

Mortality Improvement Rates: Modeling and Parameter Uncertainty

Andrew Hunt

Pacific Life Re, London, United Kingdom

Andrés M. Villegas

School of Risk and Actuarial Studies and

ARC Centre of Excellence in Population Ageing Research (CEPAR)

UNSW Sydney, Australia

Presented at the Living to 100 Symposium

Orlando, Fla.

January 4–6, 2017

Copyright © 2017 by the Society of Actuaries.

All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

Mortality Improvement Rates: Modeling and Parameter Uncertainty

Andrew Hunt^a, Andrés M. Villegas^b

^a*Pacific Life Re, London, United Kingdom*

^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. (2002), 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 U.S. population and the England and Wales population.

Keywords: Mortality improvements; Mortality forecasting; Parameter uncertainty

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 aging of populations across the world. One strand of efforts to meet 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 historical data and expert judgment.

Email addresses: a.o.d.hunt.00@cantab.net (Andrew Hunt), a.villegas@unsw.edu.au (Andrés M. Villegas)

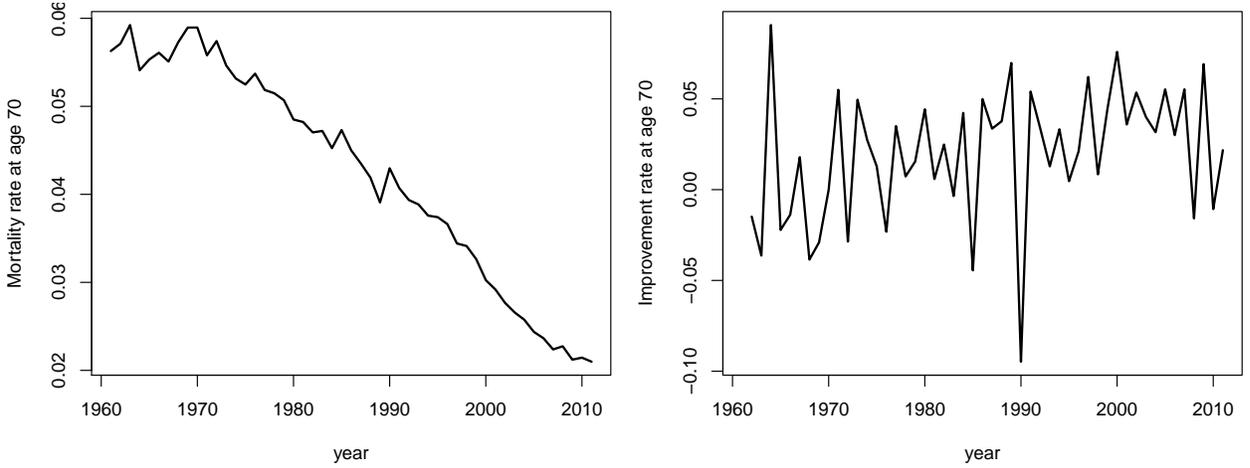
This project has received funding from the ARC Center of Excellence in Population Ageing Research (grant CE110001029). The work in this study was started before Dr. Hunt commenced work at Pacific Life Re, and any opinions expressed in this paper are held in a personal capacity and should not be construed as the views of Pacific Life Re or related companies.

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 (2015c) and implemented in Villegas et al. (2016). More specifically, as discussed in Hunt and Blake (2015c), 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 numbers 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 η_{xt} by means of a log-link function, i.e.,

$$\ln m_{x,t} = \eta_{x,t}. \tag{1}$$

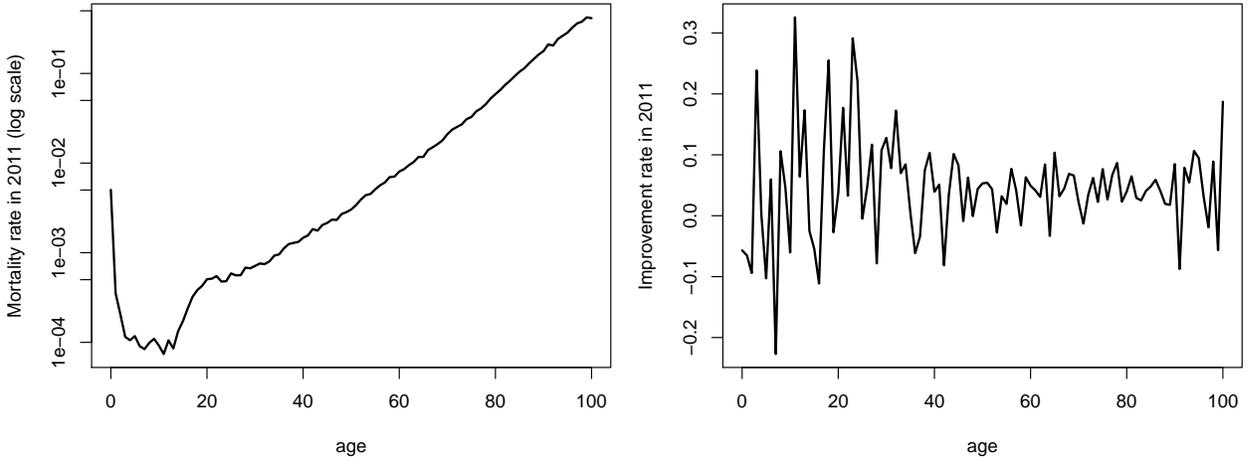
In contrast, practitioners are often interested primarily in the mortality improvement rates, usually defined by $-\ln\left(\frac{\mu_{x,t}}{\mu_{x,t-1}}\right)$, $-\ln\left(\frac{m_{x,t}}{m_{x,t-1}}\right)$ or $1 - \frac{q_{x,t}}{q_{x,t-1}}$. This is because the changes in mortality rates are what is of interest when assessing longevity risk for an insurer or pension scheme. However, improvement rates are usually estimated using the largest data set available over a long time period—often the national population—in order to give reliable estimates. Such a data set will usually have mortality rates very different from the population of interest. Nonetheless, by considering mortality improvement rates, inferences made using these large data sets can still be used for smaller sub-populations, albeit potentially subject to longevity “basis risk” (see Haberman et al. (2014)). 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 populations in different countries. In the United Kingdom, the concept of mortality improvement rates became widely adopted among actuaries as a result of the Continuous Mortality Investigation (2002) and has continued with the development 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 at 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, while the age shape of mortality improvement rates is very noisy and displays considerable heteroskedasticity across ages.



(a) Crude Mortality Rates at Age 70

(b) Crude Improvement Rates at Age 70



(c) Crude Mortality Rates in Year 2011

(d) Crude Improvement Rates in Year 2011

Fig. 1. England and Wales Male Mortality and Improvement Rates

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 nonstandard link function.

Such an approach does not present any theoretical problems; however, some practical issues need to be considered. First, the distribution of the response variables is highly

nonstandard, 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, 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, while 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 rates. 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 (see , Haberman and Renshaw (2012)).
- 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 organized 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, as well as the potential to allow for constant improvement rates in mortality models. We then develop techniques for fitting improvement rate models to data and apply them to the mortality experiences of England and Wales and of the United States in Sections 4 and 5. 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 age and period terms in improvement rate models and briefly look at projections from improvement rate models. Finally, in Section 6, 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 $m_{x,t} = \mathbb{E}(D_{x,t})/e_{x,t}$. 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

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

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

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 (2015c), 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 that modulate the corresponding period functions;² and

²These age functions can be nonparametric (having form determined entirely by the data) or parametric (having a predefined functional form), as discussed in Hunt and Blake (2015c).

