



**ARC Centre of Excellence in Population Ageing
Research**

Working Paper 2020/28

**Mortality Improvement Rates: Modeling, Parameter
Uncertainty and Robustness**

Andrew Hunt, Andrés M. Villegas

This paper can be downloaded without charge from the ARC Centre of
Excellence in Population Ageing Research Working Paper Series available at
www.cepar.edu.au

Mortality Improvement Rates: Modeling, Parameter Uncertainty and Robustness

Andrew Hunt^a, Andrés M. Villegas^b

^a*Pacific Life Re, London, UK*

^b*School of Risk and Actuarial Studies and
ARC Centre of Excellence in Population Ageing Research (CEPAR)
UNSW Sydney, Australia*

Abstract

Rather than looking at mortality rates directly, a number of recent academic studies have looked at modeling rates of improvement in mortality when making mortality projections. Although relatively new in the academic literature, the use of mortality improvement rates has a long-standing tradition in actuarial practice when allowing for improvements in mortality from standard mortality tables. However, mortality improvement rates are difficult to estimate robustly and models of them are subject to high levels of parameter uncertainty, since they are derived by dividing one uncertain quantity by another. Despite this, the studies of mortality improvement rates to date have not investigated parameter uncertainty due to the ad hoc methods used to fit the models to historical data. In this study, we adapt the Poisson model for the numbers of deaths at each age and year, proposed in Brouhns et al. [Insurance: Mathematics and Economics 3 (2002) 31] to model mortality improvement rates. This enables models of improvement rates to be fitted using standard maximum likelihood techniques and allows parameter uncertainty to be investigated using a standard bootstrapping approach. We illustrate the proposed modeling approach using data for the England and Wales population. The methods of these paper are available in the R package **iMoMo**.

Keywords: Mortality improvements; Mortality forecasting; Parameter uncertainty; Robustness

1. Introduction

Some of the most far-reaching social and economic challenges of the current age are caused by the rapid increases in longevity and ageing of populations across the world. One strand in meeting these challenges has been the development of a wide range of models in order to forecast the future evolution of mortality rates, based on a combination of statistical extrapolation of historic data and expert judgement.

Email addresses: andrew.hunt@pacificlifere.com (Andrew Hunt), a.villegas@unsw.edu.au (Andrés M. Villegas)

The work in this study was started before Dr Hunt commenced work at Pacific Life Re, and any opinions expressed in this paper are held in a personal capacity and should not be construed as the views of Pacific Life Re or related companies.

November 3, 2020

However, one of the subtle differences between academic models for forecasting and those used by actuaries in the life insurance industry is over what variable to model. Academic mortality models usually focus on modeling mortality rates at age, x , and time, t , denoted variously as $\mu_{x,t}$ (the instantaneous force of mortality), $m_{x,t}$ (the central rate of mortality) or $q_{x,t}$ (the one year probability of dying). Many of these models have been inspired by the seminal paper of Lee and Carter (1992) and operate in the generalized age/period/cohort framework described in Hunt and Blake (2020c) and implemented in Villegas et al. (2018). More specifically, as discussed in Hunt and Blake (2020c), much of the recent actuarial literature looking at the modeling and forecasting of human mortality builds on the Poisson log-bilinear modeling approach introduced in Brouhns et al. (2002), in which the number of deaths at age x and year t are modeled as independent Poisson variables and where the central rate of mortality, $m_{x,t}$, is taken as the response variable linked to a parametric predictor structure, $\eta_{x,t}$, by means of a log-link function, i.e.,

$$\ln m_{x,t} = \eta_{x,t}. \quad (1)$$

In contrast, practitioners are often interested primarily in the mortality improvement rates, usually defined by $-\ln\left(\frac{\mu_{x,t}}{\mu_{x,t-1}}\right)$, $-\ln\left(\frac{m_{x,t}}{m_{x,t-1}}\right)$ or $1 - \frac{q_{x,t}}{q_{x,t-1}}$. This is because it is the changes in mortality rates that are of interest when assessing longevity risk for an insurer or pension scheme. However, improvements rates are usually estimated using the largest dataset available over a long time period, often the national population, in order to give reliable estimates. Such a dataset will usually have very different mortality rates to the population of interest. Nonetheless, by considering mortality improvement rates, inferences made using these large datasets can still be used for smaller sub-populations, albeit potentially subject to longevity “basis risk” (see Haberman et al. (2014); Villegas et al. (2017)). Furthermore, the discussion of mortality improvement rates also allows practitioners to compare the evolution of mortality in populations with very different levels of mortality, for instance, men and women or in different countries. In the UK, the concept of mortality improvement rates became widely adopted among actuaries as a result of Continuous Mortality Investigation (2002) and has continued with the developments of the CMI Mortality Projection Model (Continuous Mortality Investigation (2009) and subsequent developments). Similarly, the Scale AA improvement rates were introduced by the Society of Actuaries in the United States in 1995, and the Scale BB improvement rates in 2012, for use when projecting mortality rates (Society of Actuaries Group Annuity Valuation Table Task Force, 1995; Society of Actuaries, 2012).

However, the modeling of improvement rates is more challenging than the modeling of mortality rates themselves. Since improvement rates are effectively the first derivatives of the mortality rates, any uncertainty in the measurement of mortality rates is magnified significantly in the measurement of improvement rates. On the one hand, as illustrated by Figures 1a and 1b, the general trend in generally improving mortality rates in the raw (or “crude”) data is far clearer when looking at mortality rates themselves than the improvement rates, where the noise around the signal is far more prominent. On the other hand, as Figures 1c and 1d illustrate, the age shape of mortality rates is very clear and well understood,

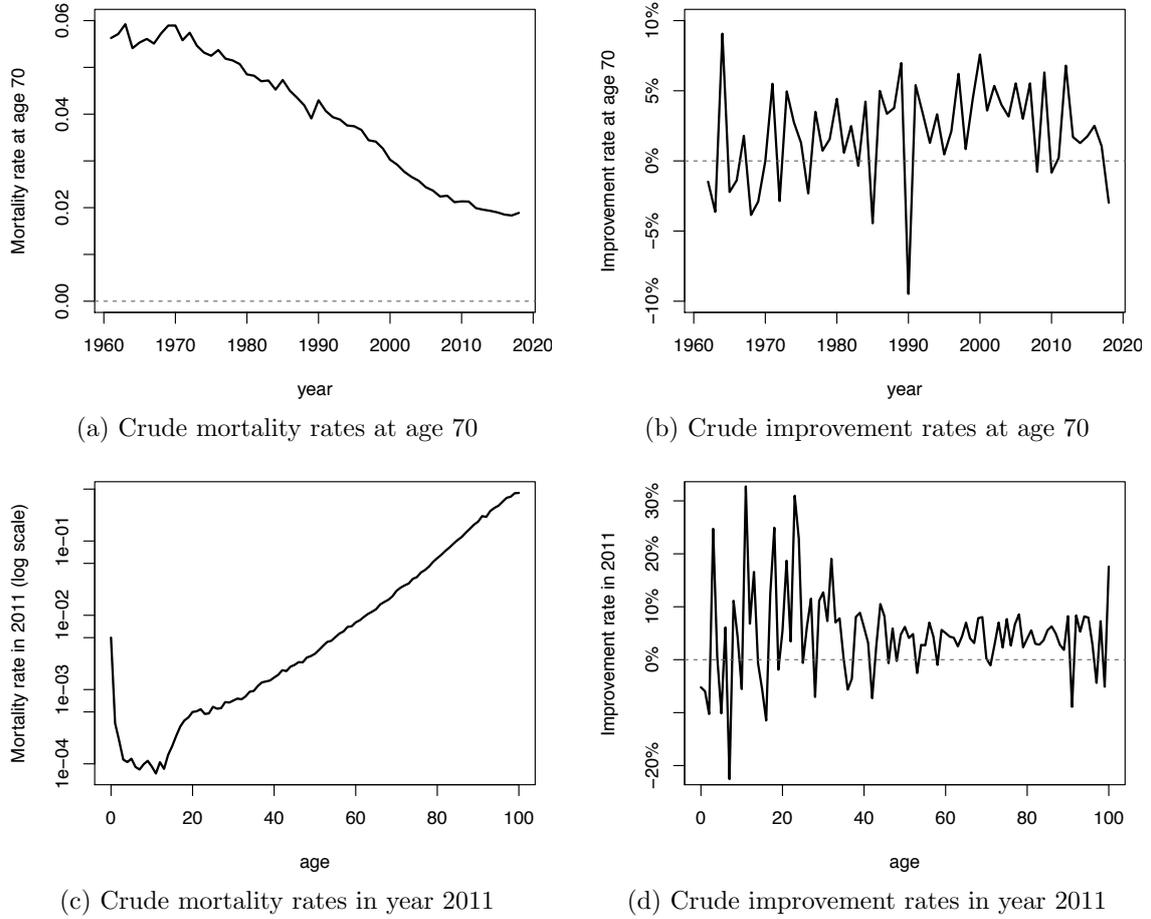


Figure 1: England and Wales male mortality and improvement rates.

while the age shape of mortality improvement rates is very noisy and displays considerable heteroscedasticity across ages.

In recent years, a number of academic studies have modified the structure in Equation (1) to look at the modeling and forecasting of mortality improvement rates. This has meant using response variables and link functions such as

$$\eta_{x,t} = \ln \left(\frac{m_{x,t+1}}{m_{x,t}} \right)$$

in Mitchell et al. (2013), and

$$\eta_{x,t} = 2 \frac{m_{x,t-1} - m_{x,t}}{m_{x,t-1} + m_{x,t}}$$

in Haberman and Renshaw (2012). This is usually thought of as using a new response variable with the log or identity link respectively, rather than keeping $m_{x,t}$ as the response variable with a non-standard link function.

Such an approach does not present any theoretical problems, however there are a number of practical issues which need to be considered. First, the distribution of the response variables is highly non-standard and so the use of the Poisson distribution is no longer appropriate. In practice, a Gaussian error structure is often assumed with suitable modifications to allow for the complex relationship between the variance of an observation and the underlying exposures.

Second, as illustrated before in Figure 1, the variance of the response variable is likely to be far higher as a proportion of the mean than when modeling mortality rates and with a high degree of heterogeneity across ages and years. The parameter error in the measurements of the free parameters in the predictor structure will therefore be far higher than for the corresponding model of mortality rates. This means we must adopt far simpler predictor structures than would be the case for models of the mortality rate. For these reasons, more research is needed before such mortality improvement models become widely adopted.

The academic studies of improvement rates to date, whilst trailblazing in their approach to the topic, have been forced to make ad hoc modeling assumptions in order to deal with the challenges associated with the direct modeling of mortality improvement rate. In contrast, a well-developed theoretical framework for the class of generalized age/period/cohort models of mortality rates has been developed. Therefore, this paper tries to apply some of the structure developed for the study of mortality rates to the modeling of mortality improvements, to reduce the need for some of the ad hoc modeling assumptions and allow a more rigorous examination of mortality improvement rates. More specifically, we adapt the Poisson model for the numbers of deaths at each age and year, proposed in Brouhns et al. (2002), to model mortality improvement rates. This approach enables models of improvement rates to be fitted using standard maximum likelihood techniques, which has several advantages:

- i. the Poisson structure for death counts accounts automatically for heterogeneity across ages due to exposures (c.f., Haberman and Renshaw (2012)), and
- ii. it allows parameter uncertainty to be investigated using the standard bootstrapping techniques considered in Brouhns et al. (2005) and Koissi et al. (2006).

The remainder of this paper is organised as follows. In Section 2 we introduce some of the notation used throughout the paper. In Section 3 we investigate the connections between models of mortality and improvement rates and develop techniques for fitting improvement rate models to data using a Poisson framework. In Section 4 we apply these techniques to the mortality experience of England and Wales. In doing so, we note some of the differences in the definition of improvement rates in previous studies, and the impact these have on the robust estimation of the parameters within improvement rate models. We also investigate the impact of parameter uncertainty on the robustness and stability of the mortality projection derived from mortality improvement rate models. Finally, in Section 5 we summarize our findings and provide some conclusions.

2. Data and notation

Throughout this paper we assume that the available data comprise a cross classified mortality experience containing observed numbers of deaths at age x in year t , $d_{x,t}$, with matching

central exposures, $e_{x,t}$. We assume that age, x , is in the range $[1, X]$, calendar year or period, t , is in the range $[0, T]$ and, therefore, that year of birth, $y = t - x$, is in the range $[-X, T - 1]$.

