality scenarios for a relatively small pension scheme with significant implications for the risk management of such a scheme.

Comparing parts (a) and (b) of Figure 7 we find that the inclusion of parameter uncertainty for the drift parameter μ adds further uncertainty about the projected mortality rates. This reflects the additional randomness from not having a sufficiently long period of observed rates. We notice that the difference of variance between including and excluding parameter uncertainty increases as time increases. Thus parameter uncertainty becomes much less important when only relatively short forecast horizons are considered. A similar results can be found in Figure 6 in Cairns et al. (2006) which shows the log scaled variance of both with and without parameter uncertainty for the survival index. This supports our conclusion that the differences in the variances are tiny when t is very small, and become more significant for long term projection.

We are also interested in how much of the forecast variation is due to the impact of sampling variation and parameter uncertainty on the covariance matrix, V , and the drift, μ . To investigate this, we consider four experiments outlined in Table 2. Note that we still projected the cohort effect, given the point estimates for population w with the method introduced in Section 4.2 and we sample from the empirical distribution of $\hat{\mu}$ and \hat{V} without considering the Bayesian posterior. The results are shown in Figure 8. We find

1. Project mortality rates for each constructed population, while fixing the parameters μ and V of the random walk to the estimates obtained from the England and Wales data.
2. Project mortality rates for each constructed population, while fixing only the drift μ to the corresponding EW estimates and sample realisations of \hat{V} from its empirical distribution.
3. Project mortality rates for each constructed population, while fixing only the variance matrix V to the corresponding EW estimates and sample the drift parameter from the empirical distribution of $\hat{\mu}$.
4. Project mortality rates when both V and μ are samples from the empirical distribution of \hat{V} and $\hat{\mu}$.

Table 2

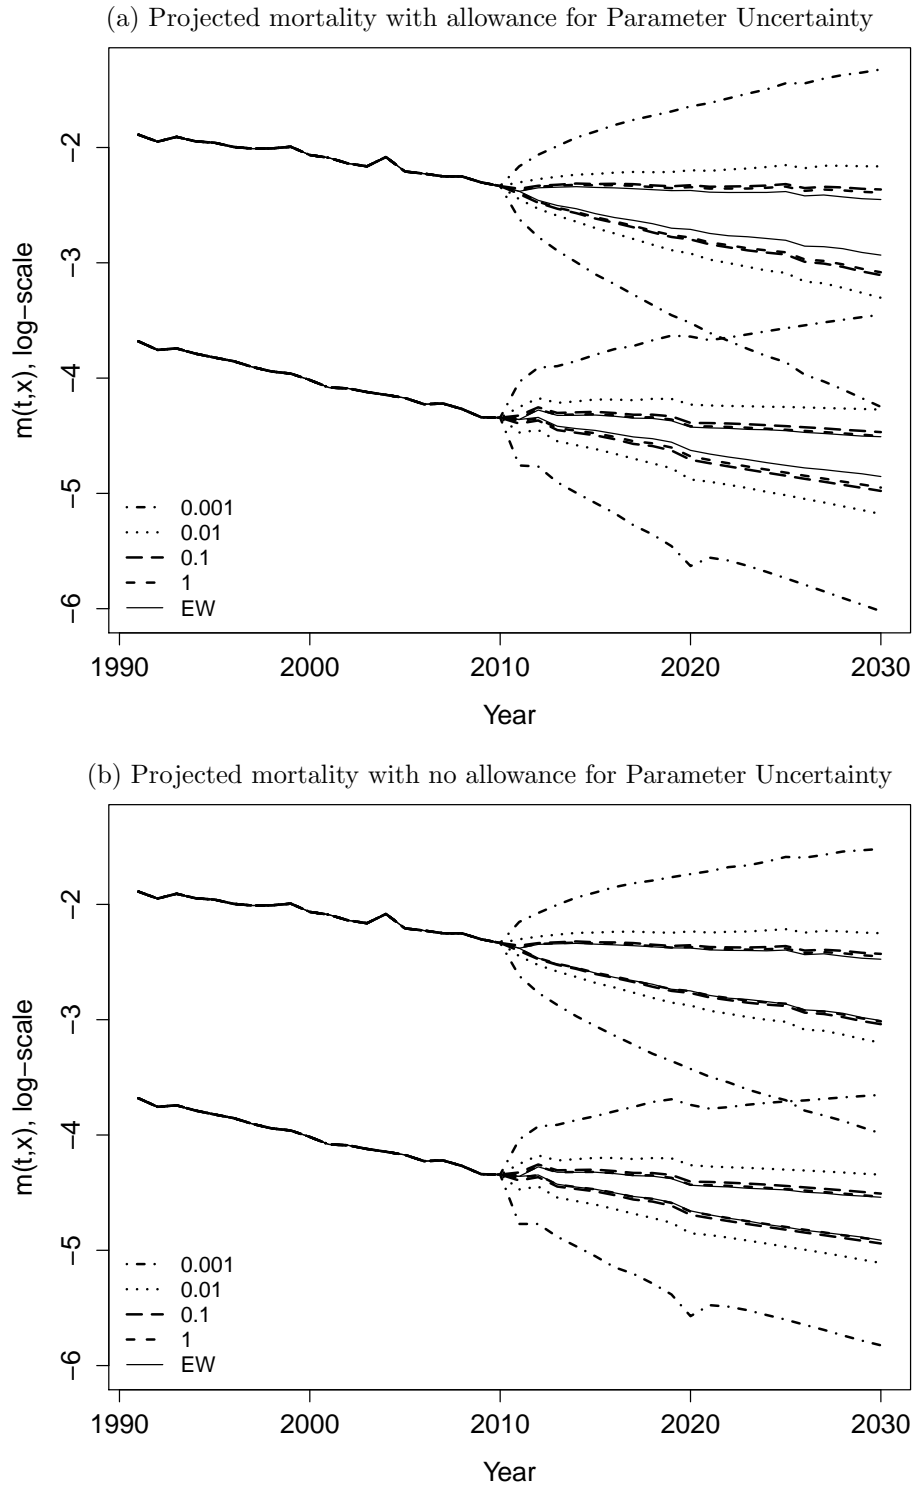


Figure 7: The log-scaled prediction intervals of twenty-year forward mortality rate projections with (upper plot) and without (lower plot) allowance for parameter uncertainty at ages 65 and 85, for population size $w = 1$ (dashed line), 0.1 (long dashed line), 0.01 (dotted), 0.001 (dot dashed line) and England and Wales (solid line). Note that the solid line at the left end is the estimated mortality rate of the England and Wales population, with length of 20 years. The upper bound of the prediction interval is the 95% quantile of the distribution and lower bound is the 5% quantile.

that fixing parameters has a significant effect on mortality forecasting when populations are very small ($w=0.001$) in Figure 8a. We can see that the widths of prediction intervals for our experiments 1 and 3 are much narrower than for experiments 2 and 4, and the difference of variance is greater for long term projections. The major difference between these two scenarios is that we fix the co-variance matrix V to its estimate obtained from England and Wales data in experiments 1 and 3. Thus we conclude that a major source of uncertainty for our mortality forecasts comes from the bias in the estimated covariance matrix for small populations.