- γ_y is a cohort function describing systematic differences in the rate of improvement which depend upon a cohort's year of birth, $y = t - x$.

Unlike Mitchell et al. (2013) and Haberman and Renshaw (2012), we do not model $-\Delta m_{x,t}$ directly, since the mortality improvement rates in this specification do not follow a standard probability distribution. They are also highly heteroskedastic, 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 rewritten 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 (2015c), 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

³One drawback of using a Poisson model for the death counts, common to models of mortality rates and improvement rates, is that it assumes that the variance of an observation is equal to its expectation. Such overdispersion can be dealt with by using an overdispersed Poisson model in a generalized nonlinear 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); 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 generalized linear model 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. Estimation and Equivalent Mortality Rate Structure

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 rewritten 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 (2015c)). Therefore, we see that conventional mortality rates models are identical to improvement rate 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 (2015c) with a term that is nonparametric in age and linear in time. Therefore, every mortality rate model discussed in Hunt and Blake (2015c) has an extended version that includes a constant improvement rate term, which is equivalent to using the same predictor structure for mortality improvement rates rather than mortality rates.

To estimate the improvement rate model in Equations (5) and (6), we can then estimate the equivalent mortality rate model defined by (7) with the constraints in (8) and recover the parameters of the improvement rate model using the relationships in (9). Hence, we can use standard techniques to fit mortality rate models to data and convert these to models of the improvement rate. In this paper, we follow such an approach and estimate the models using the R package StMoMo (Villegas et al., 2016), which enables the fitting of general age/period/cohort mortality rate models.

3.3. Models With Constant Improvement Rates

A key question when deciding on the form of the predictor structure in Equation (4) is whether to include an α_x term. This term represents the average rate of improvement at each age over the period of the historic data and, when the model is projected, will give a constant component to the rate of improvement in mortality in the future.

Some authors take issue with this and believe that it conflicts with the requirements of biological reasonableness.⁵ In particular, Haberman and Renshaw (2012) show that the α_x static age function in a standard mortality rate model disappears when the first derivative is taken to obtain an improvement rate model. Furthermore, there are legitimate questions as to what form any constant rate of improvement should take. Including a nonparametric α_x term in the predictor structure of Equation (4) will assume that the average rates of improvement observed over the period of the historic data at each age will continue indefinitely into the future. Instead, it may be desirable to impose a parametric structure on the age shape, in the same way that models from the Cairns-Blake-Dowd family impose a parametric structure on the shape of mortality rates. Such a choice can result in a more parsimonious model, making it easier to fit to data, and will have generally simpler identifiability issues. However, unlike models of mortality rates, we are unlikely to have as strong an intuition as to what the shape of mortality improvements will be a priori when deciding on an improvement rate model.

Ultimately, the decision whether to include an α_x term in the model of mortality improvement rates is largely subjective and will depend on the preferences of the model user.

3.4. Estimation and Use of Crude Mortality Rates

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 defines improvement rates in terms of the *crude* mortality rates, $\hat{m}_{x,t} = d_{x,t}/e_{x,t}$, and so uses

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

in Mitchell et al. (2013), as opposed to the model *fitted* mortality rates, $m_{x,t} = \mathbb{E}(D_{x,t})/e_{x,t}$, in Equation (3). Converting this to a Poisson formulation of the model, we see that this gives⁶

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

as opposed to

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

⁵The concept of biological reasonableness was introduced in Cairns et al. (2006b) and defined as “a method of reasoning used to establish a causal association (or relationship) between two factors that is consistent with existing medical knowledge.”

⁶Equation (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 $e_{x,t}\hat{m}_{x,t-1}\exp(-\eta_{x,t})$

from Equation (6) in our formulation. We note that the Mitchell et al. (2013) form of an improvement rate model can also be fitted using standard Poisson generalized (non-)linear modeling techniques by setting $\ln \hat{m}_{x,t-1}$ as an offset within the generalized (non-)linear model predictor structure.

The difference between formulations (11) and (12) of an improvement rate model, although subtle, has profound consequences, as we will discuss in the remainder of this paper. From now onward, when referring to the estimation of the parameters in the predictor $\eta_{x,t}$, we will say that we use a “crude” estimation approach whenever we assume (11) and a “fitted” estimation approach whenever we assume (12). This naming convention reflects the fact that formulation (11) is based on defining improvement rates in terms of crude mortality rates (see, Equation (10)) while formulation (12) is based on defining improvement rates in terms of fitted (or model-estimated) mortality rates (see, Equation (3)).

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 that 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 $m_{x,t} = 0.5\%$ per year. 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 percent 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 percent 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 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 percent. To get levels of certainty in our estimates of improvement rates that are comparable 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 percent), 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 semiparametric 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 resamples 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. (2016, 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 that used in a model of mortality rates. Therefore, it is possible to use similar time series techniques.⁷ However, since the model is now one of improvement rates rather than mortality rates, the demographic significance of the parameters is now different, which will influence our choice of projection model.⁸

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):

$$\Delta^d \boldsymbol{\kappa}_t = \mathbf{C} + \mathbf{D}t + \sum_{i=1}^p \mathbf{A}_i \Delta^d \boldsymbol{\kappa}_{t-1} + \boldsymbol{\xi}_t^\kappa, \quad \boldsymbol{\xi}_t^\kappa \sim N(\mathbf{0}, \Sigma), \quad (13)$$

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 (13).

As for the cohort effects, we can assume in general that they follow an autoregressive integrated moving average model, ARIMA(p, q, d), with drift, i.e.,

$$\Delta^d \gamma_c = \delta_0 + \phi_1 \Delta^d \gamma_{c-1} + \dots + \phi_p \Delta^d \gamma_{c-p} + \epsilon_c + \delta_1 \epsilon_{c-1} + \dots + \delta_q \epsilon_{c-q}, \quad (14)$$

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 (13) and (14) 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:

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

⁷Allowing for similar issues as described in Hunt and Blake (2015a,b) in order to obtain “well-identified” projections that do not depend on the arbitrary identifiability constraints chosen when fitting the model.

⁸Demographic significance is defined in Hunt and Blake (2015c) as the interpretation of the components of a model in terms of the underlying biological, medical or socioeconomic causes of changes in mortality rates which generate these rates.

Table 1. Model Structures Considered in This Paper

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

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, 1961–2011, Ages 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 1961–2010 and ages 20–89, obtained from the Human Mortality Database (2014). Of particular interest in this discussion are the inclusion of constant improvement rates, the implications of using a “fitted” or “crude” estimation approach, the impact of parameter uncertainty, and the choice of time series model for the period indexes.