We denote the force of mortality and the central rate of mortality by $\mu_{x,t}$ and $m_{x,t}$, respectively, with the crude (empirical) estimate of the latter being $\hat{m}_{x,t} = d_{x,t}/e_{x,t}$. Furthermore, we assume that the force of mortality is constant over each year of age x and calendar year t , implying that the force of mortality and central death rate coincide, i.e., $\mu_{x,t} = m_{x,t}$. Finally, consistent with Brouhns et al. (2002), we assume that the random number of deaths, $D_{x,t}$, at age x in year t is a Poisson distributed random variable with distribution

$$D_{x,t} \sim \text{Poisson}(e_{x,t}m_{x,t}) \quad (2)$$

and, hence, that $m_{x,t} = \mathbb{E}(D_{x,t})/e_{x,t}$. We note that the observed death counts, $d_{x,t}$, are the realization of the random variable defined in Equation (2).

3. Poisson improvement rate models

In this section we exploit the connections between improvement rate models and mortality rate models to produce a Poisson formulation of mortality improvement rate models. We then discuss how this formulation can be used to assess parameter uncertainty in mortality improvement rate models and to obtain forecasts of mortality rates.

3.1. Preliminaries

Similar to Mitchell et al. (2013), we start from a model of the annual improvement rate, given by

$$-\ln\left(\frac{m_{x,t}}{m_{x,t-1}}\right) = -\Delta \ln m_{x,t} = \eta_{x,t}, \quad (3)$$

where the minus sign is for presentational purposes to ensure that improvements (i.e., falling) in mortality rates are positive and that $\eta_{x,t}$ can be interpreted as the continuous rate of improvement at age x in year t .

In order to add structure to this, we then define the predictor structure, $\eta_{x,t}$, using the general age/period/cohort structure described in Hunt and Blake (2020c), i.e.,

$$\eta_{x,t} = \alpha_x + \sum_{i=1}^N \beta_x^{(i)} \kappa_t^{(i)} + \gamma_{t-x}, \quad (4)$$

where

- α_x is a static function of age, which gives the average (constant) rate of improvement in mortality at each age x ;
- $\kappa_t^{(i)}$ are period functions governing the change in improvement rate in year t ;

- $\beta_x^{(i)}$ are age functions which modulate the corresponding period functions;² and
- γ_c is a cohort function describing systematic differences in the rate of improvement which depend upon a cohort's year of birth, $c = t - x$.

Unlike Mitchell et al. (2013) and Haberman and Renshaw (2012), we do not model $-\Delta \ln m_{x,t}$ directly, since the mortality improvement rates in this specification do not follow a standard probability distribution. They are also highly heteroscedastic, meaning that standard estimation techniques are problematic. Instead, we iterate Equation (3) to give

$$\ln(m_{x,t}) = \ln(m_{x,0}) - \sum_{\tau=1}^t \eta_{x,\tau}.$$

By defining $A_x = \ln(m_{x,0})$ as the initial mortality curve, this can be re-written as

$$\ln(m_{x,t}) = \tilde{\eta}_{x,t} = A_x - \sum_{\tau=1}^t \eta_{x,\tau}. \quad (5)$$

In this form, it is natural to use a Poisson model for the death counts, such that the number of deaths observed at age x and for year t follows a Poisson distribution with mean $e_{x,t}m_{x,t}$. Under this assumption and with the log-link function

$$D_{x,t} \sim \text{Poisson}(e_{x,t} \exp(\tilde{\eta}_{x,t})), \quad (6)$$

as per Brouhns et al. (2002) and Hunt and Blake (2020c), but with the modified predictor structure, $\tilde{\eta}_{x,t}$, which gives us a model of mortality improvement rates directly rather than a model for mortality rates.³

We also see that, since we can use the Poisson model for the death counts in this formulation of an improvement rate model, we are able to estimate the parameters using maximum likelihood techniques and estimate their parameter uncertainty using the techniques of Brouhns et al. (2005) and Koissi et al. (2006). This, therefore, overcomes some of the key limitations of the methods in Mitchell et al. (2013) and Haberman and Renshaw (2012, 2013), which used more ad hoc fitting techniques and did not investigate parameter uncertainty.⁴

²These age functions can be non-parametric (having form determined entirely by the data) or parametric (having a pre-defined functional form), as discussed in Hunt and Blake (2020c).

³One drawback of using a Poisson model for the death counts, common to both models of mortality rates and improvement rates, is that it assumes that the variance of an observation is equal to its expectation. Such over-dispersion can be dealt with by using an over-dispersed Poisson model in a generalized non-linear modeling framework or by allowing for heterogeneity in the population via the use of the negative binomial distribution, such as in Delwarde et al. (2007); Dodd et al. (2020); Li et al. (2009). However, we do not investigate this further in this study.

⁴In the case of Mitchell et al. (2013), least squares estimation was used to fit the improvement rates, while in Haberman and Renshaw (2012), an iterated GLM procedure was used to allow for overdispersion in the observed improvement rates. However, in neither case, were these distributions selected on the basis of providing an appropriate distribution for the observed death counts. Consequently, this means that many common methods for assessing parameter uncertainty are not appropriate, as discussed in Section 3.5.

3.2. Equivalent mortality rate structure and indirect estimation of improvement rate models

We now exploit the connection between improvement rate models and traditional mortality rate models to devise an estimation approach for the Poisson improvement rate model defined by Equations (5) and (6).

From Equations (4) and (5), the predictor structure in this latter Equation can be re-written as

$$\begin{aligned}
 \ln(m_{x,t}) &= \tilde{\eta}_{x,t} \\
 \ln(m_{x,t}) &= A_x - \sum_{\tau=1}^t \eta_{x,\tau} \\
 \ln(m_{x,t}) &= A_x - \sum_{\tau=1}^t \left(\alpha_x + \sum_{i=1}^N \beta_x^{(i)} \kappa_t^{(i)} + \gamma_{\tau-x} \right) \\
 \ln(m_{x,t}) &= A_x - \alpha_x t - \sum_{i=1}^N \beta_x^{(i)} \sum_{\tau=1}^t \kappa_t^{(i)} - \sum_{\tau=1}^t \gamma_{\tau-x} \\
 \ln(m_{x,t}) &= A_x - \alpha_x t + \sum_{i=1}^N \beta_x^{(i)} K_t^{(i)} + \Gamma_{t-x} \tag{7}
 \end{aligned}$$

with

$$K_0^{(i)} = 0 \quad \text{and} \quad \Gamma_{-X} = 0, \tag{8}$$

and

$$K_t^{(i)} = - \sum_{\tau=1}^t \kappa_t^{(i)} \quad \text{and} \quad \Gamma_{t-x} = - \sum_{\tau=1}^t \gamma_{\tau-x}, \quad \text{for } 1 \leq t \leq T. \tag{9}$$

In Equation (7), it is clear that α_x is determining the constant trend rate of mortality improvement in the historic data at each age.

We also see that, if the α_x term is not included, Equation (7) is equivalent to a standard age/period/cohort model (see Hunt and Blake (2020c)). Therefore, we see that conventional mortality rates models are identical to improvement rates models without constant improvement terms, and merely differ in the presentation of the parameters (i.e., the constraints in Equation (8) as opposed to the conventional identifiability constraints $\sum_t K_t = 0$ and $\sum_c \Gamma_c = 0$).

In contrast, we see that including an α_x constant improvement term in Equation (7) extends the family of generalized age/period/cohort models discussed in Hunt and Blake (2020c) with a term that is non-parametric in age and linear in time. Therefore, every mortality rate model discussed in Hunt and Blake (2020c) has an extended version which includes a constant improvement rate term, which is equivalent to using the same predictor structure for mortality improvement rates rather than mortality rates.

This means that we can then estimate the improvement rate model in Equations (5) and (6) *indirectly* by first estimating the equivalent mortality rate model defined by Equation (7) with the constraints in Equation (8) and recover the parameters of the improvement rate model using the relationships in Equation (9). In this paper we follow such an estimation approach and develop an extension the R package **StMoMo** (Villegas et al., 2018), which enables the fitting of general age/period/cohort mortality rate models. We make this extension available through the R package **iMoMo** which is available in Github at <https://github.com/amvillegas/iMoMo>. In Appendix B we illustrate the use of **iMoMo** to implement the methods discussed in this paper.

3.3. Crude mortality rates and direct estimation of improvement rate models

One of the main differences between the formulation of mortality improvement rate models in Equation (3) and that in Mitchell et al. (2013) and Haberman and Renshaw (2012) is that the previous literature estimates the models directly on the improvement rates. This is done by first computing the improvement rates based on the *crude* mortality rates, $\hat{m}_{x,t} = d_{x,t}/e_{x,t}$, so that

$$-\ln\left(\frac{\hat{m}_{x,t}}{\hat{m}_{x,t-1}}\right) = -\Delta \ln \hat{m}_{x,t} = \eta_{x,t} \quad (10)$$

in Mitchell et al. (2013). This contrasts with our formulation in Equation (3) which uses the *fitted* mortality rates, $m_{x,t} = \mathbb{E}(D_{x,t})/e_{x,t}$, to compute the mortality improvement rates.

We can convert Equation (10) into a Poisson formulation of the number of deaths to get⁵

$$D_{x,t} \sim \text{Poisson}(e_{x,t}\hat{m}_{x,t-1} \exp(-\eta_{x,t})) \quad (11)$$

as opposed to

$$D_{x,t} \sim \text{Poisson}(e_{x,t}m_{x,t-1} \exp(-\eta_{x,t})) \quad (12)$$

from Equation (6) in our indirect formulation. From Equations (11) and (12) we also see that

$$\mathbb{E}(D_{x,t}) = \frac{e_{x,t}}{e_{x,t-1}}d_{x,t-1} \exp(-\eta_{x,t}) \quad (13)$$

in the direct formulation of mortality improvement rates whilst

$$\mathbb{E}(D_{x,t}) = \frac{e_{x,t}}{e_{x,t-1}}\mathbb{E}(D_{x,t-1}) \exp(-\eta_{x,t}) \quad (14)$$

⁵Equation (11) follows from noting that under the Mitchell et al. (2013) form of an improvement rate model the expected number of deaths at age x in year t is $e_{x,t}\hat{m}_{x,t-1} \exp(-\eta_{x,t})$

under the indirect formulation. Hence, we see that the direct formulation builds the observed idiosyncratic risk of the observed deaths, $d_{x,t-1}$, into the modeling framework, which is avoided in the indirect framework.

One key advantage of using Equation (11) to frame the direct modeling of mortality improvement rates is that we can also rely on standard Poisson generalized (non-)linear modeling techniques to estimate the Mitchell et al. (2013) form of an improvement rate model. Specifically, this is achieved by setting $\ln \hat{m}_{x,t-1}$ as an offset within the generalized (non-)linear model predictor structure so that

$$\mathbb{E}(D_{x,t}) = \ln e_{x,t} + \ln \hat{m}_{x,t-1} - \eta_{x,t}. \quad (15)$$

3.4. Direct vs. indirect estimation of mortality improvement rate models

The differences between formulating a mortality improvement rate model directly as in Equation (11) or indirectly as in Equation (12), although subtle, have profound consequences on the parameter uncertainty and robustness as we will discuss in the remainder of this paper. Figure 2 schematizes the conceptual differences between the direct and indirect estimation approaches. From now onwards, when referring to the estimation of the parameters in the predictor $\eta_{x,t}$, we will say that we use a “direct” estimation approach whenever we follow the left-red route in Figure 2 and assume Equation (11) and that we use an indirect estimation approach whenever we follow the right-gray route in Figure 2 and assume Equation (12).

We note that while the pioneering studies of Haberman and Renshaw (2012) and Mitchell et al. (2013) followed the direct modeling route, recently, other studies such as Dodd et al. (2020), Li et al. (2020) and Richards et al. (2019) have explored the use of the indirect modeling route when estimating specific mortality improvement rate structures.

3.5. Parameter uncertainty

As discussed in Section 1 one of the key problems with investigating mortality improvement rates is the level of uncertainty in estimating models for them. This is far greater than in similar models for mortality rates, and is a feature which is understated in models of improvement rates to date.

To give an example of this, consider the situation where we are trying to estimate mortality rates, when the true mortality rate is $m_{x,t} = 0.5\%$ p.a.. Using a Poisson model, the relative parameter uncertainty in our estimate is proportional to $1/\sqrt{e_{x,t}m_{x,t}} = 1/\sqrt{\mathbb{E}(D_{x,t})}$, i.e., inversely proportional to the square root of the expected number of deaths. So to obtain a relative uncertainty of 1% in our estimate of the mortality rate (i.e., a one standard deviation confidence interval for our mortality rate of (0.495%, 0.505%)) requires roughly 10,000 expected deaths or an observed population of 2 million lives.