4.4 Summary

To summarise, forecasts levels of uncertainty in future mortality are biased upwards for two reasons. First, and most obvious, the Poisson noise in the data biases up estimates of the random walk covariance matrix to a significant extent (Figure 8). Second, when we include a Bayesian analysis of parameter uncertainty, uncertainty in the random walk drift resulting from observation over a relatively small number of years is pushed up by the small population bias in the covariance matrix, V . This has its greatest impact in longer term forecasts, and less impact in the short term.

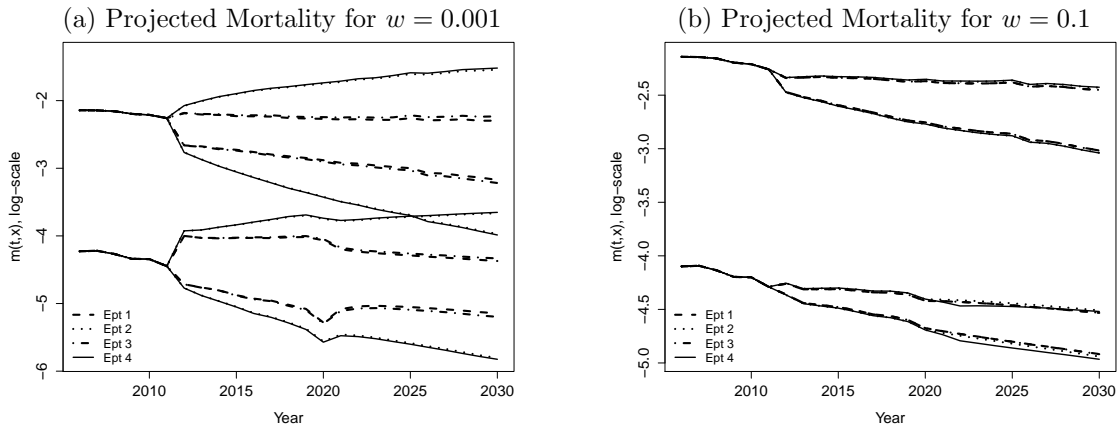


Figure 8: The projected mortality rates at age 65 and 85 for population sizes $w = 0.001$ (left), $w = 0.1$ (right) for the four experiments outlined in Table 2. The upper bound of the prediction interval is the 95% quantile of the distribution and the lower bound is the 5% quantile. Note that the solid line at the left end is the true mortality rate of the England and Wales population, up to year 2011.

5 Likelihood Ratio Test for Systematic Parameter Difference

We have seen that the size of a population has a substantial impact on the level of uncertainty about the parameters of the model in (1)–(3) when this model is fitted to the population’s mortality data. This raises the question whether the estimated period and cohort effects in $\theta = (\kappa_t^{(1)}, \kappa_t^{(2)}, \kappa_t^{(3)}, \gamma_c^{(4)})$ for a small a population are significantly different from those in a given, typically much larger, reference population. To address this question we apply a likelihood ratio test to test for significant deviations of estimated parameters from a given null hypothesis using the maximum likelihood estimator $\hat{\theta}_j^w$ defined in (7) for simulated mortality data D_j^w as in (6). We are particularly interested in the finite sample distribution of the test statistic as compared to its asymptotic distribution. As in section 3 we will use simulated deaths scenarios to investigate the finite sample distribution and the power of the likelihood ratio test (LRT) applied to mortality data. We will start with a short review of the LRT.

5.1 Review of Likelihood Ratio Test

The LRT used in this study follows the generalized form of the LRT as defined in Kendall et al. (1987). For a random variable X with a distribution that depends on a parameter vector θ , the likelihood function is defined as usual:

$$L(x|\theta) = \prod_{i=1}^n f_i(x_i|\theta)$$

where $f_i(\cdot|\theta)$ is the probability density function of X_i given the parameter vector θ . We assume that $\theta = (\theta_r, \theta_s)$ is a vector of $r + s$ parameters. The null hypothesis and alternative for the LRT concern only the parameters in θ_r , that is,

$$H_0 : \theta_r = \theta_{r0}; \quad H_1 : \theta_r \neq \theta_{r0}. \quad (14)$$

In order to calculate the test statistic, we first find the MLE of (θ_r, θ_s) , which leads to the unconditional maximum of the likelihood function

$$\hat{\theta} = (\hat{\theta}_r, \hat{\theta}_s) = \arg \max_{(\theta_r, \theta_s)} L(x | \theta_r, \theta_s). \quad (15)$$

We then find the MLE of θ_s assuming that the null hypothesis is fulfilled, that is,

$$\tilde{\theta}_s = \arg \max_{\theta_s} L(x | \theta_{r0}, \theta_s). \quad (16)$$

In general $\tilde{\theta}_s \equiv \tilde{\theta}_s(\theta_{r0}) \neq \hat{\theta}_s$.

We now define the test statistic in the usual way:

$$\Gamma = -2 \log \frac{L(x | \theta_{r0}, \tilde{\theta}_s)}{L(x | \hat{\theta}_r, \hat{\theta}_s)} \quad (17)$$

Wilks (1938) proved that when H_0 holds, Γ asymptotically follows a central χ^2 distribution with r degrees of freedom. From the central limit theorem, it follows that the χ_r^2 distribution can be approximated by a normal distribution with mean r , given r is

sufficiently large.² Thus we expect that the distribution of Γ should approximately be symmetric around r .

In the remainder of this section we will consider a null hypothesis about the entire parameter vector θ setting $s = 0$. In section 7 we will then consider a null hypothesis about the cohort effect γ only, that is $s > 0$.

5.2 Finite Sample Distribution of LRT

As mentioned above, we now consider a test for systematic parameter differences involving all period effects and the cohort effect, that is, $s = 0$ and $\theta = \theta_r = (\kappa_t^{(1)}, \kappa_t^{(2)}, \kappa_t^{(3)}, \gamma_c^{(4)})$. The null hypothesis and alternative are given in (14), and the LRT statistic is defined in (17) which simplifies to

$$\Gamma = -2\log \frac{L(x | \theta_{r0})}{L(x | \hat{\theta}_r)} \quad (18)$$

since $s = 0$.

As in section 3, we choose the male population in England and Wales as our base case and set $\theta_0 = \hat{\theta}^{\text{EW}}$.

To investigate the finite sample properties of the LRT in small populations we apply a parametric bootstrap procedure in which we simulate N mortality scenarios, estimate the parameter vector θ as in section 3 and apply the LRT in each scenario. More precisely we use the following steps to find a bootstrap approximation of the finite sample distribution of Γ : For different values of w and for each scenario $j = 1 \dots N$ we

1. simulate D_j^w as in (6),
2. find the estimate $\hat{\theta}_j^w$ as in (7),
3. calculate the realisation of the LRT statistic Γ_j^w as in (18) and
4. calculate the p -value P_j^w based on the asymptotic χ^2 -distribution as $P_j^w = \mathbf{P}[X > \Gamma_j^w]$ where X is has χ^2 -distribution with α degrees of freedom.