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

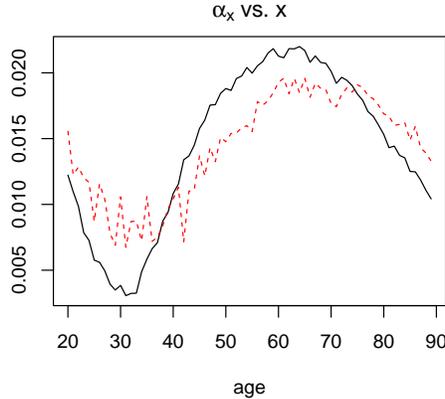


Fig. 2. Parameters for Model CI: England and Wales Males, Ages 20–89, 1961–2011

form and in Callot et al. (2014) in its mortality rate form. Model CBD is the improvement rate equivalent of the two-factor model introduced in Cairns et al. (2006a) and in mortality rate form is equivalent to the “CBDX” model discussed in Hunt and Blake (2015a). 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 a model has recently been considered by the Continuous Mortality Investigation (2016a,b) in their proposed update of the widely used CMI mortality projection model.

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 1961–2011 and ages 20–89 are shown in Figures 2–7. In these figures, black continuous lines depict parameter estimates obtained with the “fitted” estimation approach introduced in this paper, while red dashed lines depict parameter estimates obtained with the “crude” estimation approach discussed in Section 3.4. From Figures 2–7, we note the following:

- The noticeable differences in the estimated age-dependent parameters α_x and β_x under the two alternative estimation approaches, with the “crude” approach producing very ragged estimates. This is a clear reflection of the parameter uncertainty issues of working with improvement rates directly, as discussed in Section 3.5.
- The contrasting close alignment between the estimates of period indexes $\kappa_t^{(1)}$, $\kappa_t^{(2)}$ and the cohort index γ_{t-x} under the two alternative estimation approaches. This suggests that the choice of estimation approach will be of greater relevance for model structures with constant improvement rates (e.g., CI, LC-CI, CBD-CI, APC-CI) or

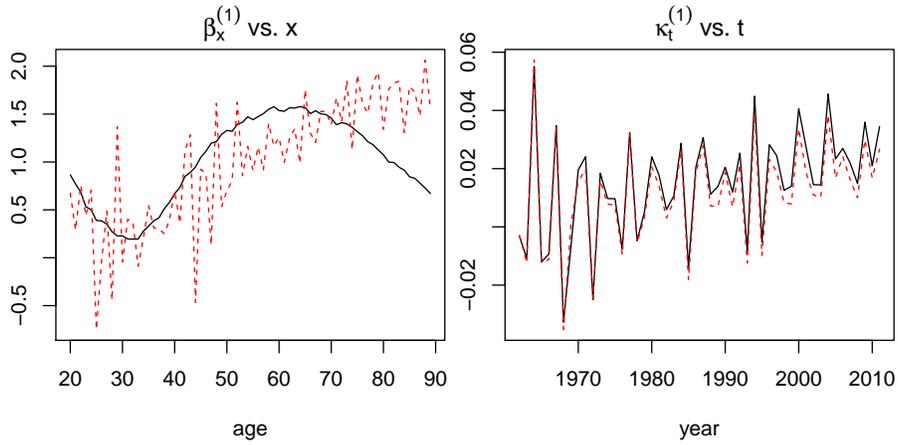


Fig. 3. Parameters for Model LC: England and Wales Males, Ages 20–89, 1961–2011

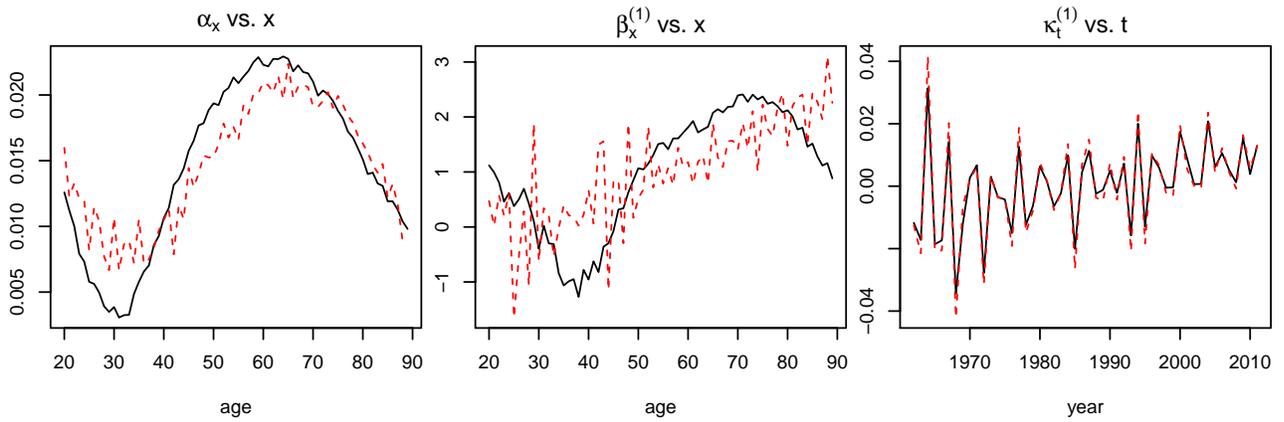


Fig. 4. Parameters for Model LC-CI: England and Wales Males, Ages 20–89, 1961–2011