If the true rate of mortality improvement is 2% over a one year period, then observing the same population in the following year will yield an estimate for the mortality rate in the second year of (0.485%, 0.495%). Therefore, although our central estimate for the annual

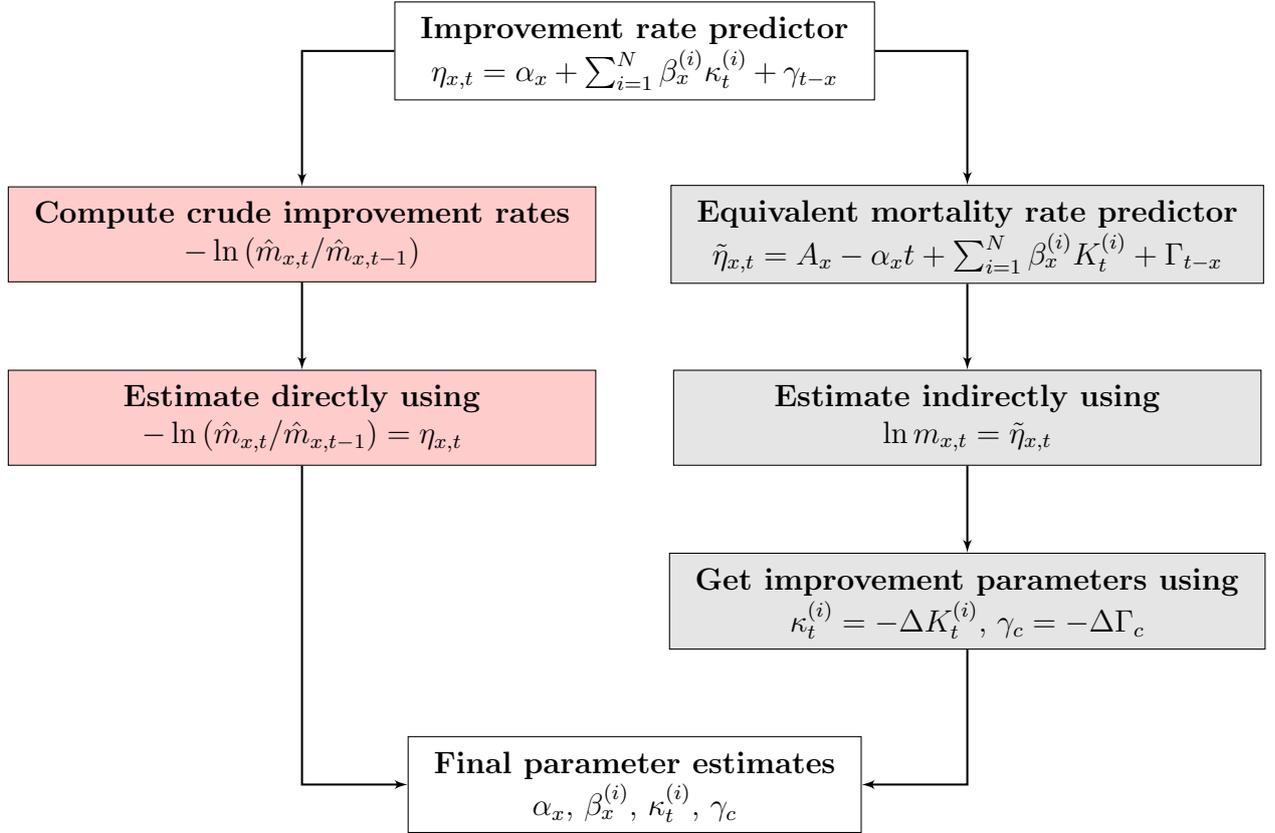


Figure 2: Direct and indirect estimation approaches for mortality improvement rate models.

improvement rate observed will be $1 - \frac{0.49\%}{0.50\%} = 2\%$, the range of our confidence interval for the annual improvement will be $(0.0\%, 4.0\%)$, i.e., a relative uncertainty in the estimate of the rate of improvement of 100%. In order to get comparable levels of certainty in our estimates of improvement rates to those obtained for mortality rates themselves, we roughly need to square the number of expected deaths being observed each year (e.g., 1 million expected deaths in order to obtain a relative uncertainty of 1%), with a corresponding increase in the number of lives under observation (e.g., 200 million lives). This is clearly impractical in almost all circumstances.

This is not a fatal limitation when using improvement rate models as long as we accept the fundamental uncertainty in our parameter estimates: however, this means that it is vital that we allow for parameter uncertainty when using improvement rate models. Because there was no clear process generating the observed numbers of deaths or improvement rates in the models of Mitchell et al. (2013) and Haberman and Renshaw (2012), this was very difficult to do systematically. However, since we assume a Poisson distribution for the death counts, we can use standard techniques for estimating parameter uncertainty in our framework. Specifically, we use the semi-parametric bootstrapping technique of Brouhns et al. (2005), which generates new death counts by sampling from the Poisson distribution with mean $d_{x,t}$, to which the model is refitted in order to give new parameter estimates. Alternatively, one

could use the residual bootstrapping technique of Koissi et al. (2006) which re-samples the deviance residuals from fitting the model to generate new death counts. In practice, however, both approaches yield qualitatively similar results. We refer the interested reader to Villegas et al. (2018, Section 8) for the specific details of our implementation of the bootstrapping approaches of Brouhns et al. (2005) and Koissi et al. (2006).

3.6. Projection of mortality and improvement rates

To project the improvement rate model to give future improvement rates (and hence future mortality rates), we project the period and cohort functions in a similar fashion to a model of mortality rates. Therefore, similar time series techniques⁶ can be used. However, since the model is now one of improvement rates rather than mortality rates, the demographic significance⁷ of the parameters is now different, which will influence our choice of projection model.

In general, we can assume that the d difference of the period index $\boldsymbol{\kappa}_t := (\kappa_t^{(1)}, \dots, \kappa_t^{(N)})'$ follows a vector autoregressive (VAR) model around a linear trend (Pfaff, 2008):

$$\Delta^d \boldsymbol{\kappa}_t = \mathbf{C} + \mathbf{D}t + \sum_{i=1}^p \mathbf{A}_i \Delta^d \boldsymbol{\kappa}_{t-1} + \boldsymbol{\xi}_t^\kappa, \quad \boldsymbol{\xi}_t^\kappa \sim N(\mathbf{0}, \Sigma), \quad (16)$$

where \mathbf{C} and \mathbf{D} are N -dimensional vectors of parameters, $\mathbf{A}_1, \dots, \mathbf{A}_p$ are $N \times N$ matrices of autoregressive parameters, and Σ is the $N \times N$ variance-covariance matrix of the multivariate white noise $\boldsymbol{\xi}_t^\kappa$. We note that the VAR(1) model used in Haberman and Renshaw (2012) and the multivariate random walk with drift are particular cases of Equation (16).

As for the cohort effects, we can assume in general that they follow a ARIMA(p, q, d) with drift, i.e.,

$$\Delta^d \gamma_c = \delta_0 + \phi_1 \Delta^d \gamma_{c-1} + \dots + \phi_p \Delta^d \gamma_{c-p} + \epsilon_c + \delta_1 \epsilon_{c-1} + \dots + \delta_q \epsilon_{c-q}, \quad (17)$$

where δ_0 is the drift parameter, ϕ_1, \dots, ϕ_p are the autoregressive coefficients with $\phi_p \neq 0$, $\delta_1, \dots, \delta_q$ are the moving average coefficients with $\delta_q \neq 0$ and ϵ_c is a Gaussian white noise process with variance σ_ϵ .

The time series models in (16) and (17) can be used to obtain projected values of the period index $\boldsymbol{\kappa}_{T+s} := (\kappa_{T+s}^{(1)}, \dots, \kappa_{T+s}^{(N)})'$ and cohort index γ_{T-1+s} , $s = 1, \dots, h$, respectively, and to derive projected values of mortality improvement rates:

$$\eta_{x,T+s} = \alpha_x + \sum_{i=1}^N \beta_x^{(i)} \kappa_{T+s}^{(i)} + \beta_x^{(0)} \gamma_{T-x+s}.$$

⁶Allowing for similar issues as described in Hunt and Blake (2020a,b) in order to obtain “well-identified” projections which do not depend on the arbitrary identifiability constraints chosen when fitting the model.

⁷Demographic significance is defined in Hunt and Blake (2020c) as the interpretation of the components of a model in terms of the underlying biological, medical or socio-economic causes of changes in mortality rates which generate them.

Table 1: Model structures considered in this paper.

Model	Equivalent Mortality Model ($\tilde{\eta}_{x,t}$)	Improvement Model ($\eta_{x,t}$)
CI	$A_x - \alpha_x t$	α_x
LC	$A_x + \beta_x^{(1)} K_t^{(1)}$	$\beta_x^{(1)} \kappa_t^{(1)}$
LC-CI	$A_x - \alpha_x t + \beta_x^{(1)} K_t^{(1)}$	$\alpha_x + \beta_x^{(1)} \kappa_t^{(1)}$
CBD	$A_x + K_t^{(1)} + (x - \bar{x}) K_t^{(2)}$	$\kappa_t^{(1)} + (x - \bar{x}) \kappa_t^{(2)}$
CBD-CI	$A_x - \alpha_x t + K_t^{(1)} + (x - \bar{x}) K_t^{(2)}$	$\alpha_x + \kappa_t^{(1)} + (x - \bar{x}) \kappa_t^{(2)}$
APC-CI	$A_x - \alpha_x t + K_t^{(1)} + \Gamma_{t-x}$	$\alpha_x + \kappa_t^{(1)} + \gamma_{t-x}$

Now, to obtain projected mortality rates we use

$$m_{x,T+s} = \hat{m}_{x,T} \exp \left(- \sum_{\tau=1}^s \eta_{x,T+\tau} \right),$$

where $\hat{m}_{x,T}$ are the last observed central mortality rates.

4. England and Wales 1961-2018, 20-89: A Poisson improvement rate approach

In this section we illustrate the discussion of Section 3 by applying a Poisson improvement rate approach to the modeling of mortality in England and Wales. In particular and similarly to Haberman and Renshaw (2012), we use historical mortality data for the England and Wales male population covering calendar years 1961-2018 and ages 20-89 obtained from the Human Mortality Database (2020). Of particular interest in this discussion are the inclusion of constant improvement rates, the implications of using a direct or indirect estimation approach and the impact of parameter uncertainty on the robustness of the projections derived from improvement rate models..

4.1. Predictor structures

We focus on the models summarized in Table 1. For each model, this table includes the predictor structure used to model improvement rates (recall Equation (4)) as well as the equivalent mortality rate predictor as per Equation (7).

Model CI represents a simple model including only constant improvement rates, whose equivalent mortality rate model is similar to the generalized linear model of mortality rates discussed in Renshaw and Haberman (2003). Model LC is the celebrated Lee and Carter (1992) model, in both mortality rate and improvement rate form, while the LC-CI structure corresponds to Lee-Carter model with added constant mortality improvement rates. This latter predictor structure was considered in Mitchell et al. (2013) in its improvement rate form and in Callot et al. (2016) in its mortality rate form. Model CBD is the improvement rate

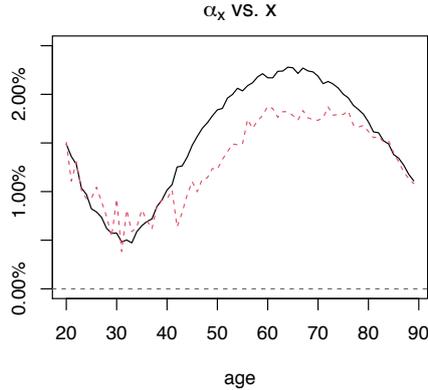


Figure 3: Parameters for model CI. England and Wales males, age 20-89, period 1961-2018.

equivalent of the two-factor model introduced in Cairns et al. (2006) and in mortality rate form is equivalent to the “CBDX” model discussed in Hunt and Blake (2020a) and Dowd et al. (2020). Similar to the LC-CI model, the CBD-CI stands for the CBD model including constant mortality improvements. The APC-CI structure is the improvement rate version of the classical Age-Period-Cohort model. Such model has been considered by the (Continuous Mortality Investigation, 2016a,b) in the latest versions of the widely used CMI mortality projection model. Such structure has also recently been investigated in Richards et al. (2019) and Dodd et al. (2020).