The degrees of freedom of the χ^2 distribution in step 4 should be the effective number of parameters denoted by α , which is the total number of parameters r less the number of constraints, that is

$$\alpha = 3n_y + n_c - 3$$

where n_y is the number of years, and $n_c = n_y + n_a - 1$ is number of cohorts in a given dataset. In our case, $n_y = 51$, $n_c = 51 + 40 - 1 = 90$, hence $\alpha = 240$. After applying the parametric bootstrap method we can generate the distribution of the test statistic. We expect that the distribution of Γ^w should be approximately symmetric around 240.

For any population size w we can now find the empirical distribution function of Γ^W based on the sample $\Gamma_1^w, \dots, \Gamma_N^w$. Furthermore, if the asymptotic χ^2 approximation is accurate, the p -values P_1^w, \dots, P_N^w should be independent and uniformly distributed on $[0, 1]$. The cumulative distribution of the test statistic Γ^w and the p -values P^w for all considered population sizes w are shown in Figure 9 for $N = 1000$. Figure 9a shows that the empirical distribution of Γ^w is indeed centred around $\alpha = 240$. We also observe in Figure 9b that the cumulative distribution function of the p -values resembles the

²See Kendall et al. (1987) for more details about the likelihood ratio test and the asymptotic distribution of the LRT statistic

distribution function of the uniform distribution on $[0, 1]$. Both results indicate that the χ^2 approximation for the distribution of Γ^w under the null hypothesis is very good for all values of w considered.




Figure 9: Likelihood ratio test: (a) empirical CDF's of test statistics for sample size $N = 1000$. (b) empirical CDF's of asymptotic p-values. Results shown for populations of $w = 1$ (solid line), $w = 0.1$ (dashed line), $w = 0.01$ (dotted line) and $w = 0.001$ (dot dashed line). The mean of the asymptotic χ^2_{240} distribution is also shown as the vertical dashed line in plot (a).

5.3 Power of the Likelihood Ratio Test

In last section, we carried out the likelihood ratio test for the parameter difference and found that the χ^2 approximation does not fail to capture the feature of the test statistic Γ^w when H_0 holds. We will now investigate how the population size affects the power of LRT.

As usual the power of a test is defined as

$$\text{Prob (Reject } H_0 \mid H_1 \text{ is True)}.$$

To evaluate the power of the LRT with a parametric bootstrap procedure similar to the one used in the previous section we need to generate scenarios under the alternative. So far we have considered a very general alternative $\theta_r \neq \theta_{r0}$. We will now need to specify this alternative further. To this end we define four alternative models and investigate the power assuming that the “true” data generating model is one of those alternatives. The four models we consider for the alternative shift or scale one of the period effects or the cohort effect estimated from the England and Wales data.

More specifically, the alternatives we consider are:

- $\theta^{(1)} = (\hat{\kappa}_0^{(1)} + \lambda, \hat{\kappa}_0^{(2)}, \hat{\kappa}_0^{(3)}, \hat{\gamma}_0^{(4)})$
- $\theta^{(2)} = (\hat{\kappa}_0^{(1)}, \hat{\kappa}_0^{(2)} + \lambda, \hat{\kappa}_0^{(3)}, \hat{\gamma}_0^{(4)})$
- $\theta^{(3)} = (\hat{\kappa}_0^{(1)}, \hat{\kappa}_0^{(2)}, \hat{\kappa}_0^{(3)} + \lambda, \hat{\gamma}_0^{(4)})$
- $\theta^{(4)} = (\hat{\kappa}_0^{(1)}, \hat{\kappa}_0^{(2)}, \hat{\kappa}_0^{(3)}, \lambda \hat{\gamma}_0^{(4)})$

We then evaluate the power of the LRT against each of those alternatives with different values of λ . Note that we scaled the cohort effect by λ units instead of shifting it since shifting the cohort effect would result in the same fitted mortality rates as shifting $\kappa_t^{(1)}$ in $\theta^{(1)}$.

We can now proceed as in the previous section with simulating death counts and then apply the LRT for different alternatives and different values of λ . We define $D_j^{w,(i)}(t, x)$ to be the simulated deaths in scenario $j = 1 \dots N$ for the population of size $wE_0(t, x)$ using the parameter $\theta^{(i)}$ in our model, that is,

$$D_j^{w,(i)}(t, x) \sim \text{Pois}(m(\theta^{(i)}, t, x)wE_0(t, x))$$

for any $i = 1, 2, 3, 4$, where m is defined in (2) and (3). Note that death counts also depend on λ .

Using the simulated death counts $D_j^{w,(i)}$ we obtain the MLE $\hat{\theta}_j^{w,(i)}$ as in (7). We then use the asymptotic χ^2 -distribution to test the null hypothesis that the parameters of our model are equal to the parameters obtained from the England and Wales populations. The p -values $P_j^{w,(i)} = P_j^{w,(i)}(\lambda)$ are then calculated as in step 4 in the previous section, and the null hypothesis is rejected in any scenario j for which $P_j^{w,(i)} < 0.05$, that is, the significance level of the test is 0.05.

The power of the LRT for any fixed alternative i , relative population size w and fixed λ is the proportion of the simulated p -values which are less than 0.05, that is, we count the number of scenarios for which the null hypothesis is rejected. More specifically, we define the random variables

$$\begin{aligned} R_j^{w,(i)}(\lambda) &= \begin{cases} 1 & \text{if } P_j^{w,(i)}(\lambda) < 0.05 \text{ (} H_0 \text{ rejected)} \\ 0 & \text{otherwise} \end{cases} \\ R^{w,(i)}(\lambda) &= \frac{1}{N} \sum_{j=1}^N R_j^{w,(i)}(\lambda) \end{aligned} \quad (19)$$

so that $R^{w,(i)}(\lambda)$ is the proportion of scenarios in which the null hypothesis is rejected among N simulated scenarios. We call $R^{w,(i)}(\lambda)$ the empirical rejection rate. Since we are considering independent scenarios, $R^{w,(i)}(\lambda)$ has a Binomial distribution,

$$NR^{w,(i)}(\lambda) \sim \text{Bin}\left(N, p^{w,(i)}(\lambda)\right) \quad (20)$$

where $p^{w,(i)}(\lambda)$ is the (unknown) power of the LRT if alternative $\theta^{(i)}$ with parameter λ is the true parameter set for the simulated death counts. Therefore, the empirical rejection rate $R^{w,(i)}(\lambda)$ is an unbiased estimator for the power $p^{w,(i)}(\lambda)$ and the estimated standard deviation of $R^{w,(i)}(\lambda)$ can easily be found from (20) in the usual way.

Then we investigate sensitivity of the power with respect to the size of λ and the size of population w . for each of the four cases, $\theta^{(1)}, \dots, \theta^{(4)}$, we consider a set of values for λ that are regularly spaced.

Figure 10 shows the obtained estimates $R_j^{w,(i)}(\lambda)$ for the power as a function of λ for different relative population sizes w . Note that for each alternative $\theta^{(i)}$ and any fixed λ we have simulated $N = 100$ scenarios, which is less than in the previous section. The reason is that we need to simulate those scenarios for each combination of i (alternative) and λ , which makes the total number of simulated scenarios very large.