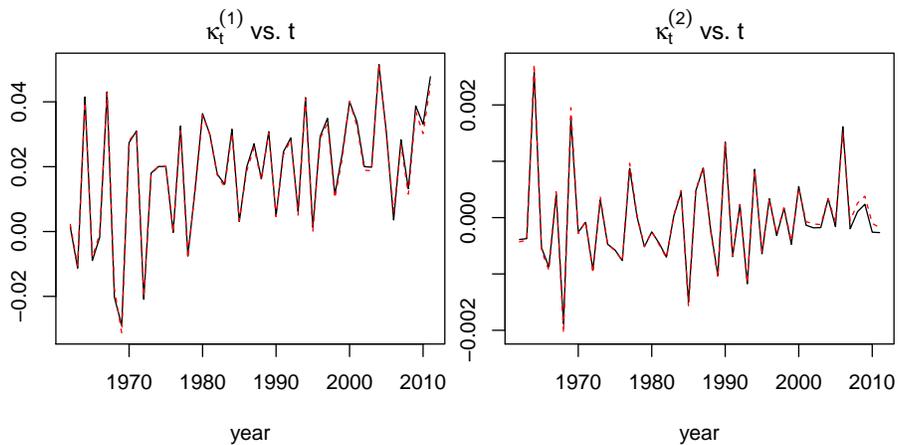


Fig. 5. Parameters for Model CBD: England and Wales Males, Ages 20–89, 1961–2011

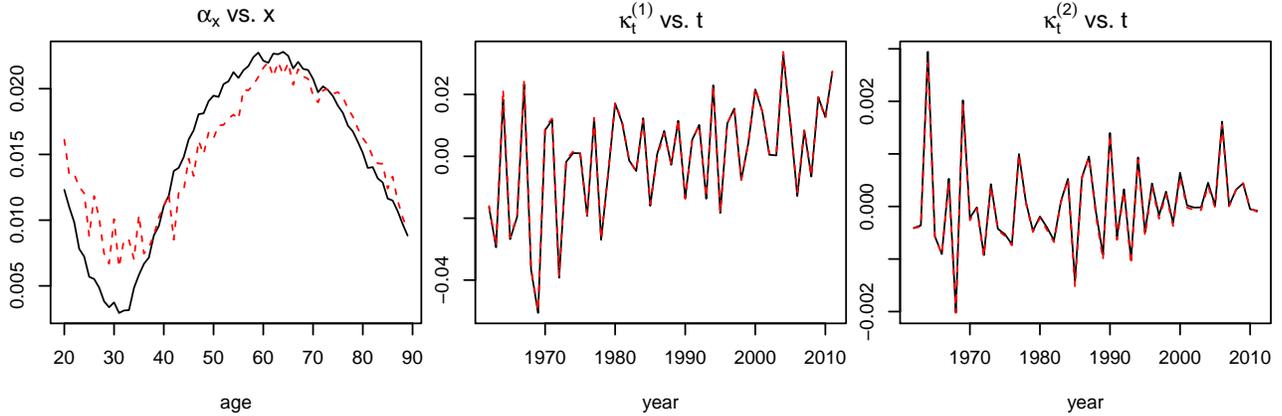


Fig. 6. Parameters for Model CBD-CI: England and Wales Males, Ages 20–89, 1961–2011

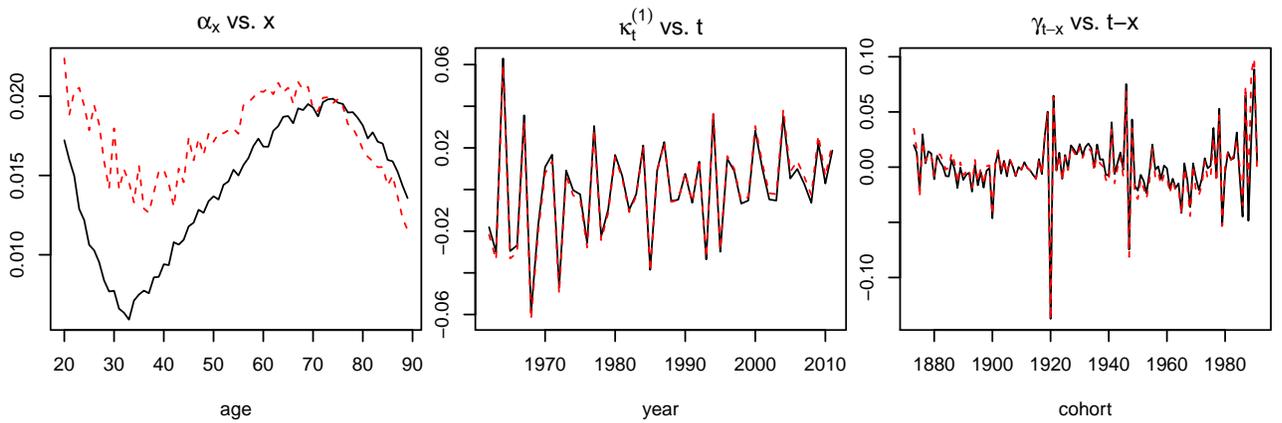


Fig. 7. Parameters for Model APC-CI: England and Wales Males, Ages 20–89, 1961–2011

“non-parametric” age modulating function (e.g., LC and LC-CI) than for purely “parametric” structures such as the CBD.

- The clear interpretation of the α_x term as average mortality improvements, indicating that in England and Wales, mortality improvement rates over the 1961–2011 period have ranged from about 0.75 percent per year between ages 20 and 30 and about 2 percent per year at ages 60 to 70.
- The similarity in the α_x parameters of the CI, LC-CI and CBD-CI models using the “fitted” approach. From the interpretation above, this should not be surprising and is comparable to the similarity in static age functions in different models. However, it is worth noting that this similarity does not follow through to the parameter values fitted using the “crude” approach.
- The ease of interpretation of the primary period index $\kappa_t^{(1)}$, 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. This is in contrast with mortality level models where it is difficult to link the value of $\kappa_t^{(1)}$ to quantities with intuitive practical relevance.

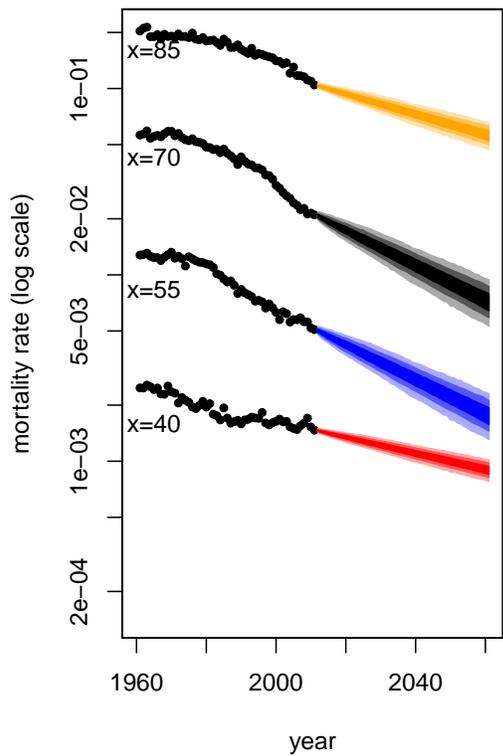
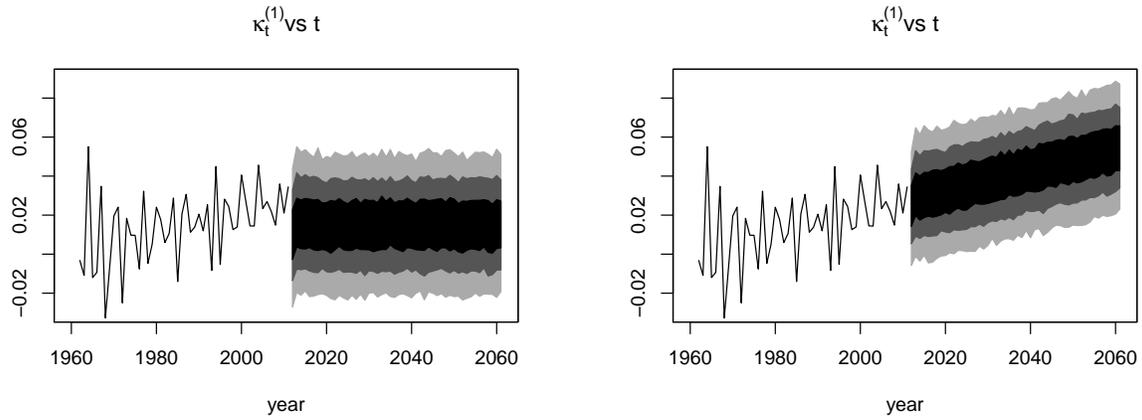
- The interpretation of the cohort parameters γ_{t-x} in Figure 7 as average deviations in improvement rates. In particular, we see that the so-called golden generation born in the inter-war period (see Willets (2004) and Murphy (2009)) has experienced mortality improvements of around 1 to 2 percent per year greater 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.
- The noticeable upward trend in the primary period index $\kappa_t^{(1)}$ in most models, reflecting the faster improvements observed in recent years in England and Wales.⁹ This trend raises questions as to how to model the period indexes appropriately while ensuring the plausibility of the projected mortality rates.