In estimating the parameters of the models in Table 1, we impose where necessary the standard parameter constraints. However, for the LC and LC-CI, we deviate from the standard $\sum_x \beta_x^{(1)} = 1$ constraint and impose instead $\sum_x \beta_x^{(1)} = X$, so that the period index $\kappa_t^{(1)}$ can roughly be interpreted as average improvement rates in year t (or average deviations from the constant improvement rates in the LC-CI structure). The specific parameter constraints applied in estimating the models in Table 1 are discussed in Appendix A.

The parameter estimates of all the models in Table 1 applied to male data for England and Wales over the period 1961-2018 and ages 20-89 are shown in Figures 3-8. In these figures, black continuous lines depict parameter estimates obtained with the indirect estimation approach introduced in this paper while red-dashed lines depict parameter estimates obtained with the direct estimation approach discussed in Section 3.3. From Figures 3-8, we note the following:

- The noticeable differences in the estimated age-dependent parameters α_x and $\beta_x^{(1)}$ under the direct and indirect estimation approaches, with the direct approach producing estimates that are considerably less smooth across ages. This reflects the fact that the direct approach builds the idiosyncratic sampling risk from the observed deaths into the parameter estimates in a fashion that is avoided by the indirect approach. Biological reasonableness would suggest that improvement rates should be continuous across ages, both to avoid discontinuities in projected mortality rates in future and because the underlying drivers of aging are unlikely to change rapidly with age. This

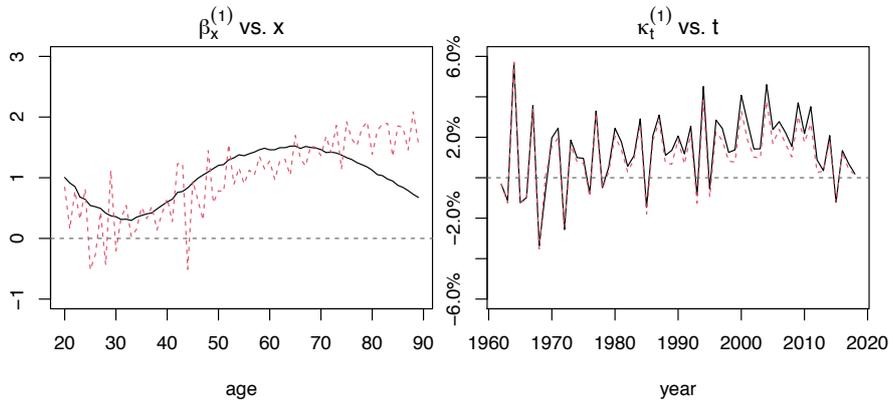


Figure 4: Parameters for model LC. England and Wales males, age 20-89, period 1961-2018.

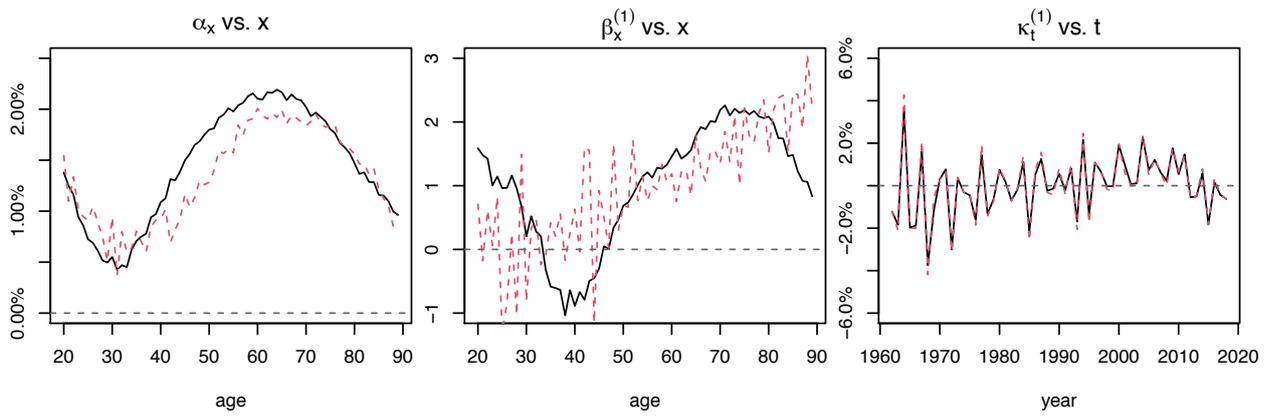


Figure 5: Parameters for model LC-CI. England and Wales males, age 20-89, period 1961-2018.

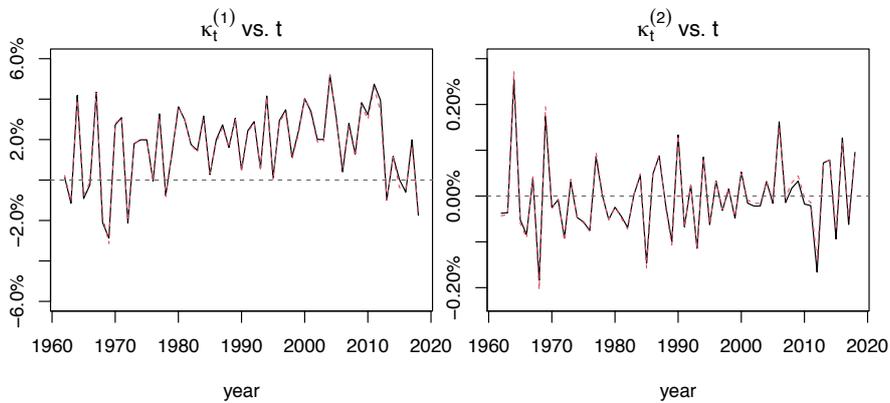


Figure 6: Parameters for model CBD. England and Wales males, age 20-89, period 1961-2018.

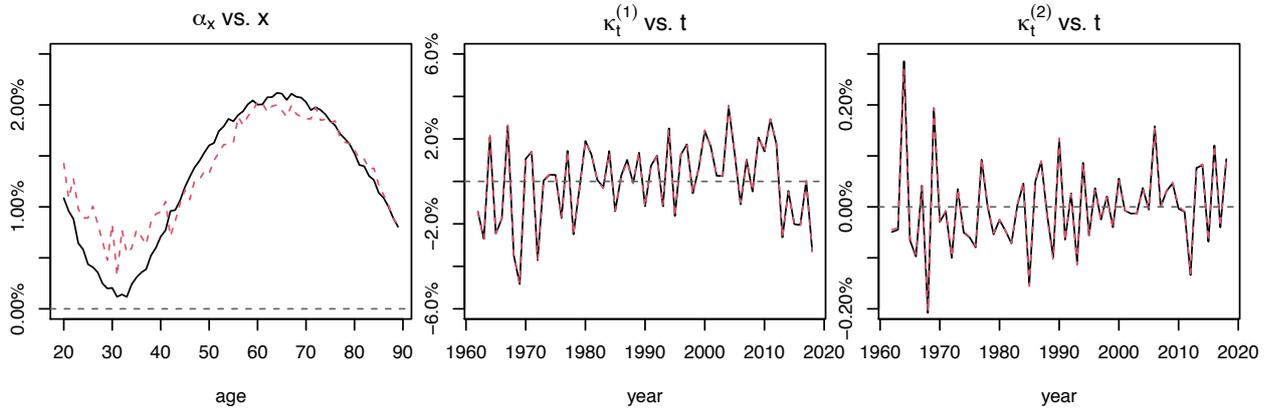


Figure 7: Parameters for model CBD-CI. England and Wales males, age 20-89, period 1961-2018.

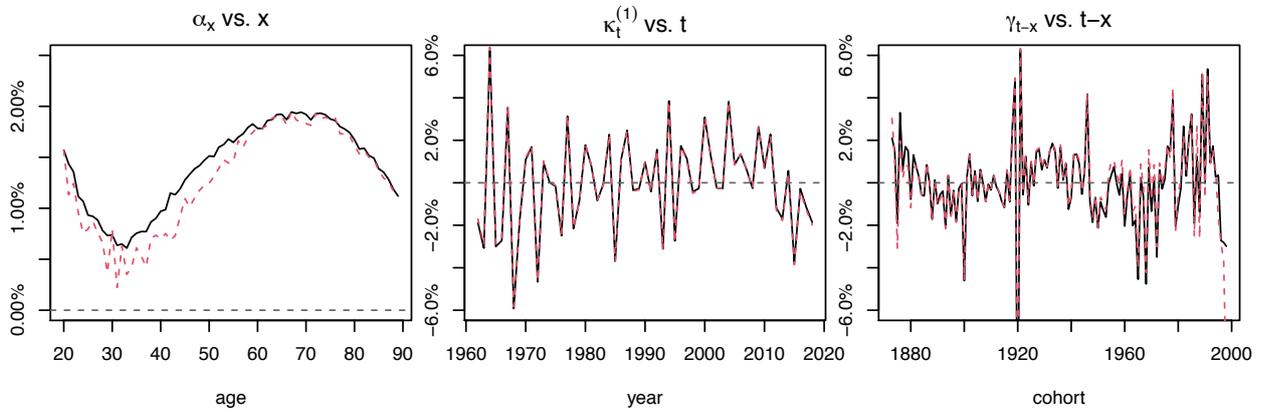


Figure 8: Parameters for model APC-CI. England and Wales males, age 20-89, period 1961-2018.

cautions us against using the direct approach to make projections unless the age-dependent parameter estimates are smoothed.

- The contrasting close agreement between the estimates of the period improvements $\kappa_t^{(1)}$, $\kappa_t^{(2)}$ and cohort improvements γ_{t-x} under the alternative estimation approaches. This suggests that the idiosyncratic sampling error being built into the estimates of the age-dependent parameters using the direct approach is less important for the estimates of period and cohort improvement rates. It also suggests that there is less to choose between estimation techniques for models without or with fewer age-dependent parameters (for instance, the CBD model).
- The clear interpretation of the α_x term as average mortality improvements, indicating that in England and Wales mortality improvement rates over the 1961-2018 period have ranged from about 0.75% p.a. between ages 20-30 and about 2% p.a. at ages 60-70.
- The similarity in the α_x parameters under the indirect approach for all models with such a term. Since the interpretation above gives these parameters a clear demographic

significance, this should not be surprising since it relates these parameters directly to quantities that could be estimated from the model in a model-independent fashion. This is comparable to the static age function in traditional mortality models. However, this demographic significance is lost in the direct estimation method.

- The ease of interpretation of the primary period improvement $\kappa_t^{(1)}$ whose numerical value can be thought of as the average improvement rate observed in the year t for models without constant improvement rate and as the average deviation from the constant improvement rates for predictors with an α_x term. For example, from the parameter value of $\kappa_{2010}^{(1)}$ in the CBD model in Figure 6 we can say that in 2010 mortality improvements were 3.2% on average across all ages. Similarly, the parameter value of $\kappa_{2010}^{(1)}$ the CBD-CI model in Figure 7 indicate that in 2010 mortality improvements were 1.4% higher than the average improvement rate observed over the 1961-2018 period. This clear interpretation of the primary period index is in contrast with traditional mortality rate models where it is difficult to link the value of the period effects to quantities with intuitive practical relevance.
- The interpretation of the cohort improvements, γ_{t-x} , in Figure 8, as average deviations in improvement rates. In particular, we see that the so-called golden generation born in the inter-war period (see Willets (2004) and Murphy (2009)) has experienced mortality improvements of around 1%-2% p.a. higher than the average. It is also interesting to note the existence of a “tarnished” cohort born after the Second World War who, in contrast, appear to be experiencing worse than average mortality improvements.

4.2. Impact of parameter uncertainty

We now turn our attention to the investigation of the impact of parameter uncertainty on the estimation of the parameters of improvement rate models. To do so, for each of the six predictor structures in Table 1 and for the two parameter estimation approaches, we have generated 1,000 bootstrapped samples of the model parameters using the semi-parametric bootstrapping approach introduced in Brouhns et al. (2005). Figure 9 presents fan charts depicting the 50%, 80% and 95% bootstrapped confidence intervals of the parameter of the CI, LC and CBD models. The results for these three models are representative of the results for all the six models considered in this paper.