Unsurprisingly, the power of the LRT is increasing in λ for any $\theta^{(i)}$ and relative population size w ; the more we shift/scale the null hypothesis, the easier it is for the

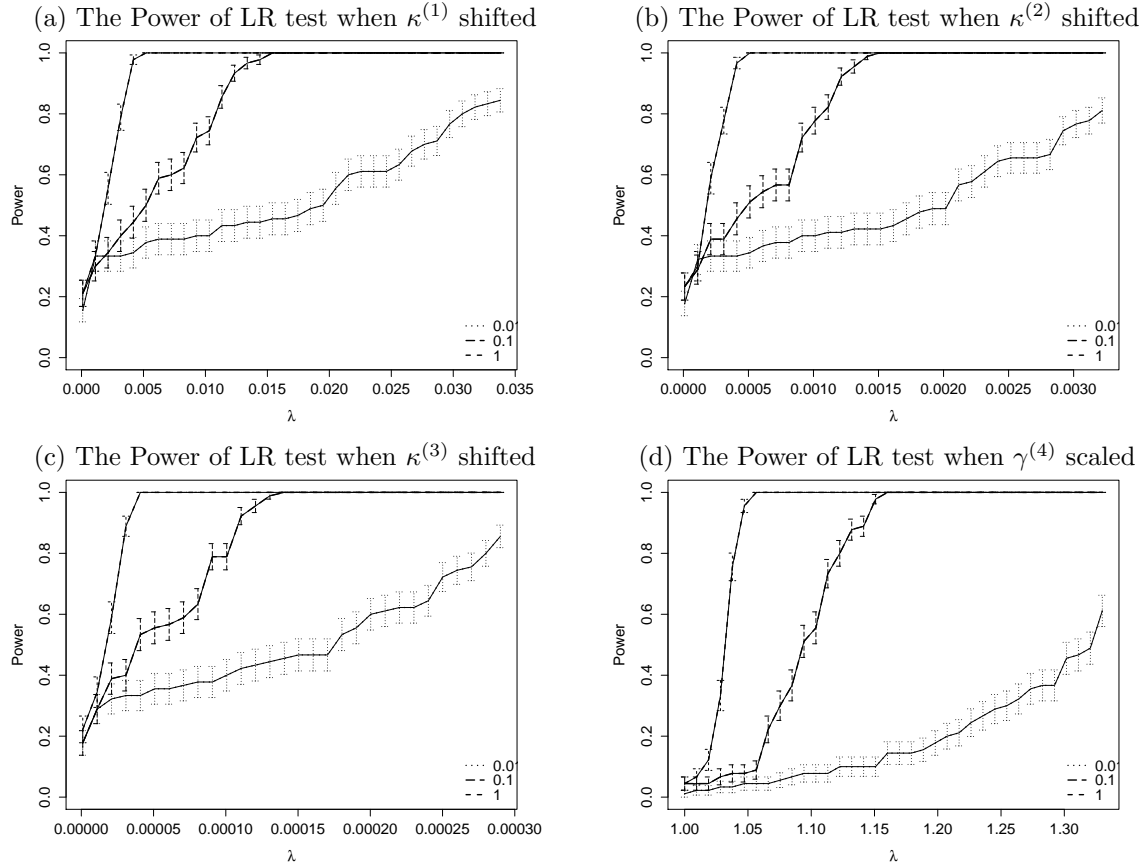


Figure 10: The empirical rejection rates $R^{w,(i)}(\lambda)$ under the LRT together with error bars for relative population sizes $w = 1$ (dashed line), $w = 0.1$ (long dashed line) and $w = 0.01$ (dotted line). The width of the error bars is one standard error based on (20).

test to detect any shift/scaling. For the three period effects, decreasing the population size will greatly reduce the capability of LRT to detect the same amount of shift to a single parameter. We can also compare these plots with the earlier Figure 2 which includes distributions of parameter estimates resulting from sampling variation. By way of example, for $w = 0.01$ the width of the confidence interval in Figure 2 (e) for $\kappa_{t,w}^{(3)}$ is about 0.005. This is much larger than the shifts that are considered in the power plot in Figure 10. The reason why the latter values are so much lower is because we apply a systematic adjustment to all of the $\kappa_{t,w}^{(3)}$, in contrast to random adjustments (due to sampling variation) in the former.

6 Impact of Parameter Misspecification on Mortality Rates and Annuities

We now investigate how significant the impact of shifting and scaling parameters is on the fitted mortality rates and corresponding annuity prices. We consider again the four alternatives in the previous section. For each of those and for each relative population size w we determine the value of λ that results in a power of 50% of the LRT, that is, there is a 50% probability that the LRT will detect the wrong model and reject the null hypothesis. Those values are denoted by $\lambda_{0.5}^{w,(i)}$ and shown in Table 3.

Parameter shifted	$w = 1$	$w = 0.1$	$w = 0.01$
$\lambda_{0.5}^{w,(1)}$	0.003	0.006	0.02
$\lambda_{0.5}^{w,(2)}$	0.0003	0.0006	0.002
$\lambda_{0.5}^{w,(3)}$	0.0000025	0.000005	0.00018
$\lambda_{0.5}^{w,(4)}$	1.03	1.09	1.32

Table 3: The table contains the size of shift required for 50% power when each parameter is shifted separately, with respect to population $w = 1, 0.1, 0.01$

We then calculate fitted mortality rates using the model in (2) and (3) with the following parameter constellations:

- $\theta_{0.5}^{w,(1)} = (\hat{\kappa}_0^{(1)} + \lambda_{0.5}^{w,(1)}, \hat{\kappa}_0^{(2)}, \hat{\kappa}_0^{(3)}, \hat{\gamma}_0^{(4)})$
- $\theta_{0.5}^{w,(2)} = (\hat{\kappa}_0^{(1)}, \hat{\kappa}_0^{(2)} + \lambda_{0.5}^{w,(2)}, \hat{\kappa}_0^{(3)}, \hat{\gamma}_0^{(4)})$
- $\theta_{0.5}^{w,(3)} = (\hat{\kappa}_0^{(1)}, \hat{\kappa}_0^{(2)}, \hat{\kappa}_0^{(3)} + \lambda_{0.5}^{w,(3)}, \hat{\gamma}_0^{(4)})$
- $\theta_{0.5}^{w,(4)} = (\hat{\kappa}_0^{(1)}, \hat{\kappa}_0^{(2)}, \hat{\kappa}_0^{(3)}, \lambda_{0.5}^{w,(4)} \hat{\gamma}_0^{(4)})$

To quantify the change in fitted mortality rates we calculate the following ratio

$$\rho^{w,(i)} = \frac{m(\theta_{0.5}^{w,(i)}, t, x)}{m(\hat{\theta}_0, t, x)}$$

for each $i = 1, 2, 3, 4$ and different values of w . We expect that shifting $\kappa_t^{(1)}$, $\kappa_t^{(2)}$ and $\kappa_t^{(3)}$ will result in a parallel shift upwards, tilting rates in an anti clockwise direction and add some concavity to the rates respectively. This can indeed be seen in Figure 11 where we plot the ratio $\rho^{w,(i)}$ for the year $t = 2011$. Figure 11d suggests that scaling $\gamma_c^{(4)}$ tilts and introduces more fluctuation to the ratio. For all the four parameters, reducing the relative population size w increased the relative change $\rho^{w,(i)}$ since $\lambda_{0.5}^{w,(i)}$ increases. This confirms the intuitive idea that even misspecified parameters which produce significant changes in the fitted mortality rates are hard to detect with an LRT when the exposures are small.