4.2. Modeling the Period Indexes

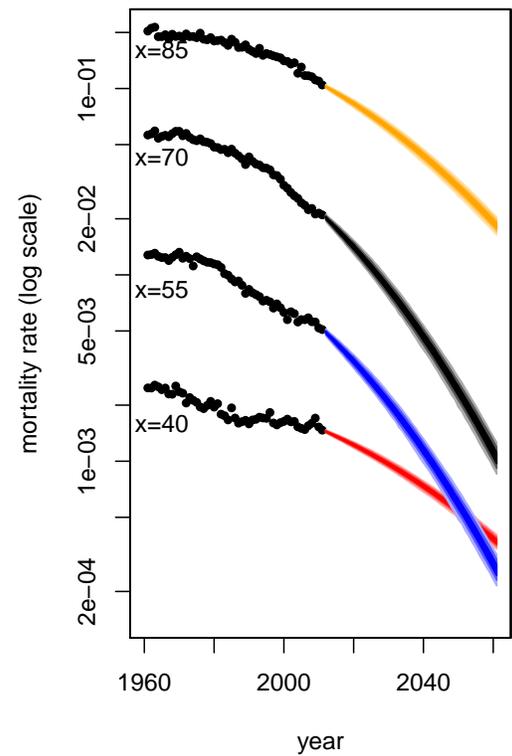
We consider two alternative time series models to project the period indexes: a VAR(1) model as in Haberman and Renshaw (2012), and a VAR(1) around a linear trend to acknowledge the upward trend seen in the primary period index for all models. To illustrate the implications of the choice of time series model, in Figure 8 we plot for both time series specifications 95 percent prediction intervals for the period index $\kappa_t^{(1)}$ of the LC model along with the matching prediction intervals for mortality rates at ages 40, 55, 70 and 85. Recalling that the $\kappa_t^{(1)}$ can be interpreted as average improvement rates, we see that under the VAR(1) model (Figure 8a), improvement rates are predicted to range from -2 to 6 percent per year, while under the VAR(1) model around a linear trend (Figure 8b), improvement rates are forecast to increase steadily to reach a value of between 2 and 8 percent per year by 2060. This upward trend induces a quadratic behavior in the log-mortality rate forecasts and accelerating improvements in mortality rates. In some models, the faster improvement rates at some ages than others result in “mortality crossovers” (as shown in Figure 8b), where younger ages are forecast to have higher mortality than older ages. This may be felt to be unrealistic and conflicts with our desire for “biologically reasonable” projections. Hence, we restrict our projections to the VAR(1) model to avoid implausible projection patterns.¹⁰

⁹We note that since 2011, there have been several years of slower improvements than were expected, so this trend may not have continued beyond our data.

¹⁰However, mortality crossovers were not observed in models such as the CBD and CDB-CI models, indicating that they may be dealt with by imposing some form of parametric structure on the age-shape of the improvements. It is also possible that mortality crossovers could occur using the VAR(1) model, e.g., as when extending projections further into the future.



(a) LC, VAR(1)



(b) LC, VAR(1) With Trend

Fig. 8. Fan Charts of the Primary Index $\kappa_t^{(1)}$ and of Mortality Rates $m_{x,t}$ at Ages $x = 40, 55, 70, 85$ From the LC Model: England and Wales Males, Ages 20–89, 1961–2011, Using the “Fitted” Estimation Approach and Different Time Series Models

Note: Shades in the fans represent prediction intervals at the 50%, 80% and 95% levels.

4.3. 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 semiparametric bootstrapping approach introduced in Brouhns et al. (2005). Figure 9 presents fan charts depicting the 50, 80 and 95 percent 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 six models considered in this paper.

From Figure 9a, we see that using a “crude” 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 rate parameters α_x than using the “fitted” estimation approach introduced in this paper. For instance, under the “crude” approach, the 95 percent confidence interval of the improvement rate at age 40, α_{40} , is (0.81%, 1.26%), which is three times wider than the equivalent (1.00%, 1.16%) under the “fitted” approach. Similarly, Figure 9b shows that nonparametric age-modulating parameters $\beta_x^{(1)}$ also suffer from considerable parameter uncertainty under the “crude” estimation approach. It is also interesting to note that, in many cases, the confidence intervals of the age parameters under the “crude” approach do not contain the “fitted” parameter estimates. Therefore, it is not simply a case that the “crude” 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 “fitted” 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 “crude” 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.