From Figure 9a we see that using a direct parameter estimation approach akin to the one used in Mitchell et al. (2013) and Haberman and Renshaw (2012) results in significantly higher uncertainty in estimates of constant improvement rates parameters, α_x , than using the indirect estimation approach introduced in this paper. For instance, under the direct approach the 95% confidence interval of the improvement rate at age 40, α_{40} , is (0.71%, 1.13%), which is 3.3 times wider than the equivalent (0.95%, 1.08%) under the indirect approach. This is the result of building the observed idiosyncratic sampling errors into the parameter estimates in the direct approach.

Similarly, Figure 9b shows that non-parametric age-modulating parameters, $\beta_x^{(1)}$, also suffer from considerable parameter uncertainty under the direct estimation approach. It is also

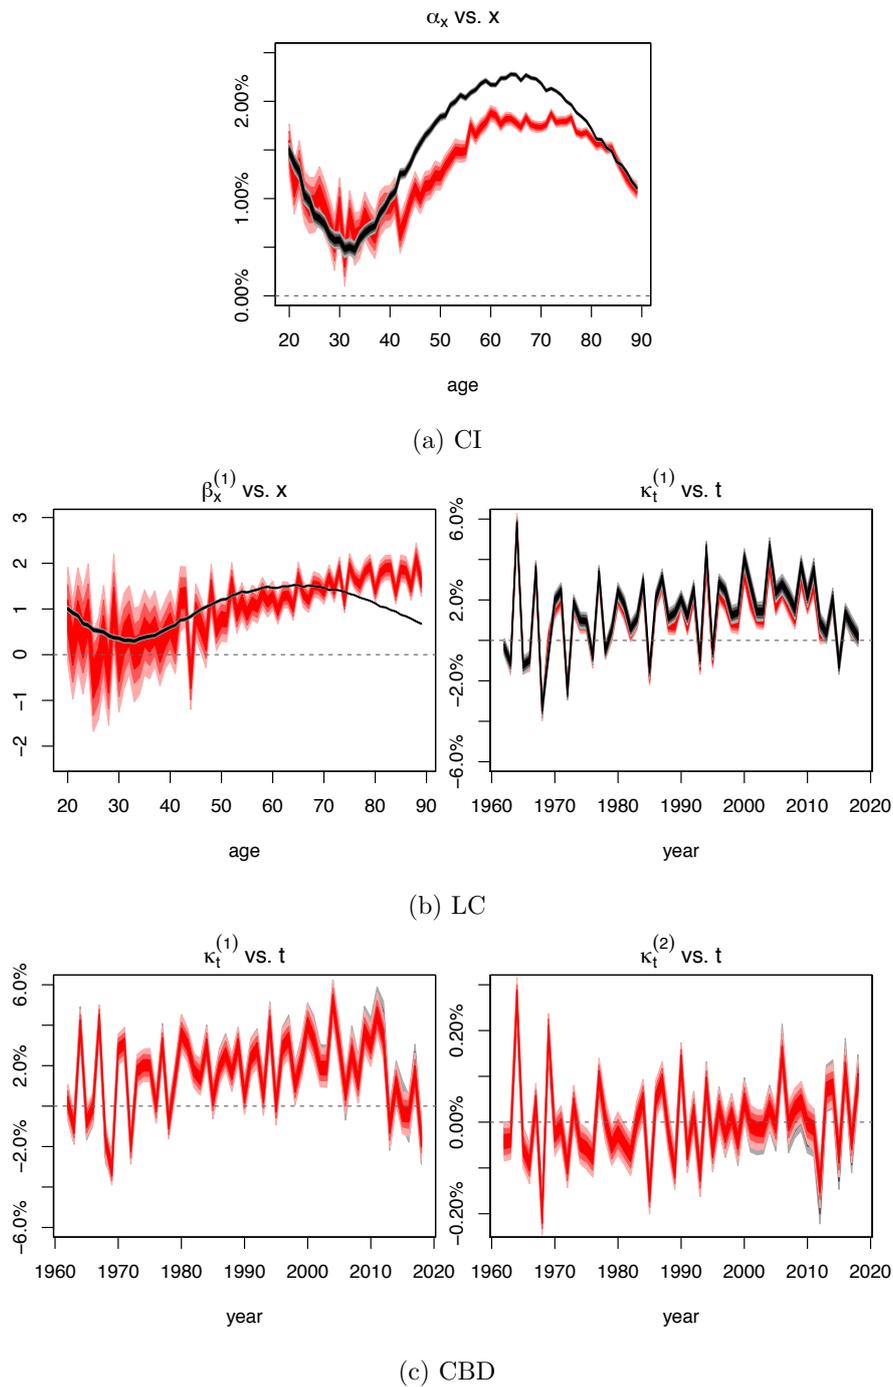


Figure 9: Parameters for models CI, LC and CBD with parameter uncertainty. England and Wales males, age 20-89, period 1961-2018. Shades in the fan represent confidence intervals at the 50%, 80% and 95% level. Black fans correspond to the indirect estimation approach and red fans to the direct estimation approach.

interesting to note that, in many cases, the confidence intervals of the age parameters under the direct approach do not contain the indirect parameter estimates. Therefore, it is not simply a case that the direct approach is estimating the same parameter values but with less

precision. In contrast, Figure 9c indicates that period indexes, $\kappa_t^{(i)}$, are in general robust with negligible differences in levels of uncertainty between the two estimation approaches.

To understand the differences in uncertainty levels produced by the two estimation approaches, it is instructive to consider in more detail the CI model. Specifically, note that under the indirect estimation approach the constant improvement rate at age x , α_x , is estimated as the slope of the (Poisson) linear regression

$$\ln \hat{m}_{x,t} = A_x - \alpha_x t + \epsilon_{x,t},$$

which depends on the whole historical mortality profile, $\{\hat{m}_{x,0}, \hat{m}_{x,1}, \dots, \hat{m}_{x,T}\}$. By contrast, under the direct estimation approach the estimate of α_x is approximately the average of the observed improvement rates at age x over the investigation period. That is,

$$\begin{aligned} \alpha_x &\approx -\frac{1}{T} \sum_{t=1}^T \Delta \ln \hat{m}_{x,t} \\ \alpha_x &\approx \frac{\ln \hat{m}_{x,0} - \ln \hat{m}_{x,T}}{T}, \end{aligned}$$

which depends only on the observed mortality rates at the start and at the end of the investigation period. Clearly, using only the first and last observations, as opposed to all the historical observations, will result in more uncertain estimates which are less robust to the addition of new data or to parameter uncertainty.

4.3. Mortality rate projections

The differences in central estimates and levels of uncertainty of the model parameters produced by the two estimation approaches can have an important impact on the mortality projections produced by the models. To investigate this potential issue, in Figure 10 we present for models CI, LC, CBD and APC-CI fan charts of mortality rate forecasts at selected ages. For each model we consider the following four types of forecasts:

- i. Forecast produced by the indirect approach without allowance for parameter uncertainty;
- ii. Forecast produced by the indirect approach with allowance for parameter uncertainty;
- iii. Forecast produced by the direct approach without allowance for parameter uncertainty;
and
- iv. Forecast produced by the direct approach with allowance for parameter uncertainty.

From Figure 10 we note the following:

- The noticeable impact of considering parameter uncertainty under the direct estimation approach for models CI at all ages and for the LC at younger ages. In particular, we note that due to the absence of a period index in the CI structure, this model only provides point forecasts when parameter uncertainty is ignored.

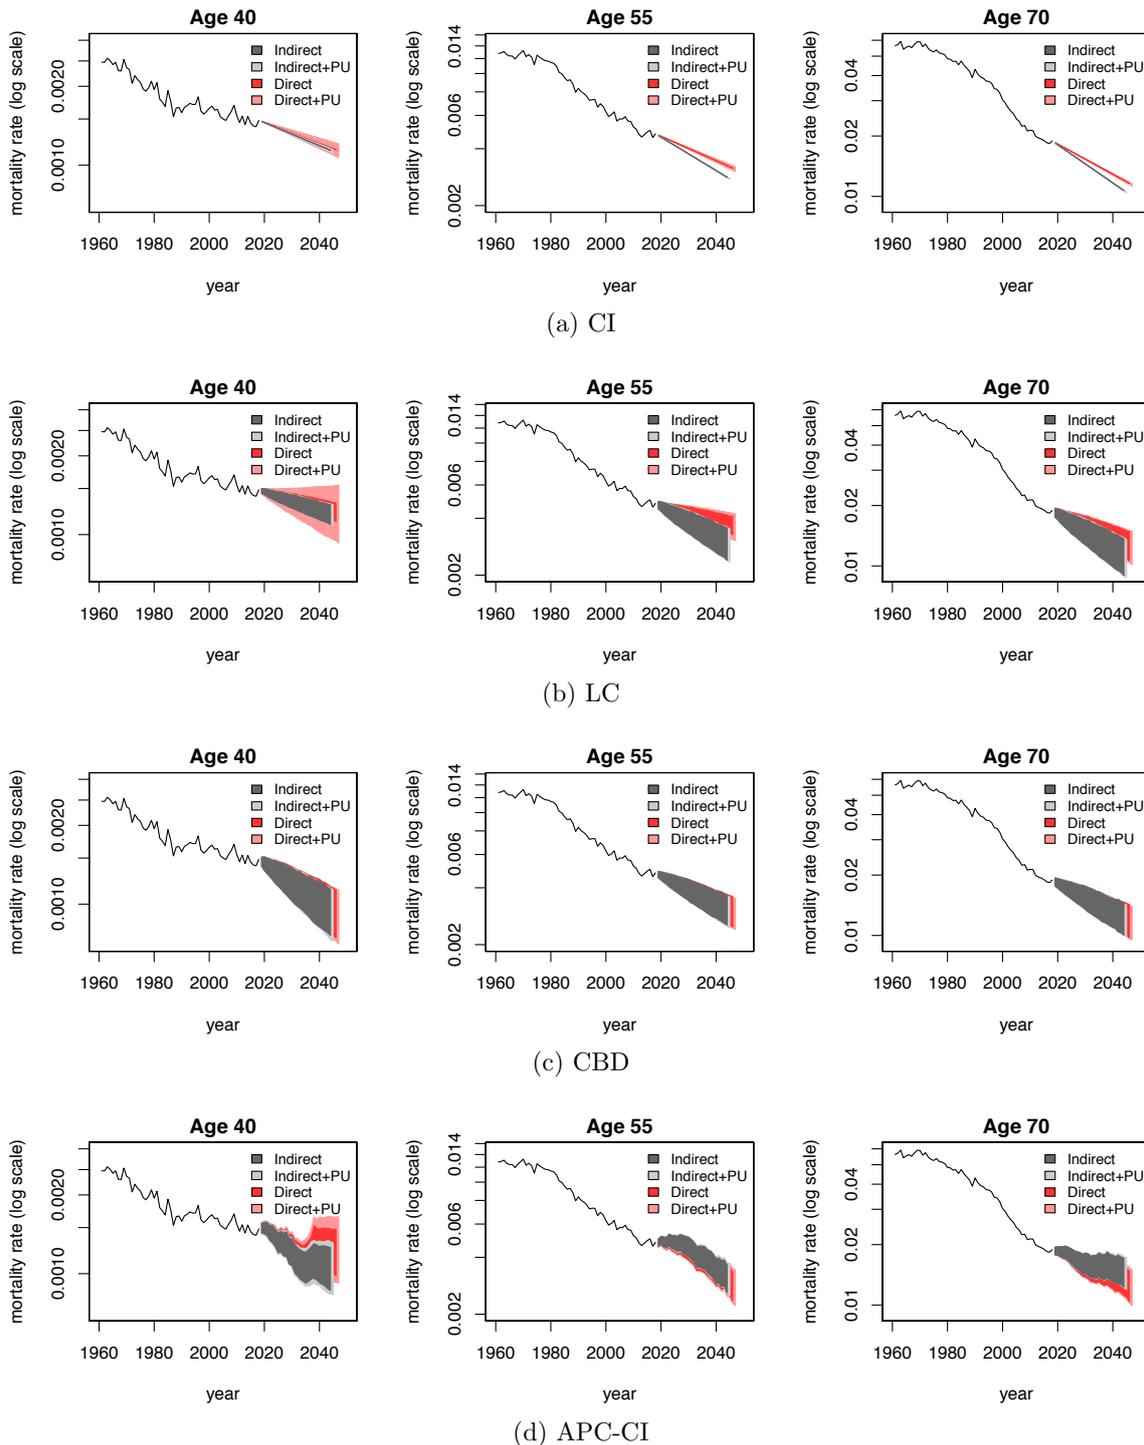


Figure 10: Fan charts for mortality rates $m_{x,t}$ at ages $x = 40, 55, 70$ from the CI, LC, CBD and APC-CI models applied to the England and Wales males population for ages 20-89 and the period 1961-2018. The solid lines show historical mortality rates for the period 1961-2018. Shades in the fan represent prediction intervals at the 95% level.