From a financial point of view the effect on fitted mortality rates is only relevant in so far as annuity prices are affected. We will therefore consider the following annuities and discuss the effect of the four alternatives specified above on their values:

- A temporary annuity of £1 per annum payable annually in arrears to a life now aged 65 exactly, starting at the beginning of year 2012 with term of 25 years. Its expected present value is calculated as:

$$a_{65:\overline{25}|} = \sum_{j=1}^{25} S(T + j, 65)v^j$$

- An annuity of £1 per annum payable annually in arrears to a life now aged 55 exactly, deferred for 10 years, starting at the beginning of year 2012 with term of 25 years. Its expected present value is:

$${}_{10|}a_{55:\overline{25}|} = \sum_{j=11}^{35} S(T + j, 55)v^j$$

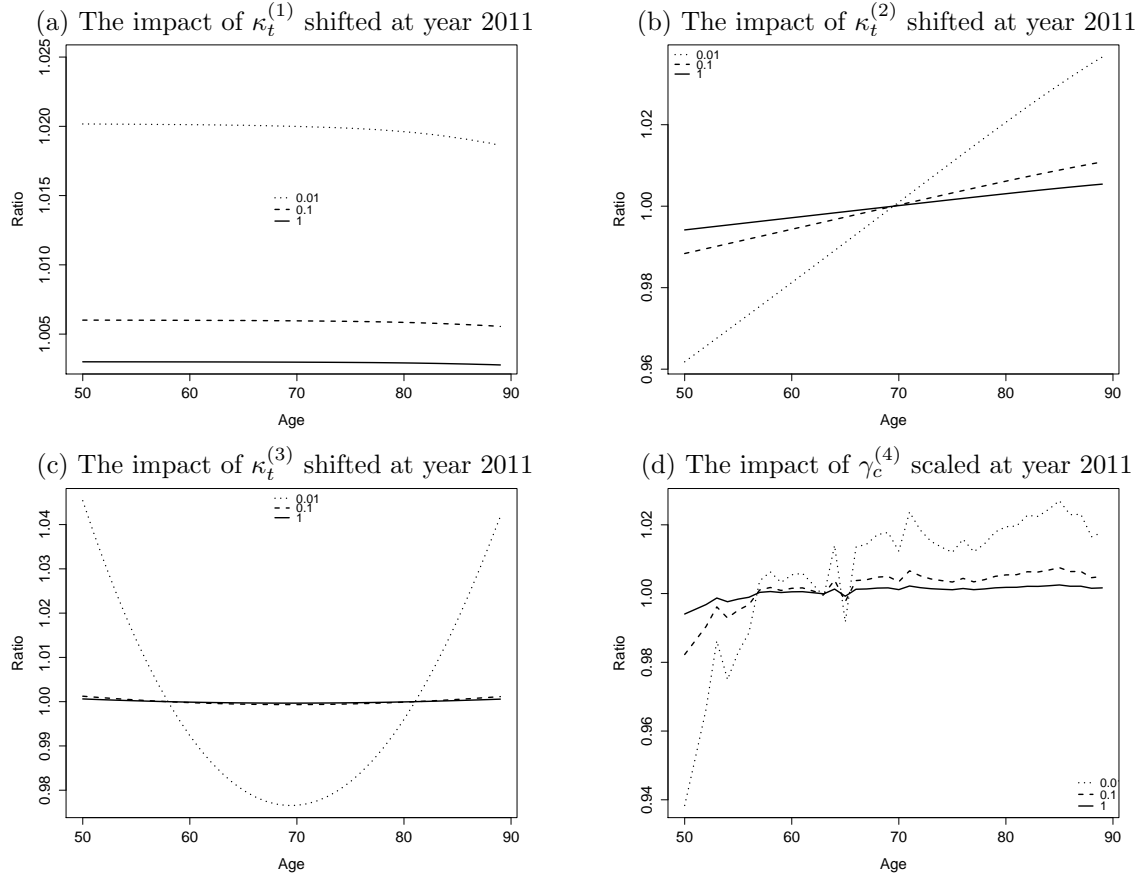


Figure 11: The impact of shifting each parameter separately on the estimated death rate of England and Wales. The shift is determined when it results in 50% power for each population $w = 1$ (solid line), $w = 0.1$ (dashed line) and $w = 0.01$ (dotted line).

where v is the discount factor, $S(T + t, x)$ is the survival index for the probability of an individual aged x exactly at the start of year T , that will survive for the next t years. We assume the interest rate of $i = 2\%$. The reason for investigating the deferred annuity is that Figure 2g suggests that the estimates of cohort effect at $c = 1946$ is approximately zero and the effect of scaling cohort estimates may not be obvious on the annuity price $a_{65:\overline{25}|}$ but more obvious for $10|a_{55:\overline{25}|}$.

We project the period and cohort effects in $\theta_{0.5}^{w,(i)}$ ($i = 1, 2, 3, 4$) and $\hat{\theta}^{EW}$ forward for 35 years as in Section 4 where we use the point estimates defined in (9) and (10) for the parameters of the random walk for the shifted period effects, that is, we do not consider uncertainty about the drift and variance matrix of the random walk. Annuity prices are calculated for each sample path and we then calculate the average annuity price for each w with the i^{th} parameter shifted or scaled. The results are shown in Table 4 and Table 5.

The effects of shifting the period effects and scaling the cohort effect are somewhat varied. As might be expected, the impact on prices is most obvious for $w = 0.01$. The impact on both types of annuity is straightforward to see for $\kappa^{(1)}$: the shift pushes up mortality rates at all ages and lowers prices. For $\kappa^{(2)}$ there is more impact on the age-65 annuity than the age-55 deferred annuity as the shift lowers mortality at younger ages and raises it at higher ages. For $\kappa^{(3)}$, also, the impact is different at different ages. Finally, for $\gamma^{(4)}$, the impact of scaling simply depends on the sign and magnitude of the value of

$\gamma^{(4)}$ for the cohort being priced.

Generally shifting or scaling the parameter estimates has no obvious effect on the annuity price and smaller populations can be affected more. Thus for testing a null hypothesis $H_0 : \theta_r^{w=0.01} = \hat{\theta}^{EW}$, if we accept H_0 when, in fact, they are actually different (type II error) the financial consequence of this type II error will be small in our case. In other words, the fact that we have accepted H_0 means that $\theta_r^{w=0.01}$, while not identical, must be very close to $\hat{\theta}^{EW}$, and that, therefore, any error in pricing will also be very small.

Parameter shifted	England and Wales	$w = 1$	$w = 0.1$	$w = 0.01$
$\kappa^{(1)}$	14.67466	14.66393	14.65318	14.60280
$\kappa^{(2)}$	14.67466	14.66887	14.66307	14.63588
$\kappa^{(3)}$	14.67466	14.67500	14.67534	14.69850
$\gamma^{(4)}$	14.67466	14.66997	14.66056	14.62441

Table 4: The impact of shifting each parameter separately on the price of a twenty five-year temporary annuity for an individual aged at 65. The shift is determined when it results in 50% power for each population $w = 1, 0.1, 0.01$, which are shown in Table 3. We assume an interest rate of 2%.