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

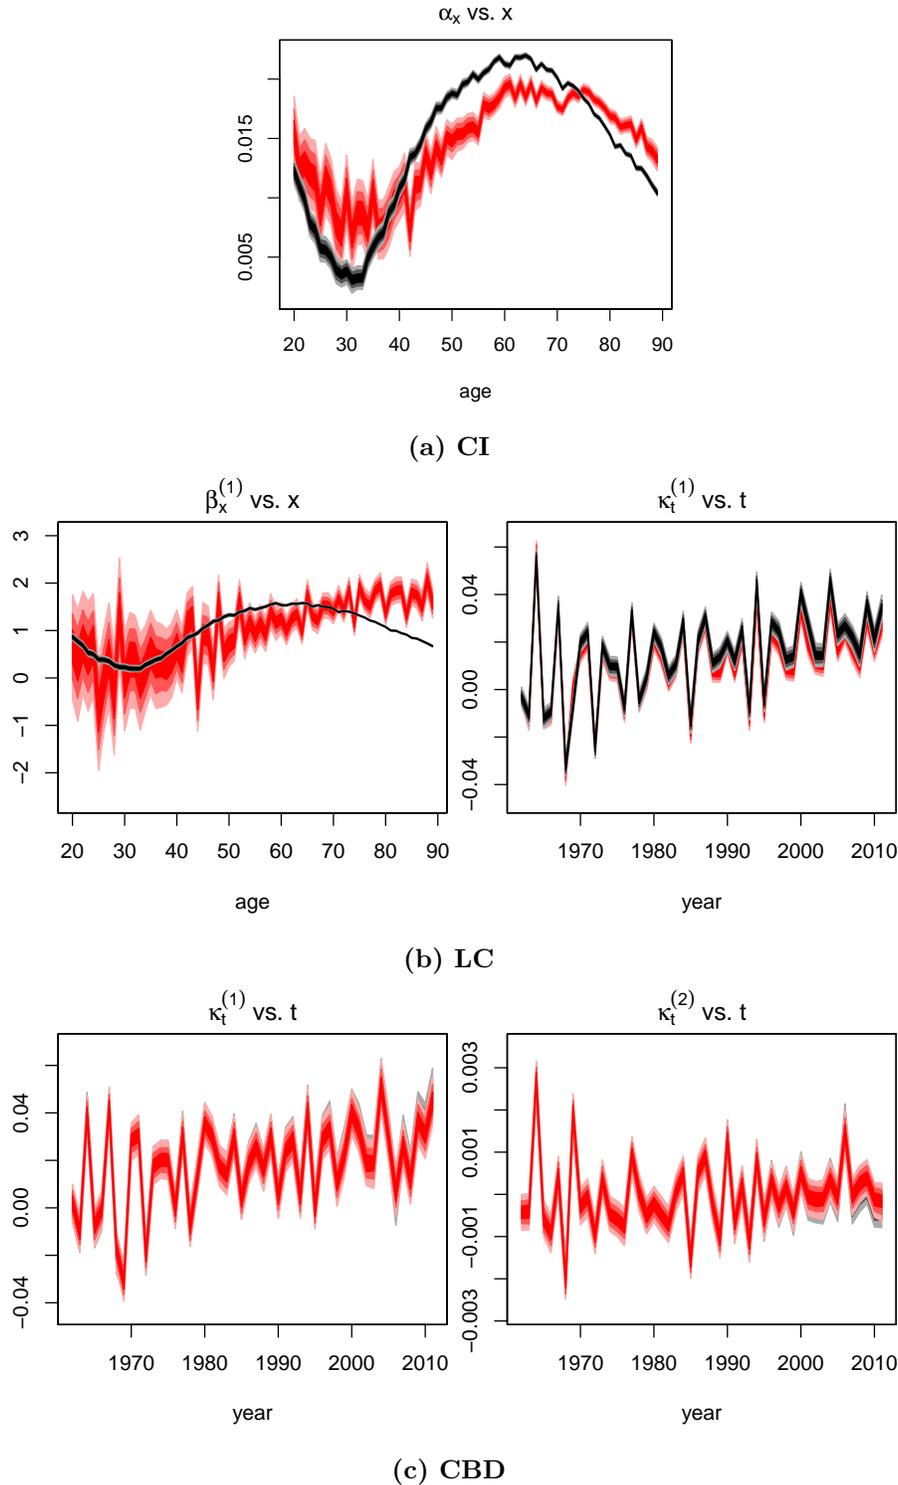


Fig. 9. Parameters for Model CI With Parameter Uncertainty: England and Wales Males, Ages 20–89, 1961–2011.

Note: Shades in the fan represent confidence intervals at the 50%, 80% and 95% level. Black fans correspond to the “fitted” estimation approach, and red fans to the “crude” estimation approach.

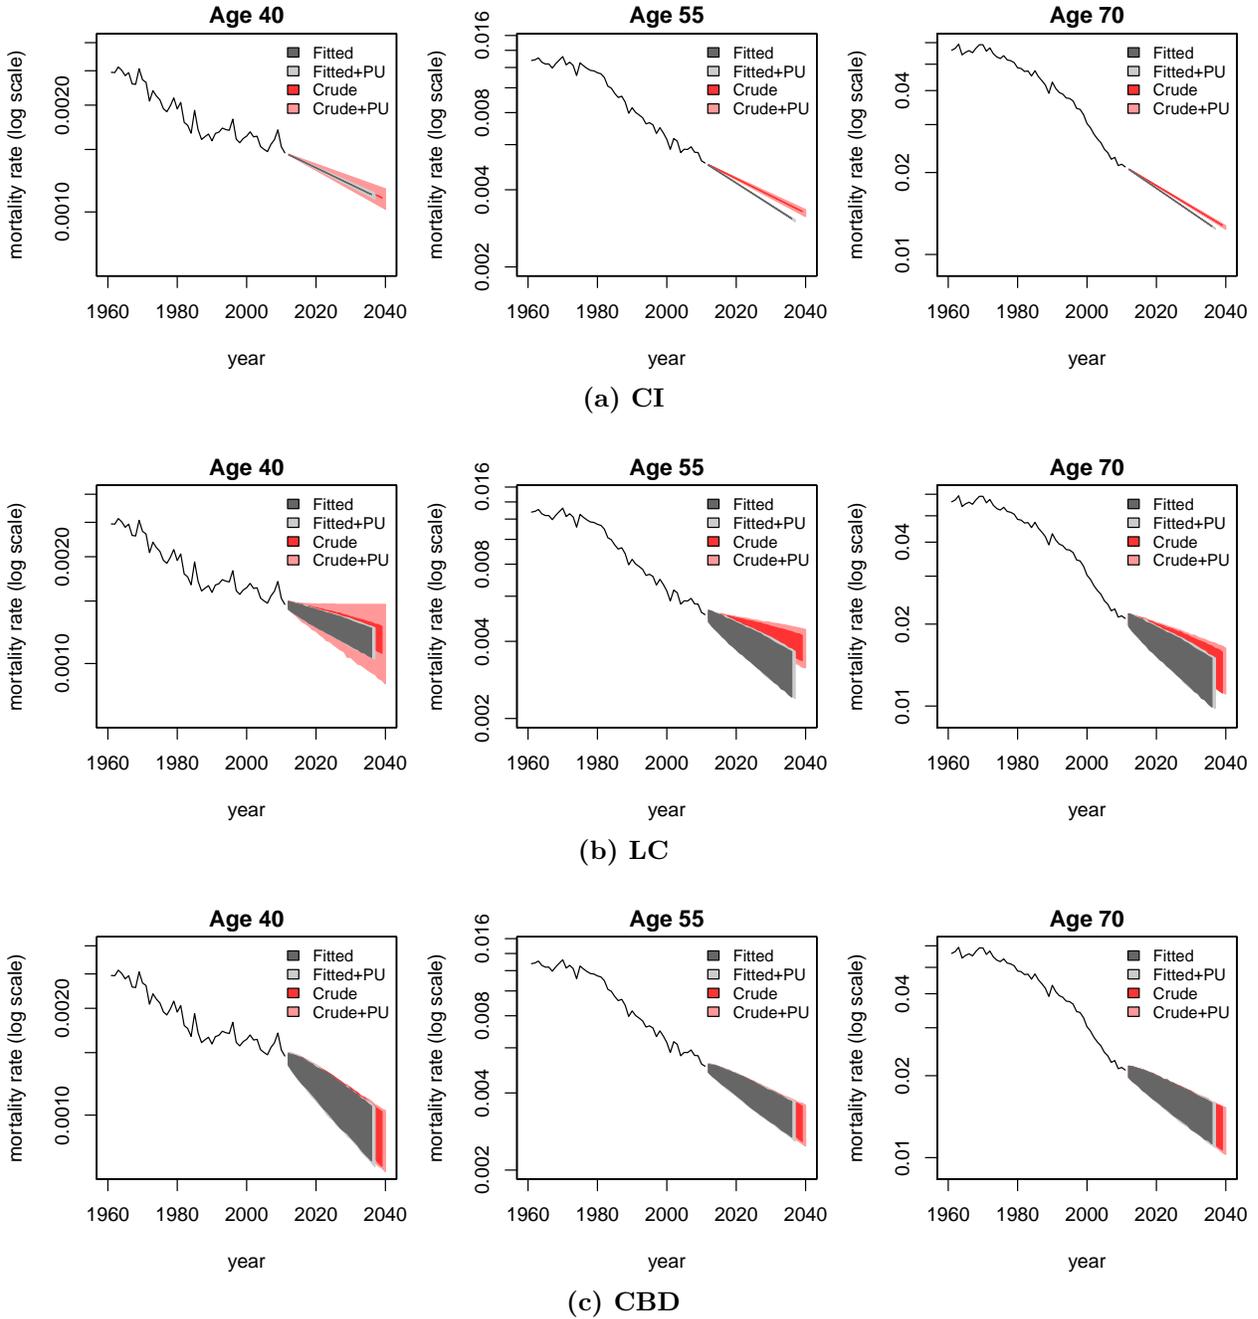


Fig. 10. Fan Charts for Mortality Rates $m_{x,t}$ at Ages $x = 40, 55, 70$ from the CI, LC and CBD Models: England and Wales Males, Ages 20–89, 1961–2011.