- The noticeable differences between the central projections in the indirect and direct approaches for the LC model. This is particularly visible at age 55 where the indirect approach projects a much more steeper decline in mortality than the direct fitting approach. These differences in central forecasts can be linked back to the differences in the estimates of $\beta_x^{(1)}$ produced by the two estimation approaches (recall Figure 4).
- The noticeable differences between the central projection and prediction interval at age 40 in the indirect and direct approaches for the APC-CI model. This is in part the result of the difference in central estimates of constant improvement rates at age 40, α_{40} , produced by the two estimation methods (recall Figure 8).
- The contrasting similarity in central forecasts and levels of uncertainty for the CBD under both estimation approach and with or without parameter uncertainty (see Figure 10c).

The visual inspection of the fan charts indicates that the choice of estimation approach has a material impact on the central forecasts produced by a mortality improvement model. To examine this further, we plot in Figure 11 median forecasts of mortality rates for selected cohorts produced by the indirect and direct estimation approaches applied to the CI, LC, CBD and APC-CI models. We see that, with the exception of the CBD model which produces essentially the same forecasts under the two approaches, for all other models both estimation approaches result in significantly different central forecasts. This is particularly noticeable for the younger generation born in 1998 for which the forecasts under the direct and indirect estimation approaches are very different. It is also worth highlighting that the LC model stands out as the model with the highest discrepancies between estimation approaches with the direct estimation approach producing a very unsmooth mortality schedule. As discussed before, this cautions against using the direct estimation to make projections unless the age-dependent parameter estimates are smoothed.

4.4. Robustness and stability of projections

The considerable parameter uncertainty seen for some models discussed in the previous section may have important implication for the robustness of parameter estimates as we change the period of data used in the estimation. This in turn may result in potentially unstable projections.

To investigate this potential issue, we consider the stability of forecasts over fixed horizon periods as the estimation period rolls forward through time. In each subplot in Figure 12 we show the average ten year ahead projected age-profile of mortality improvements using different 20-year rolling estimation periods. For instance, the dark purple lines labelled as stepping off year 1980 correspond to each model fitted to data from 1961-1980 and the quantity being plotted against age is the average improvement rate at each age for the next ten years, i.e, for the period 1981-1990. Similarly, the bright yellow lines labelled as stepping off year 2018 correspond to each model fitted to data from 1989-2018 to obtain the average improvement rate projection for the period 2019-2028. For a stable model, projections should progress smoothly as we change the data window. From Figure 12 we note the following:

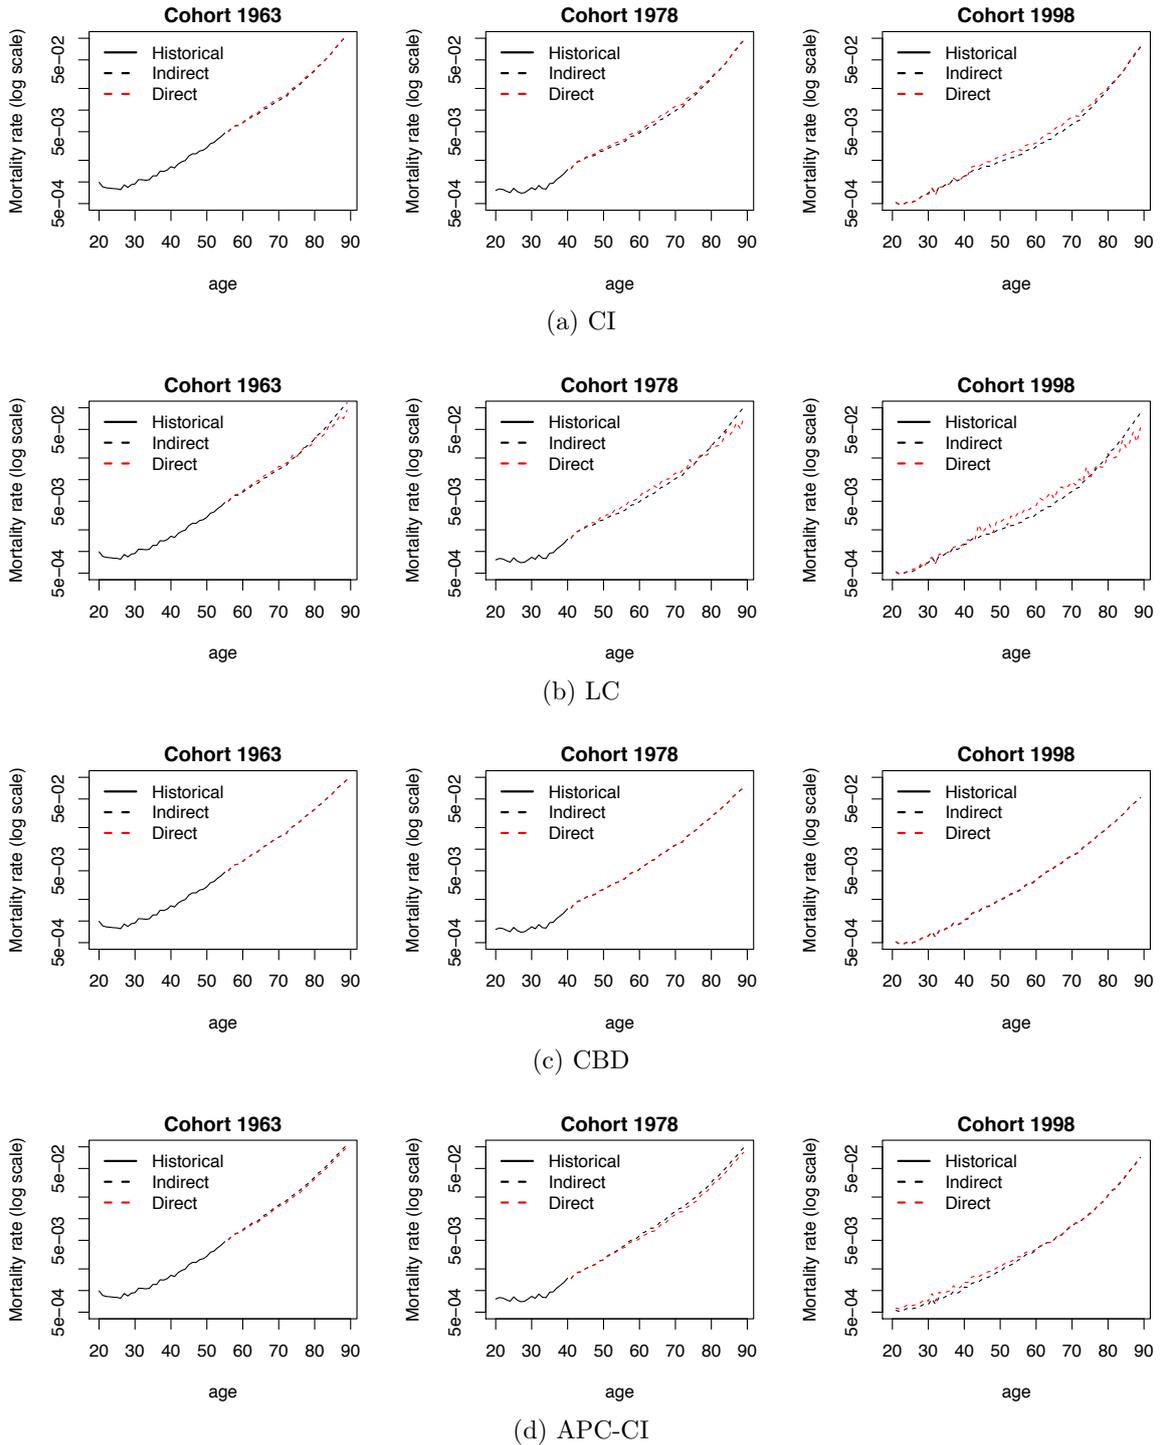


Figure 11: Median forecast of mortality rates $m_{x,t}$ for cohorts born in 1963, 1978 and 1998 obtained using an indirect and a direct estimation approach from the CI, LC, CBD and APC-CI models applied to the England and Wales males population for ages 20-89 and the period 1961-2018. The solid lines show historical mortality rates and dashed lines median forecasts.

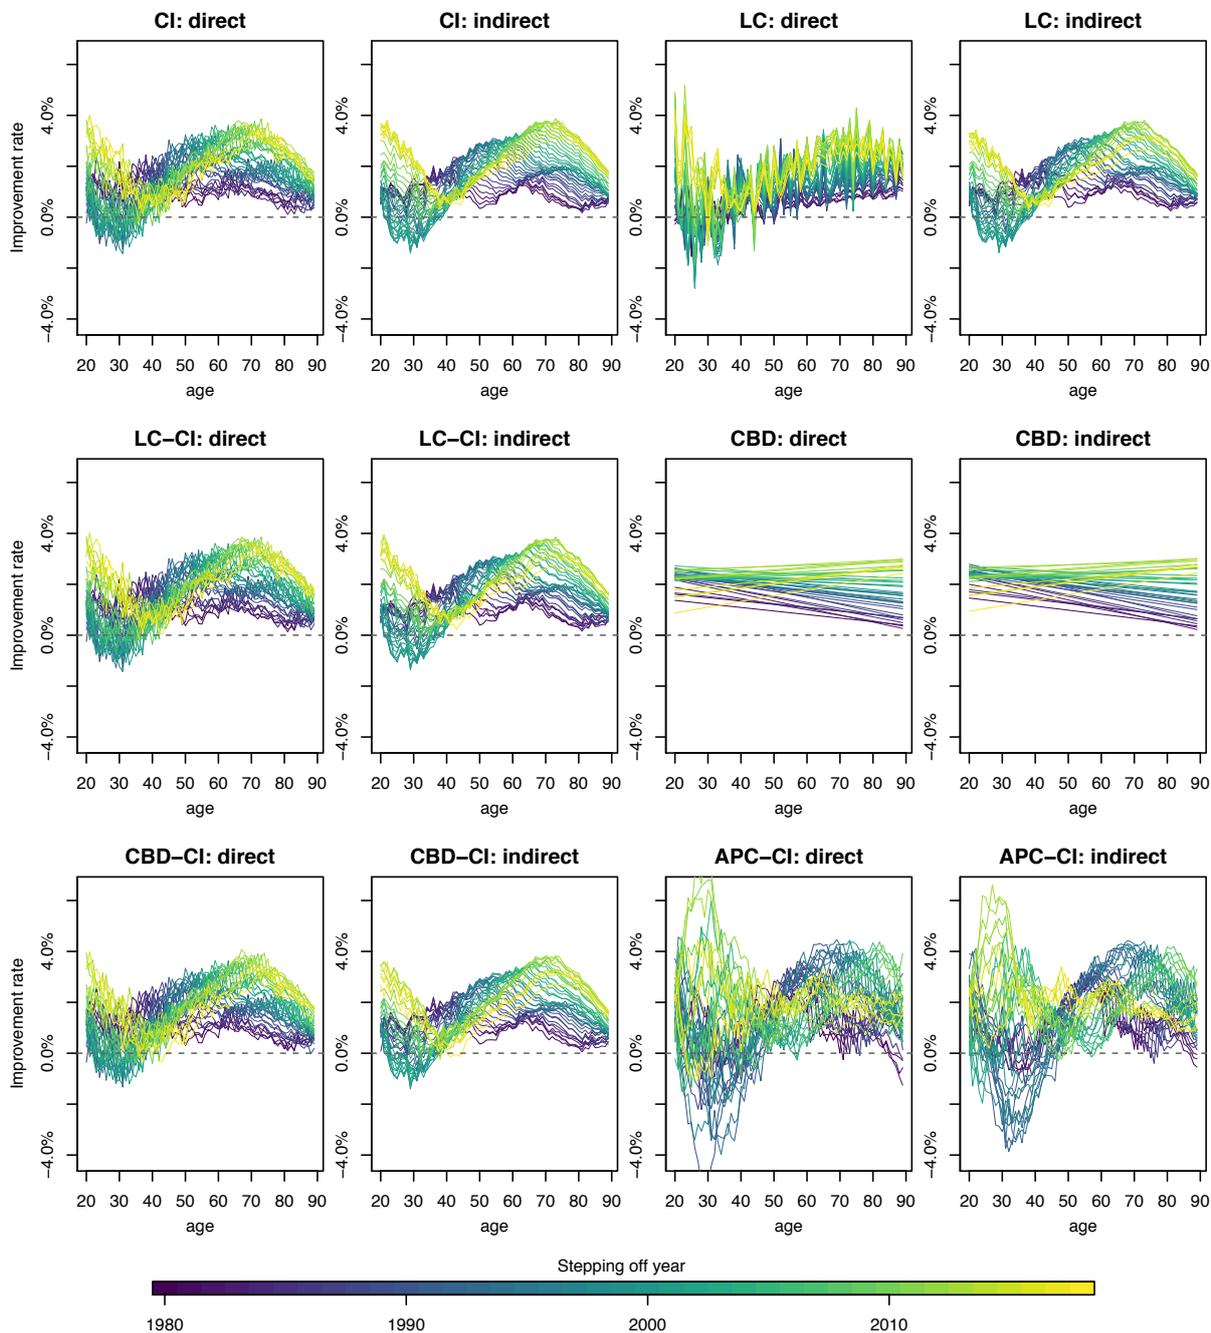


Figure 12: Average ten year ahead projected improvement rate with different stepping-of-year (20-year Rolling window).