Parameter shifted	England and Wales	$w = 1$	$w = 0.1$	$w = 0.01$
$\kappa^{(1)}$	11.96545	11.95599	11.94652	11.84214
$\kappa^{(2)}$	11.96545	11.96358	11.96169	11.95266
$\kappa^{(3)}$	11.96545	11.96565	11.96584	11.97920
$\gamma^{(4)}$	11.96545	11.96815	11.97355	11.99411

Table 5: The impact of shifting each parameter separately on the price of a ten-year deferred twenty five-year temporary annuity for an individual aged at 55. The shift is determined when it results in 50% power for each population $w = 1, 0.1, 0.01$, which are shown in Table 3. We assume an interest rate of 2%.

7 Likelihood Ratio Test for the Cohort Effect

The general form of the LRT as reviewed in section 5.1 allows us to test a null hypothesis about parts of the parameter vector θ rather than the entire $\theta = (\kappa_t^{(1)}, \kappa_t^{(2)}, \kappa_t^{(3)}, \gamma_c^{(4)})$. Testing parts of θ is particularly relevant if mortality rates in a rather small population are modelled using estimated period or cohort effects from a larger population. Setting one or more of the components of θ equal to the corresponding parameters estimated from the large population reduces the dimension of the parameter vector which needs to be estimated from the small population where parameter uncertainty is rather strong as we have seen in section 3. The example we have in mind is a pension fund that uses national mortality data to improve its mortality models, or when the mortality experience in a small country is modelled based on the combined experience of other similar countries.

In the remainder of this section we will use the LRT to test a null hypothesis about the cohort effect γ . In our general setting of section 5.1 this means that

$$\theta_r = \gamma \text{ and } \theta_s = (\kappa_t^{(1)}, \kappa_t^{(2)}, \kappa_t^{(3)}).$$

Our null hypothesis is then that $\gamma = \gamma_0$ where γ_0 is a given vector of cohort effects, for which we later use an estimated cohort effect from a different population. We can now write the hypotheses as in (14) and proceed as in section 5.2 to find the distribution of the LRT statistic in (17) for a finite sample of death counts from a small population.

For practical relevance we base our simulation study on the female and male populations in England and Wales. We choose $\gamma_0 = \hat{\gamma}^{EW}$, which is the estimated cohort effect from the mortality data for males in England and Wales. To investigate finite sample properties of Γ we will need to specify a full parameter vector θ to simulate scenarios for the death counts. Having fixed the cohort effect γ_0 we choose the period effects to be the estimated period effects from data for the female population in England and Wales assuming that the cohort effect for those data is actually γ_0 . As we are mainly interested in small populations we will consider deaths count scenarios for populations which have exposures equal to wE_0 where E_0 is here the exposure for the female population in England and Wales.

More specifically, we first find the MLE $\tilde{\theta}_s = \arg \max_{\theta_s} L(x | \theta_{r0} = \gamma_0, \theta_s)$ of the period effects $\theta_s = (\kappa_t^{(1)}, \kappa_t^{(2)}, \kappa_t^{(3)})$ from data for females assuming that the cohort effect is indeed γ_0 (which is the estimated cohort effect for males), see (16). We then generate N realisations of the value of the test statistic Γ^w for different values of the relative population size w using the following algorithm:

1. Simulate death counts D_j^w as in (6) using the parameter vector

$$\tilde{\theta} = (\tilde{\theta}_s, \theta_{r0}) = (\tilde{\kappa}_t^{(1)}, \tilde{\kappa}_t^{(2)}, \tilde{\kappa}_t^{(3)}, \gamma_0)$$

to obtain scenarios D_j^w for different values of the relative population size w . The period effects $\tilde{\kappa}$ are estimated from data for females with the cohort effect fixed to γ_0 . The exposure is wE_0 where E_0 is the exposure for the female population in England and Wales.

2. Find the MLE $\tilde{\theta}_{s,j}$ of period effects κ in scenario j assuming that the null hypothesis holds, as in (16).
3. Find the unrestricted MLE $\hat{\theta}_j$ as in (15).
4. Calculate the value of the test statistic Γ_j^w in (17) in each scenario j .
5. Calculate the p -values P_j^w based on the asymptotic χ^2 -distribution with α degrees of freedom, where α is the number of parameters (cohorts) r minus the number of constraints as in section 5.2. For our data set we obtain $\alpha = 87$.

The simulated distribution functions of the LRT statistic Γ^w and the p -values P^w are shown in Figure 12. The results suggest that changing the size of the population has no significant impact on the distribution of Γ^w and that the p -values are roughly uniformly distributed for all w , which is an indication that the χ^2 -approximation works well for our data set as we have also found in section 5.2 where the full parameter vector was tested.

8 Empirical Examples

We apply the LRT for the cohort effect in two empirical studies.

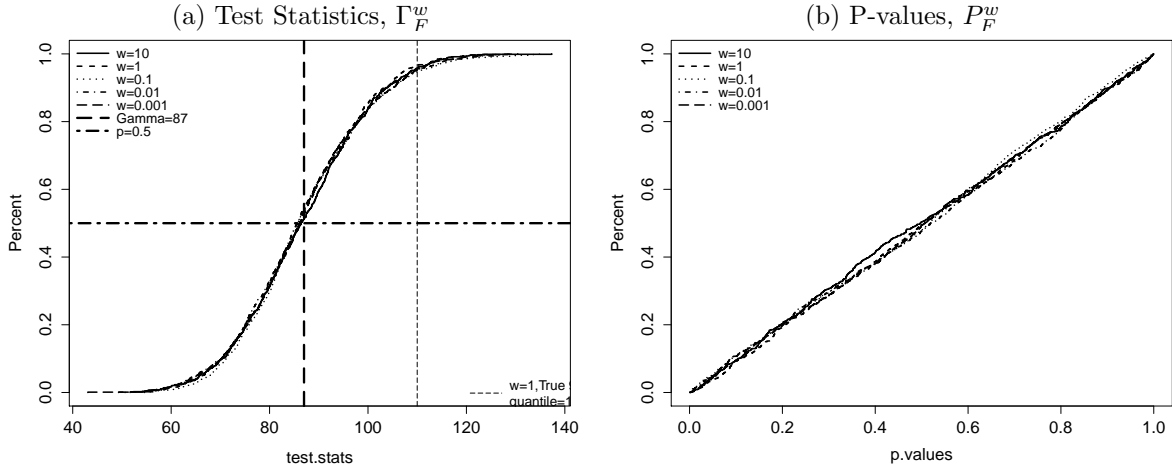


Figure 12: The results of likelihood ratio test, with distributions of test statistics (a) and p-values (b), for the population of $w = 10$ (solid line), $w = 1$ (dashed line), $w = 0.1$ (dotted line), $w = 0.01$ (dashed dotted line) and $w = 0.001$. The left vertical dashed line is the mean of normal approximation for the χ_{87}^2 , at $x = 87$. The right dashed line at $x = 110$ is the true 95% quantile of population $w = 1$.