Note: The solid lines show historical mortality rates for the period 1961–2011. Shades in the fan represent prediction intervals at the 95% level.

present, for models CI, LC and CBD, 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 “fitted” approach without allowance for parameter uncertainty;
- ii. Forecast produced by the “fitted” approach with allowance for parameter uncertainty;

- iii. Forecast produced by the “crude” approach without allowance for parameter uncertainty; and
- iv. Forecast produced by the “crude” approach with allowance for parameter uncertainty.

From Figure 10, we note the following:

- The noticeable impact of considering parameter uncertainty under the “crude” 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 provides point forecasts only when parameter uncertainty is ignored.
- The noticeable differences between the central projections in the “fitted” and “crude” approaches for the LC model. This is particularly visible at age 55, where the “fitted” approach projects a much steeper decline in mortality than the “crude” 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 3).
- The contrasting similarity in central forecasts and levels of uncertainty for the CBD under both estimation approaches and with or without parameter uncertainty (see Figure 10c).

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 in a more quantitative manner, we consider the relative difference between the median forecasts of mortality rates in year 2040 produced by the “fitted” and “crude” estimation approaches. Formally, if $m_{x,2040}^{f,med}$ and $m_{x,2040}^{c,med}$ denote the median forecasts for age x in 2040 produced by the “fitted” and “crude” estimation approaches, respectively, then the relative difference at age x , RD_x , is given by

$$RD_x = \frac{m_{x,2040}^{f,med} - m_{x,2040}^{c,med}}{m_{x,2040}^{f,med}}.$$

Figure 11 presents plots of these relative difference for all six models of Table 1. We see that, with the exception of the CBD model, for all other models, both estimation approaches result in significantly different central forecasts. The LC model stands out as the one with the highest discrepancies between estimation approaches. It is also interesting to note that the shape of these differences in the median mortality rate forecasts closely matches the shape of the age and cohort parameters in the difference models, indicating that the differing ability of the approaches to estimate these parameters is what drives the forecast differences.

Now, to investigate further the differences in forecast levels of uncertainty produced by the models under both parameter estimation approaches, in Figure 12 we plot, on a log scale, the standard deviation of the mortality rate forecasts at each age in year 2040 produced by the “fitted” and “crude” estimation approaches, with and without parameter uncertainty. From this figure, we note the following:

- The close alignment between the solid black lines and dashed black lines in Figures 12b–12f, indicating that parameter uncertainty has very small impact on the prediction

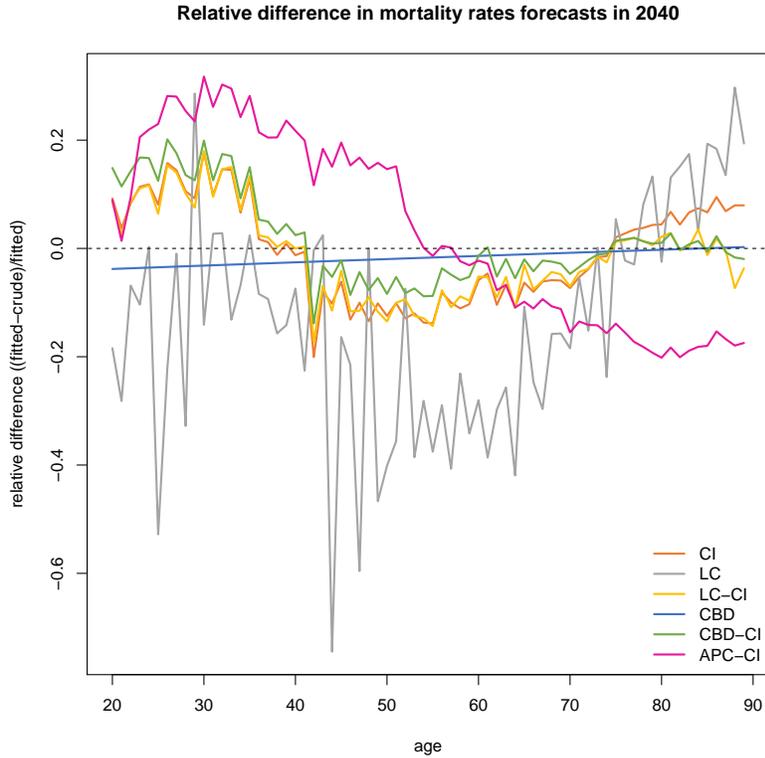


Fig. 11. Relative Difference Between the Mortality Rates Forecast in 2040 for England and Wales Males, Obtained Using a “Fitted” and a “Crude” Estimation Approach

levels of uncertainty under the “fitted” estimation approach for all models.¹¹ We note that in Figure 12a, there are no solid lines, since the CI model provides central forecasts only in the absence of parameter uncertainty.

- The contrasting significant impact of parameter uncertainty under a “crude” approach for models CI, LC, LC-CI and CBD-CI. These models have the characteristics of having a constant improvement rate term, α_x , and/or a nonparametric age modulating parameter $\beta_x^{(1)}$, which are difficult to estimate robustly under the “crude” estimation approach.
- The close similarity of the standard deviations for the CBD model under the four sets of forecasts, highlighting the considerable stability of this predictor structure. This stability is a result of the CBD model being the sole model that only requires estimates of period indexes.

4.4. Robustness and Stability of Projections

The considerable parameter uncertainty seen for some models discussed in the previous

¹¹However, it is worth noting that Figure 12 is on a log scale, so the impact of parameter uncertainty may not be negligible for practical purposes.