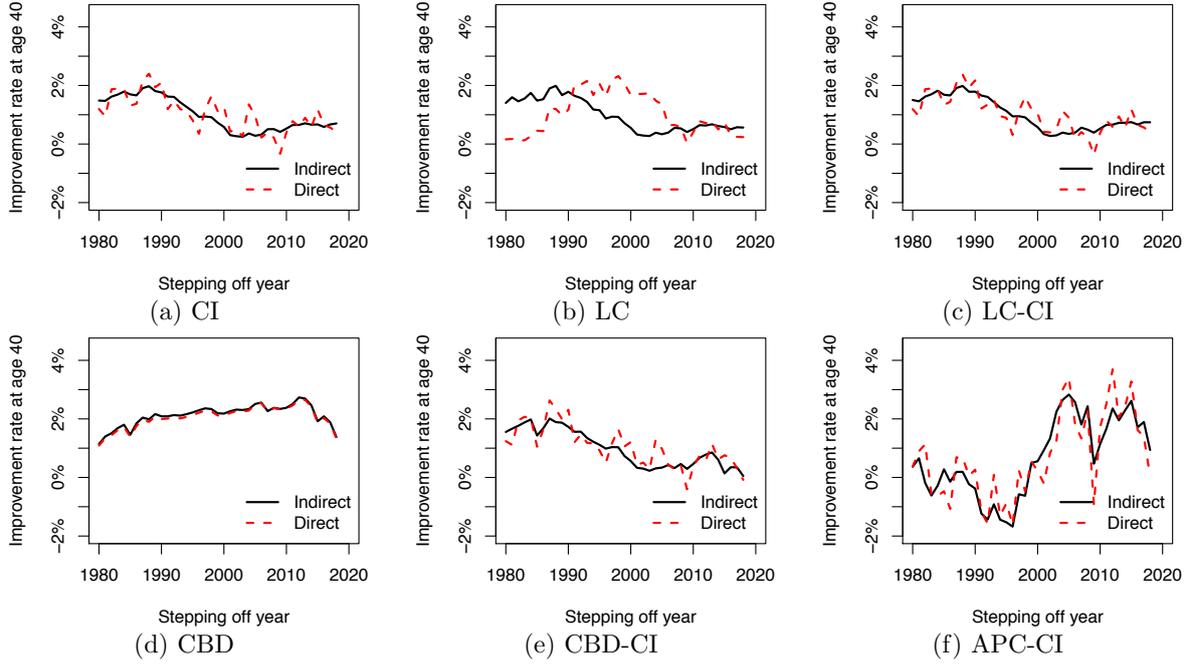


Figure 13: Average ten year ahead projected improvement rate at age 40 with different stepping-of-year (20-year Rolling window).

- The close similarity in the improvement rates projected by models CI, LC-CI and CBD-CI in spite of the varying complexity of the predictor structures. This indicates that for models with constant improvement rates and without cohort parameters, it is the estimates of α_x which dominate the central projections.
- The general lack of smoothness of the projections of models with age-related parameters under the direct modeling approach. This is particularly noticeable for the LC model which shows very unsmooth projections. This lack of smoothness can translate into instability of the projection as we change the data window. To illustrate this more clearly, in Figure 13 we show the average ten year ahead projected improvement rate at age 40 as we change the 20-year rolling estimation period. We see that with the exception of the CBD model, all other models show a general lack of stability under the direct estimation approach. This instability can be very significant for some models. For example, in the case of CI model changing the estimation period from 1984-2003 to 1985-2004 results in the projected ten year ahead average improvement passing from 0.26% to 1.35% under the direct estimation approach (see Figure 13a). This contrasts with the stability of the projections of the CI model under the indirect estimation approach in which the ten year ahead average improvement passes from 0.25% when estimated using data for the period 1984-2003 to 0.36% for data for the period 1985-2004.
- The contrasting stability of the projections of the CBD approach under both estimation approaches. This is not a surprise as the CBD is the only model which does not involve any age-related parameters.

- The noticeable different behavior of the APC-CI model as compared to the rest of the models. This is explained by the fact that this is the only model including a cohort term.
- The acceleration of improvement at older ages with time matched by a deceleration of improvements at middle ages consistent with mortality improvement rates reducing at younger ages through the 1990s.

5. Conclusions

Rates of improvement in mortality are a very natural and intuitive way of interpreting mortality data, which has led to them being widely used practically for setting and communicating assumptions regarding changes in longevity. However, they have not been studied in much depth in an academic context, possibly due to the difficulties in defining improvement rates and in fitting models robustly to data.

In this study, we have developed a more rigorous framework for the study of mortality improvement rates and its fundamental connection to models of mortality rates. This means that we can draw of the large amounts of work done to model mortality rates to obtain robust and stable estimates of improvement rates without requiring the ad hoc modeling frameworks that have been a feature of some previous studies. Furthermore, in our systematic comparison of the direct and indirect approaches for the estimation of mortality improvement rate models we have found that the direct estimation used in previous studies produces in many cases parameter estimates that are subject to considerable parameter uncertainty potentially leading to unstable projections. We have found that this robustness and instability issues are particular prevalent for model structures that include age-dependent parameters. We thus caution against the use of a direct estimation approach for predictor structures such as the LC and LC-CI models. However, if we insist in estimating these predictors structures using a direct estimation approach, it is important to smooth any age-dependent parameters (i.e., α_x and $\beta_x^{(i)}$) so that data across ages are pooled and thus reduce possible robustness and instability issues. By contrast, we find that there is less to choose between estimation techniques for predictors structures without or with fewer age-dependent parameters (for instance, the CBD model).

In summary, we find that the indirect approach to modeling mortality improvement rates is a flexible and versatile method for investigating the pattern of mortality changes in the past and for projecting mortality rates into the future. We believe it can give modelers a new perspective on existing models and potential avenues to develop new models. Perhaps most importantly however, it may allow for a common language to communicate theoretical and academic results to a wider audience of practitioners.

Acknowledgments

This project has received funding from the ARC Center of Excellence in Population Ageing Research (grant CE170100005) and from the Society of Actuaries Center of Actuarial

Excellence Research Grant 2017–2020: Longevity Risk: Actuarial and Predictive Models, Retirement Product Innovation, and Risk Management Strategies. A previous version of this paper was presented at the 2017 International Symposium “Living to 100” organised by the Society of Actuaries in Orlando, FL. We thank Ward Kingkade and the other participants of the symposium for useful comments that have helped improved this manuscript.

References

- Brouhns, N., Denuit, M.M., Van Keilegom, I., 2005. Bootstrapping the Poisson log-bilinear model for mortality forecasting. *Scandinavian Actuarial Journal* , 212–224.
- Brouhns, N., Denuit, M.M., Vermunt, J., 2002. A Poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics and Economics* 31, 373–393.
- Cairns, A.J., Blake, D., Dowd, K., 2006. A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk and Insurance* 73, 687–718.
- Callot, L., Haldrup, N., Kallestrup-Lamb, M., 2016. Deterministic and stochastic trends in the Lee–Carter mortality model. *Applied Economics Letters* 23, 486–493.
- Continuous Mortality Investigation, 2002. Working Paper 1 - An interim basis for adjusting the “92” series mortality projections for cohort effects.
- Continuous Mortality Investigation, 2009. Working Paper 38 - A prototype mortality projection model: Part one - an outline of the proposed approach.
- Continuous Mortality Investigation, 2016a. Working Paper 90 – CMI Mortality Projections Model consultation .
- Continuous Mortality Investigation, 2016b. Working Paper 91 – CMI Mortality Projections Model consultation - thecnical paper .
- Delwarde, A., Denuit, M.M., Partrat, C., 2007. Negative binomial version of the Lee-Carter model for mortality forecasting. *Applied Stochastic Models in Business and Industry* 23, 385–401.
- Dodd, E., Forster, J.J., Bijak, J., Smith, P.W.F., 2020. Stochastic modelling and projection of mortality improvements using a hybrid parametric/semiparametric age-period-cohort model. *Scandinavian Actuarial Journal* , 24.
- Dowd, K., Cairns, A.J.G., Blake, D., 2020. CBDX: a workhorse mortality model from the Cairns–Blake–Dowd family. *Annals of Actuarial Science* 14, 445—460.
- Haberman, S., Kaishev, V.K., Millosovich, P., Villegas, A.M., Baxter, S.D., Gaches, A.T., Gunnlaugsson, S., Sison, M., 2014. Longevity basis risk: A methodology for assessing basis risk. Technical Report. Cass Business School, City University of London and Hymans Robertson LLP.

- Haberman, S., Renshaw, A., 2012. Parametric mortality improvement rate modelling and projecting. *Insurance: Mathematics and Economics* 50, 309–333.
- Haberman, S., Renshaw, A., 2013. Modelling and projecting mortality improvement rates using a cohort perspective. *Insurance: Mathematics and Economics* 53, 150–168.
- Human Mortality Database, 2020. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany) URL: www.mortality.org.
- Hunt, A., Blake, D., 2020a. Identifiability in age/period mortality models. *Annals of Actuarial Science* 14, 461—499.
- Hunt, A., Blake, D., 2020b. Identifiability in age/period/cohort mortality models. *Annals of Actuarial Science* 14, 550—536.
- Hunt, A., Blake, D., 2020c. On the Structure and Classification of Mortality Models. *North American Actuarial Journal* 0, 1–20.
- Hyndman, R.J., 2014. demography: Forecasting mortality, fertility, migration and population data. URL: <http://cran.r-project.org/package=demography>.
- 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, 1–20.
- Lee, R.D., Carter, L.R., 1992. Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association* 87, 659– 671.
- Li, J.S., Zhou, R., Liu, Y., Graziani, G., Hall, R.D., Haid, J., Peterson, A., Pinzur, L., 2020. Drivers of Mortality Dynamics: Identifying Age/Period/Cohort Components of Historical U.S. Mortality Improvements. *North American Actuarial Journal* 24, 228–250.
- Li, J.S.H., Hardy, M.R., Tan, K.S., 2009. Uncertainty in mortality forecasting: An extension to the classical Lee-Carter approach. *ASTIN Bulletin* 39, 137–164.
- Mitchell, D., Brockett, P.L., Mendoza-Arriaga, R., Muthuraman, K., 2013. Modeling and forecasting mortality rates. *Insurance: Mathematics and Economics* 52, 275–285.
- Murphy, M., 2009. The “golden generations” in historical context. *British Actuarial Journal* 15, 151–184.
- Pfaff, B., 2008. Var, svar and svec models: Implementation within R package vars. *Journal of Statistical Software* 27. URL: <http://www.jstatsoft.org/v27/i04/>.
- Renshaw, A., Haberman, S., 2003. Lee-Carter mortality forecasting: a parallel generalized linear modelling approach for England and Wales mortality projections. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 52, 119–137.
- Richards, S.J., Currie, I.D., Kleinow, T., Ritchie, G.P., 2019. A stochastic implementation of the APCI model for mortality projections. *British Actuarial Journal* 24, 1–26.

Society of Actuaries, 2012. Mortality Improvement Scale BB Report .