8.1 Females vs. Males in England and Wales

The population for which we wish to test the cohort effect first is the female population in England and Wales that we already considered in our simulation study. Our null hypothesis is therefore that the true cohort effect for the female population in England and Wales is equal to the estimated cohort effect for males in England and Wales. Note that this is different from testing the hypothesis that the male and female population share the same (true) cohort effect since we ignore the uncertainty about the estimated cohort effect for males.

To illustrate the difference between the two cohort effects we plot in figure 13 the estimated cohort effects for females and males. We notice that males and females have a similar pattern after about 1900.

We also find visually in figure 13 that substantial differences between the two cohort effects exist. This difference can be confirmed more formally using the LRT with the null hypothesis that the females have the same cohort effect as the previously estimated males cohort effect. The test statistic Γ is approximately 6311, which is an extremely high value for a χ^2 -distribution with 87 degrees of freedom and is also very high compared to the values of Γ observed in our simulation study, see Figure 12. The p -value is therefore very close to zero, and we reject the null hypothesis that the cohort effect from the mortality of the female population is the same as the previously estimated cohort effect for the male population.

8.2 Male Mortality in Scotland vs. England and Wales

A second, and more intriguing, empirical example concerns the cohort effects estimated from mortality data for the male population in England and Wales versus the male population in Scotland. Figure 14 compares the independently-estimated cohort effects with a confidence interval added around the Scottish estimates. Compared to Figure 13,




Figure 13: The estimates of cohort effect, for England and Wales males (solid line) and females (dashed line), age 50 to 89 last birthday, over year 1961 to 2011

the two curves here look much more similar. On the other hand, we find that most of the cohort effects for males in England and Wales lie outside of the confidence interval calculated for Scottish males. This suggests that although the two populations have similar pattern for the cohort estimates, the difference might still be significant.

For the LRT we again choose $\gamma_0 = \hat{\gamma}^{EW}$ and then test the hypothesis that the true cohort effect for Scottish males is equal to γ_0 . The 99% quantile of a χ^2 -distribution with 87 degrees of freedom is approximately 121. For the test statistic we find $\Gamma = 193.37$ and we therefore reject the null hypothesis and conclude that the cohort effect in Scotland is significantly different from the estimated cohort effect for England and Wales. This indicates that there might be factors in the Scottish male population that result in significant differences throughout time. However, we might speculate that there is a common cohort effect, that is, for some reason, magnified in Scotland. Investigating this in detail is beyond the scope of this paper, but we speculate that a magnified effect might be the result of socio-economic differences between the two populations: for example, cohort effects might be greater in lower socio-economic groups.

9 Conclusion

In this paper, we investigated the finite sample distribution of the maximum likelihood estimators for the parameters of a stochastic mortality model. We found that the size of a population has a significant effect on the uncertainty about the estimated parameters and mortality projections. In particular, we found that there exists a bias in the estimated

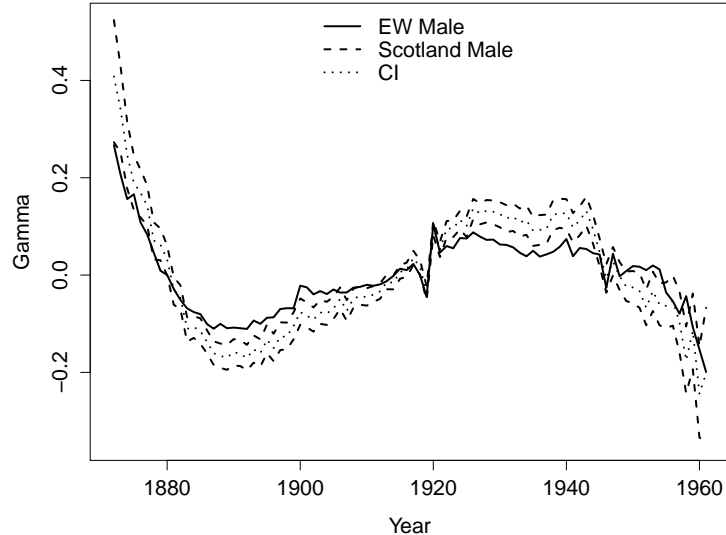


Figure 14: The estimates of cohort effect, for the males of England and Wales (solid line) and Scotland (dotted line), age 50 to 89 last birthday, over year 1961 to 2011. The dashed lines are the CI for the cohort effect of Scotland. The upper bound is 95% quantile of the distribution and the lower bound is 5% quantile.

covariance matrix of the random walk fitted to the period effects when the size of the underlying population is small. As a consequence, prediction intervals are rather wide for small populations even when parameter uncertainty is ignored.

To investigate if parameters estimated from larger populations can be used to generate scenarios for smaller populations we investigated how a likelihood ratio test performs when applied to the mortality experience of a small population. We found that the finite sample distribution of the test statistic is very close to the asymptotically correct χ^2 distribution and, therefore, the observed rejection rates are close to the chosen significance level. However, we also found that the power of the test depends strongly on the population size with the ability of the test to detect deviations from the null hypothesis being significantly reduced when the size of the underlying populations is small.

In our empirical analysis we then applied the LRT, and found that neither of the mortality rates of the female population in England and Wales and the male population in Scotland should be modelled with a cohort effect estimated from the male population in England and Wales.

10 Acknowledgements

Liang Chen is in receipt of an Actuarial Research Centre PhD scholarship funded by the Faculty of Actuaries Endowment Fund and the Institute and Faculty of Actuaries. Andrew Cairns and Torsten Kleinow acknowledge financial support from Netspar under project LMVP 2012.03.

A Derivation of the asymptotic distribution of θ_w

Recall the log likelihood function

$$l(\theta_w, D_{t,x}^w, E_{t,x}^w) = \sum_{t,x} f(t, x, \theta_w) - g(t, x, \theta_w) + h(t, x)$$

where

$$\begin{aligned} f(t, x, \theta_w) &= D_{t,x}^w \log[m(t, x, \theta_w)] \\ g(t, x, \theta_w) &= w E_{t,x}^w m(t, x, \theta_w) \end{aligned}$$

and $h(t, x)$ is the function with no respect to θ_w . The form of the column vector θ_w with $4ny + na - 1$ dimensions is

$$\theta_w = (\kappa_{t_1,w}^{(1)}, \dots, \kappa_{t_{ny},w}^{(1)}, \kappa_{t_1,w}^{(2)}, \dots, \kappa_{t_{ny},w}^{(2)}, \kappa_{t_1,w}^{(3)}, \dots, \kappa_{t_{ny},w}^{(3)}, \gamma_{c_1,w}^{(4)}, \dots, \gamma_{c_{ny+na-1},w}^{(4)})^T$$

The second derivative of $l(\theta_w, D_{t,x}^w, E_{t,x}^w)$ is

$$\frac{\partial^2 l}{\partial \theta_w^2} = \sum_{t,x} \frac{\partial^2 f}{\partial \theta_w^2} - \frac{\partial^2 g}{\partial \theta_w^2}$$