Society of Actuaries Group Annuity Valuation Table Task Force, 1995. 1994 group annuity mortality table and 1994 group annuity reserving table. Transactions of the Society of Actuaries 47, 865–915.

Villegas, A.M., Haberman, S., Kaishev, V.K., Millossovich, P., 2017. A Comparative Study of Two-Population Models for the Assessment of Basis Risk in Longevity Hedges. ASTIN Bulletin 47, 631–679.

Villegas, A.M., Millossovich, P., Kaishev, V.K., 2018. StMoMo: An R Package for Stochastic Mortality Modelling. Journal of Statistical Software 84.

Willets, R., 2004. The cohort effect: Insights and explanations. British Actuarial Journal 10, 833–877.

Appendix A. Parameter constraints

Table A.2 presents the parameter constraints used in estimating the models in Table 1. In Table A.2 we note that the constraints in the second column are applied when using the direct approach to the fitting of a mortality improvement rate model while the constraints in the third column are applied when using the indirect approach to model fitting. We also note that the “level” constraints in the improvement rate predictor structure,

$$\sum_t \kappa_t^{(i)} = 0,$$

become constraints of the deterministic trends in the mortality rate predictor structure,

$$\sum_t (t - \bar{t}) K_t^{(i)} = 0.$$

These both sets of constraints make sense intuitively, since the $\alpha_x t$ term in mortality improvement rate model explains any constant improvements in the historical data so the $K_t^{(i)}$ are constrained to only explain possible deviations from this constant improvement. Similarly, for the APC-CI the improvement rate predictor constraints on the cohort effect,

$$\sum_y \gamma_y = 0, \sum_y (y - \bar{y}) \gamma_y = 0,$$

become constraints

$$\sum_y (y - \bar{y}) \Gamma_y = 0, \sum_y (y - \bar{y})^2 \Gamma_y = 0$$

in the mortality rate predictor structure.

Table A.2: Parameter constraints for the structures considered in this paper.

Model	Improvement Model ($\eta_{x,t}$)	Equivalent Mortality Model ($\tilde{\eta}_{x,t}$)
CI	-	-
LC	$\sum_x \beta_x^{(1)} = X$	$\sum_x \beta_x^{(1)} = X, K_0^{(1)} = 0$
LC-CI	$\sum_x \beta_x^{(1)} = X, \sum_t \kappa_t^{(1)} = 0$	$\sum_x \beta_x^{(1)} = X, K_0^{(1)} = 0,$ $\sum_t (t - \bar{t}) K_t^{(1)} = 0$
CBD	-	$K_0^{(1)} = 0, K_0^{(2)} = 0$
CBD-CI	$\sum_t \kappa_t^{(1)} = 0, \sum_t \kappa_t^{(2)} = 0$	$K_0^{(1)} = 0, K_0^{(2)} = 0, \sum_t (t - \bar{t}) K_t^{(1)} = 0,$ $\sum_t (t - \bar{t}) K_t^{(2)} = 0$
APC-CI	$\sum_t \kappa_t^{(1)} = 0, \sum_y \gamma_y = 0,$ $\sum_y (y - \bar{y}) \gamma_y = 0$	$K_0^{(1)} = 0, \sum_t (t - \bar{t}) K_t^{(1)} = 0,$ $\Gamma_{-X} = 0, \sum_y (y - \bar{y}) \Gamma_y = 0,$ $\sum_y (y - \bar{y})^2 \Gamma_y = 0$

Appendix B. iMoMo: An R package for estimating mortality improvement rate models

We have implemented the estimation methods presented in this paper in the R package **iMoMo** which is available in Github at <https://github.com/amvillegas/iMoMo>. **iMoMo** is an extension of the R package **StMoMo** and as such provides tools for estimating, forecasting, simulating and bootstrapping improvement rate models. The interested reader is referred to Villegas et al. (2018) for an extensive discussion and illustration of the capabilities of **StMoMo**, most of which extend naturally to **iMoMo**. Here we illustrate some key functions in **iMoMo** and show how they can be used to replicate some of the key results in the body of the paper.

The development version of the package can be installed using the following commands:

```
install.packages("devtools")
devtools::install_github("amvillegas/iMoMo")
```

The code below defines the six predictor structures in Table 1:

```
library(iMoMo)
#CI: next = ax
CIId <- iMoMo(staticAgeFun = TRUE, periodAgeFun = NULL, type = "direct")
CIi <- iMoMo(staticAgeFun = TRUE, periodAgeFun = NULL, type = "indirect")
```

```

#LC:  $nxt = bx*kt$ 
LCd <- StMoMo2iMoMo(lc(), type = "direct")
LCi <- StMoMo2iMoMo(lc(), type = "indirect")

#LC-CI:  $nxt = ax + bx*kt$ 
LC_CId <- lci(type = "direct")
LC_CII <- lci(type = "indirect")

#CBD:  $nxt = kt^1 + (x - \bar{x})kt^2$ 
CBDd <- cbdi(type = "direct")
CBDi <- cbdi(type = "indirect")

#CBD-CI:  $nxt = ax + kt^1 + (x - \bar{x})kt^2$ 
constCBDx <- function(ax, bx, kt, b0x, gc, wxt, ages){
  #\sum kt[i, ] = 0
  ci <- rowMeans(kt, na.rm = TRUE)
  ax <- ax + ci[1] + ci[2] * bx[, 2]
  kt[1, ] <- kt[1, ] - ci[1]
  kt[2, ] <- kt[2, ] - ci[2]
  list(ax = ax, bx = bx, kt = kt, b0x = b0x, gc = gc)
}
f2 <- function(x, ages) x - mean(ages)
CBD_CId <- iMoMo(periodAgeFun = c("1", f2), constFun = constCBDx,
  type = "direct")
CBD_CII <- iMoMo(periodAgeFun = c("1", f2), constFun = constCBDx,
  type = "indirect")

#APC-CI:  $nxt = ax + kt + gc$ 
APC_CId <- apci(type = "direct")
APC_CII <- apci(type = "indirect")

```

In the above code, for each predictor, we have defined two versions depending on whether the model is to be fitted using a direct approach or an indirect approach. The code above also illustrates the three ways in which a improvement rate model can be defined:

1. Using the `iMoMo` function which works in a similar manner to the `StMoMo` function in package `StMoMo`.
2. By transforming a mortality rate model into an improvement rate model using function `StMoMo2iMoMo`.
3. By invoking a predefined model. `iMoMo` provides predefined functions for the LC-CI, CBD and APC-CI predictors which are implemented in the functions `lci`, `cbdi` and `apci`, respectively.

To fit the models to data we first need to extract the data for England and Wales from the Human Mortality Database (2020) using function `hmd.mx` of the demography package (Hyndman, 2014) with the code:

```

library(demography)
EWdata <- hmd.mx(country = "GBRTENW", username = username,
                 password = password, label = "England Wales")
dataStMoMo <- StMoMoData(EWdata, series = "male")
dataStMoMo$Dxt <- round(dataStMoMo$Dxt)

```

We note that the `username` and `password` above are for the Human Mortality Database and should be replaced appropriately.

We can now fit, for example, the LC-CI model for ages 20-89 and years 1961-2018 as follows:

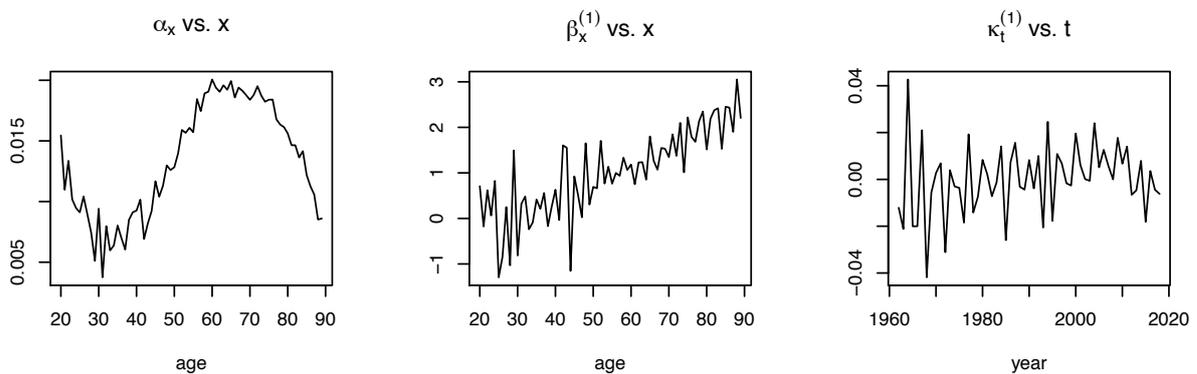
```

LC_CId_fit <- fit(LC_CId, data = dataStMoMo, ages.fit = 20:89,
                 years.fit = 1961:2018)
LC_CII_fit <- fit(LC_CII, data = dataStMoMo, ages.fit = 20:89,
                 years.fit = 1961:2018)

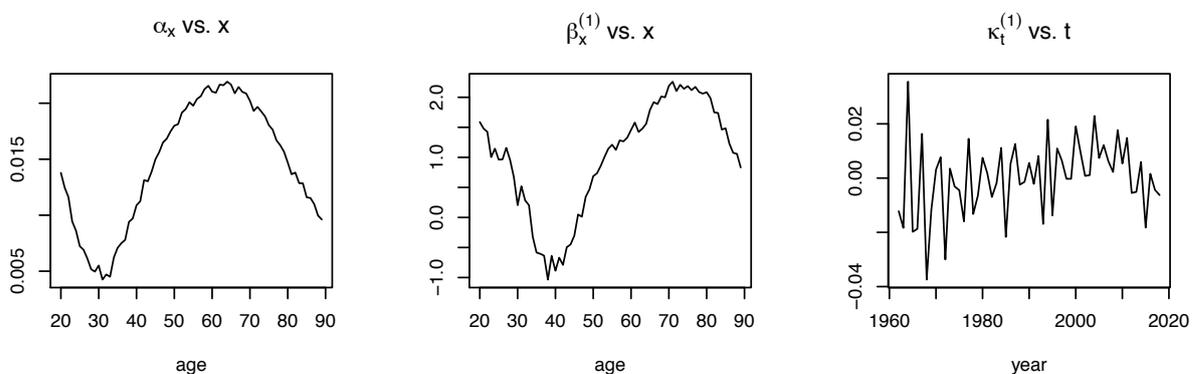
```

The other models can be fitted similarly. We can then plot the parameter estimates as follows:

```
plot(LC_CId_fit, nCol = 3)
```



```
plot(LC_CII_fit, nCol = 3)
```

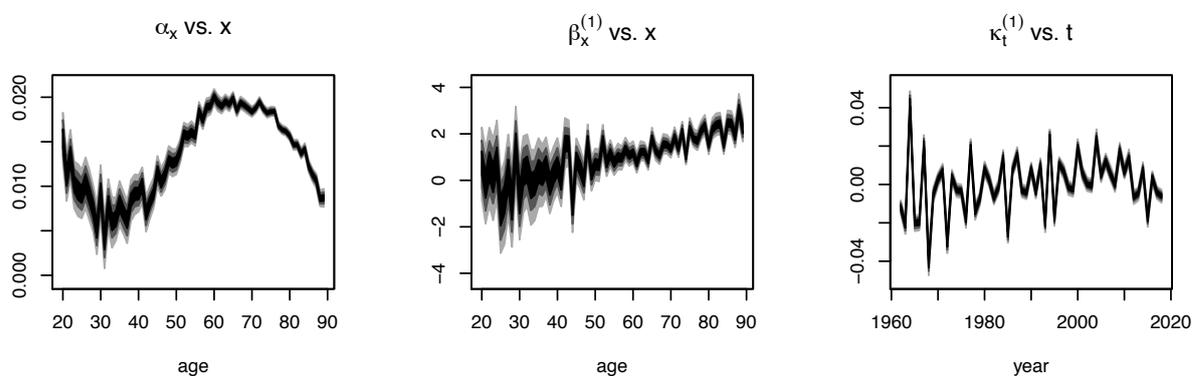


Finally, 1000 bootstrap samples of the LC-CI model can be produced with the code:

```
LC_CId_boot <- bootstrap(LC_CId_fit, nBoot = 1000)
set.seed(bootSeed)
LC_CII_boot <- bootstrap(LC_CII_fit, nBoot = 1000)
```

We note that the bootstrap is a computationally intensive procedure so the code above can take several minutes to run. The bootstrapped model can then be plotted as follows:

```
plot(LC_CId_boot, nCol = 3)
```



```
plot(LC_CII_boot, nCol = 3)
```