For every pair of (t, x) , $m(t, x, \theta_w)$ is a single value, thus the second derivative of f and g with respect to θ_w is a Hessian matrix with $4ny + na - 1$ rows and columns, with the form as

$$\frac{\partial^2 f}{\partial \theta_w^2} = \begin{pmatrix} \frac{\partial^2 f}{\partial \kappa_{t_1,w}^{(1)2}} & \cdots & \frac{\partial^2 f}{\partial \kappa_{t_1,w}^{(1)} \partial \kappa_{t_n,w}^{(i)}} & \cdots & \frac{\partial^2 f}{\partial \kappa_{t_1,w}^{(1)} \partial \gamma_{c_{ny+na-1},w}^{(4)}} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{\partial^2 f}{\partial \gamma_{c_{ny+na-1},w}^{(4)} \partial \kappa_{t_1,w}^{(1)}} & \cdots & \frac{\partial^2 f}{\partial \gamma_{c_{ny+na-1},w}^{(4)} \partial \kappa_{t_n,w}^{(i)}} & \cdots & \frac{\partial^2 f}{\partial \gamma_{c_{ny+na-1},w}^{(4)2}} \end{pmatrix}$$

Thus the form of the element at row m and column n is

$$\frac{\partial^2 f}{\partial \kappa_{m,w}^{(i)} \partial \kappa_{n,w}^{(j)}} = -\frac{D_{t,x}^w}{m(t, x, \theta_w)^2} \frac{\partial m(t, x, \theta_w)}{\partial \kappa_{m,w}^{(i)}} \frac{\partial m(t, x, \theta_w)}{\partial \kappa_{n,w}^{(j)}} + \frac{D_{t,x}^w}{m(t, x, \theta_w)} \frac{\partial^2 m(t, x, \theta_w)}{\partial \kappa_{m,w}^{(i)} \partial \kappa_{n,w}^{(j)}}$$

where $i, j \in (1, 2, 3, 4)$ and are not necessarily the same. Note that $\kappa^{(4)}$ represents the cohort effect for convenience. Same derivation can be done for g , and we have

$$\frac{\partial^2 g}{\partial \kappa_{m,w}^{(i)} \partial \kappa_{n,w}^{(j)}} = w E_{t,x}^w \frac{\partial^2 m(t, x, \theta_w)}{\partial \kappa_{m,w}^{(i)} \partial \kappa_{n,w}^{(j)}}$$

Thus for each pair of (t, x) , the expected value of the element of the second derivative of l is

$$E \left[\frac{\partial^2 (f - g)}{\partial \kappa_{m,w}^{(i)} \partial \kappa_{n,w}^{(j)}} \right] = -w \frac{E_{t,x}^w}{m(t, x, \theta_w)} \frac{\partial m(t, x, \theta_w)}{\partial \kappa_{m,w}^{(i)}} \frac{\partial m(t, x, \theta_w)}{\partial \kappa_{n,w}^{(j)}}$$

Thus we have the fisher information matrix given $\hat{\theta}^{EW}$ as

$$I(\hat{\theta}^{EW}) = w \sum_{t,x} \frac{E_{t,x}^{EW}}{m(t, x, \hat{\theta}^{EW})} \left[\frac{\partial m(t, x, \theta_w)}{\partial \theta_w} \left(\frac{\partial m(t, x, \theta_w)}{\partial \theta_w} \right)^T \right] \Bigg|_{\theta_w = \hat{\theta}^{EW}}.$$

Further, given $t = t_m, x = x_n$,

$$\frac{\partial m(t_m, x, \theta_w)}{\partial \theta_w} = \left(\frac{\partial m(t_m, x_n, \theta_w)}{\partial q(t_m, x_n, \theta_w)} \frac{\partial q(t_m, x_n, \theta_w)}{\kappa_{t,w}^{(1)}} \mathbb{1}_{t=t_m} \dots \frac{\partial m(t_m, x_n, \theta_w)}{\partial q(t_m, x_n, \theta_w)} \frac{\partial q(t_m, x_n, \theta_w)}{\kappa_{t,w}^{(4)}} \mathbb{1}_{t=t_m} \right)$$

where

$$\mathbb{1}_{t=t_m} = \begin{cases} 1 & t = t_m \\ 0 & \text{otherwise.} \end{cases}$$

References

- Andreev, K. (2002), *Evolution of the Danish Population from 1835 to 2000*, Vol. 9, University Press of Southern Denmark.
- Anisimova, M., Bielawski, J. P. and Yang, Z. (2001), ‘Accuracy and power of the likelihood ratio test in detecting adaptive molecular evolution’, *Molecular biology and evolution* **18**, 1585–1592.
- Booth, H., Hyndman, R. J., Tickle, L., De Jong, P. et al. (2006), Lee-carter mortality forecasting: a multi-country comparison of variants and extensions, Technical report, Monash University, Department of Econometrics and Business Statistics.
- Brouhns, N., Denuit, M. and Van Keilegom, I. (2005), ‘Bootstrapping the poisson log-bilinear model for mortality forecasting’, *Scandinavian Actuarial Journal* **2005**(3), 212–224.
- Brouhns, N., Denuit, M. and Vermunt, J. K. (2002), ‘A poisson log-bilinear regression approach to the construction of projected lifetables’, *Insurance: Mathematics and Economics* **31**(3), 373–393.
- Cairns, A. J. G., Blake, D. and Dowd, K. (2006), ‘A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration’, *Journal of Risk and Insurance* **73**(4), 687–718.
- Cairns, A. J. G., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D. and 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. G., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., Ong, A. and 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.
- Cox, D. R. and Hinkley, D. V. (1979), *Theoretical statistics*, CRC Press.
- DAmato, V., Haberman, S. and Russolillo, M. (2009), ‘Efficient bootstrap applied to the poisson log-bilinear lee carter model’, *Proceedings of the Applied Stochastic Models and Data Analysis (ASMDA09)* pp. 374–377.
- Davison, A. C. (1997), *Bootstrap methods and their application*, Vol. 1, Cambridge university press.

- Gelman, A., Carlin, J. B., Stern, H. S. and Rubin, D. B. (1995), *Bayesian data analysis*, Chapman and Hall: London.
- Huelsenbeck, J. P. and Crandall, K. A. (1997), ‘Phylogeny estimation and hypothesis testing using maximum likelihood’, *Annual Review of Ecology and Systematics* pp. 437–466.
- Jarner, S. F. and Kryger, E. M. (2011), ‘Modelling adult mortality in small populations: The saint model’, *Astin Bulletin* **41**(02), 377–418.
- Kendall, M., Stuart, A. and Ord, J. (1987), ‘Kendall’s advanced theory of statistics’.
- Lee, R. D. and Carter, L. R. (1992), ‘Modeling and forecasting us mortality’, *Journal of the American statistical association* **87**(419), 659–671.
- Liu, X. and Braun, W. J. (2011), ‘Investigating mortality uncertainty using the block bootstrap’, *Journal of Probability and Statistics* **2010**.
- Mood, A., Graybill, F. and Boes, D. (1963), ‘Introduction into the theory of statistics’.
- Neyman, J. and Pearson, E. S. (1992), *On the problem of the most efficient tests of statistical hypotheses*, Springer.
- Renshaw, A. E. and Haberman, S. (2003), ‘Lee–carter mortality forecasting with age-specific enhancement’, *Insurance: Mathematics and Economics* **33**(2), 255–272.
- Wilks, S. S. (1938), ‘The large-sample distribution of the likelihood ratio for testing composite hypotheses’, *The Annals of Mathematical Statistics* **9**(1), 60–62.