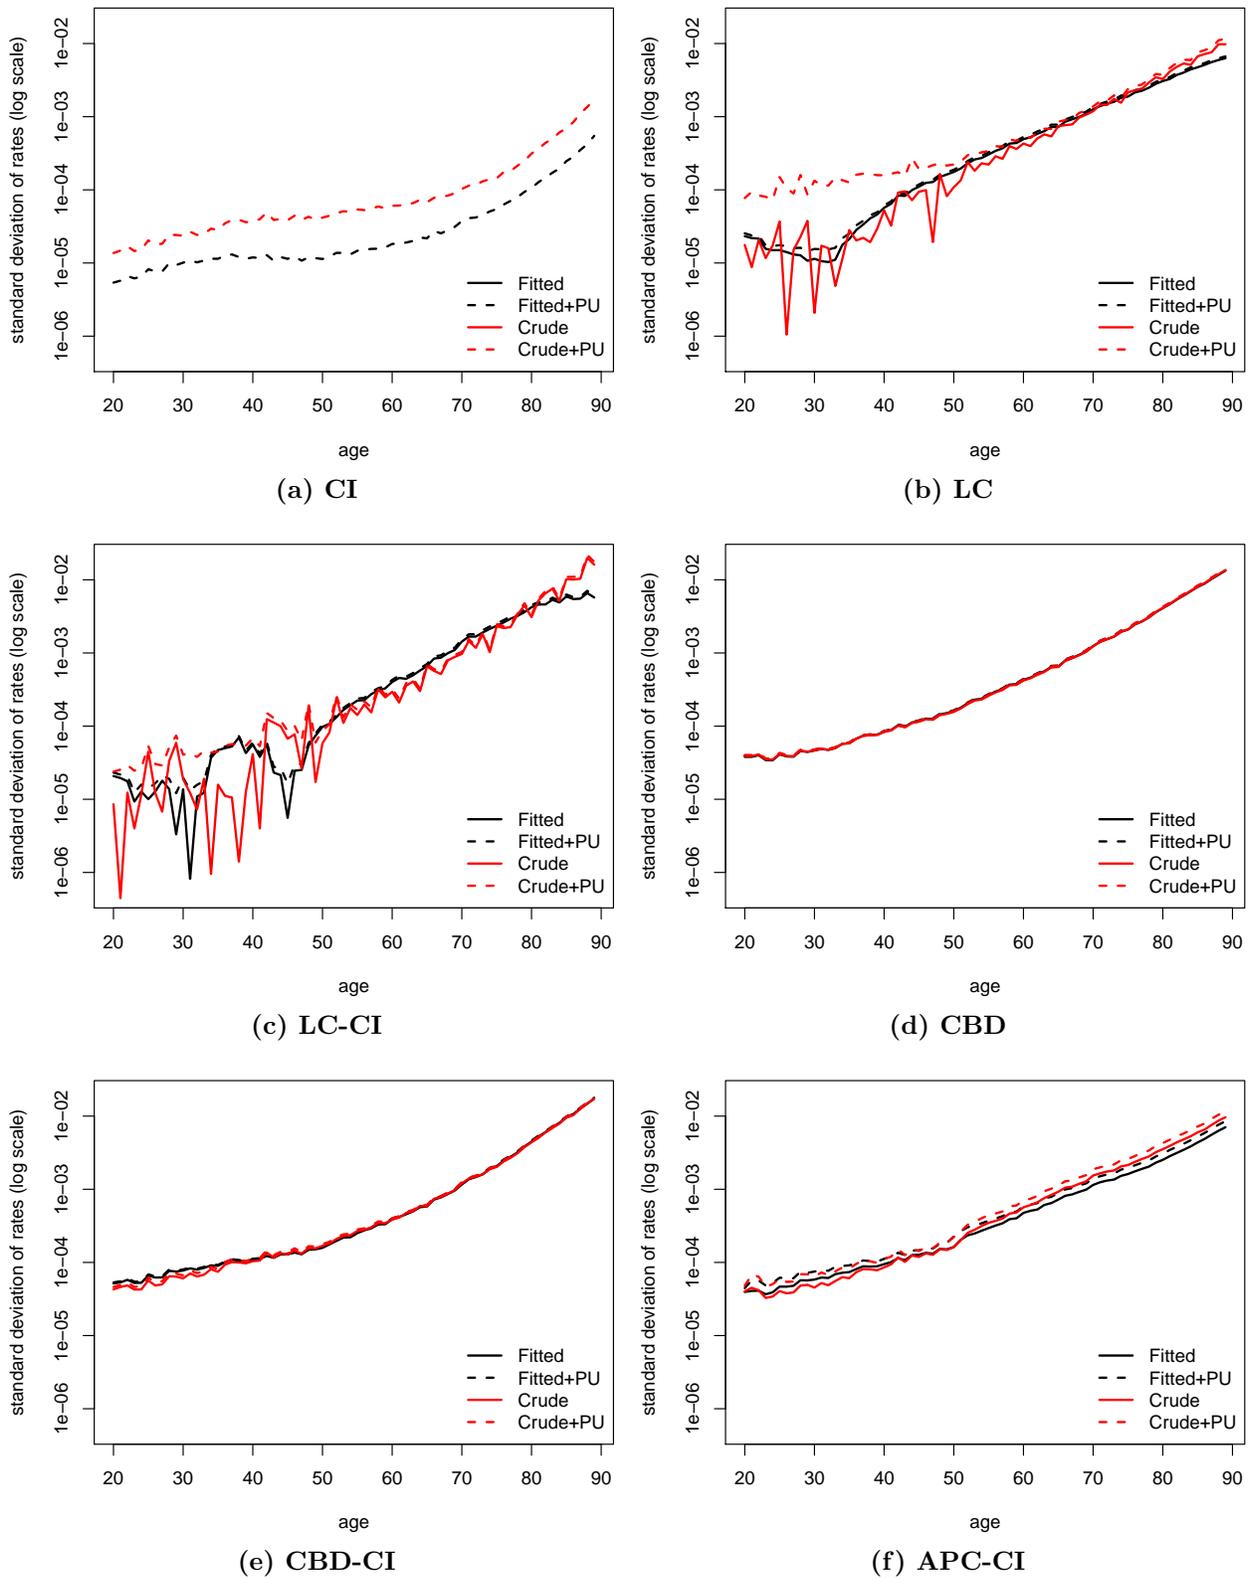


Fig. 12. Standard Deviation of the Mortality Rates Forecast in 2040 for England and Wales Males

section may have important implications 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 13, we show the average 10-year-ahead projected improvement rate at age 40, using different 20-year rolling estimation periods. For instance, the points labeled as stepping-off year 1980 (2011) correspond to each model fitted to data from 1961 to 1980 (1982 to 2011), and the quantity in the vertical axis is the average improvement rate at age 40 for the next 10 years, i.e, for the period 1981–1990 (2012–2021). For a stable model, projections should progress smoothly as we change the data window. From Figure 13, we note the following:

- The general lack of stability of the “crude” estimation approach for all models except for the CBD model. The instability of the “crude” approach for some models is very significant. For instance, in the case of the CI model, changing the estimation period from 1984–2003 to 1985–2004 results in projected 19-year-ahead average improvement passing from 0.19 percent to 1.35 percent (see Figure 13a). The contrasting stability of the CBD approach under both estimation approaches is not a surprise, as the CBD is the only model that does not involve any age improvement terms (see Figure 13d).
- The same pattern of falling 10-year forecasts at age 40 for four out of six models, consistent with the mortality improvement rate reducing at younger ages through the 1990s.
- The noticeable different behavior of the APC-CI model as compared with the rest of the models. This is explained by the fact that this is the only model including a cohort term (see Figure 13f).

5. United States, 1968–2014, Ages 20–89

To test the robustness of the approach to different data sets, we have repeated our analysis using data for men in the United States over the period 1968 to 2014 from the Human Mortality Database (2014). In so doing, we find results that are broadly comparable to those for England and Wales.

However, it is worth noting that the observed parameter uncertainty in the age functions estimated using the “crude” approach is smaller than for England and Wales but still significantly larger than when using the “fitted” approach. This indicates that, although increasing the population size can help with parameter uncertainty under this approach, it does not remove the issue, even for a very large population such as that of the United States.

Figure 14 shows the fitted parameters for the CI, LC and CBD models, with parameter uncertainty, for the U.S. data. It is interesting to see that there is little visible upward trend in $\kappa_t^{(1)}$ for the LC and CBD models, which indicates that, unlike for England and Wales, there has been little acceleration in the rate of improvement in mortality over the period. It is also interesting to note the behavior of improvement rates using the CBD model during the late 1990s, when the average rate of improvement across all ages ($\kappa_t^{(1)}$) increased

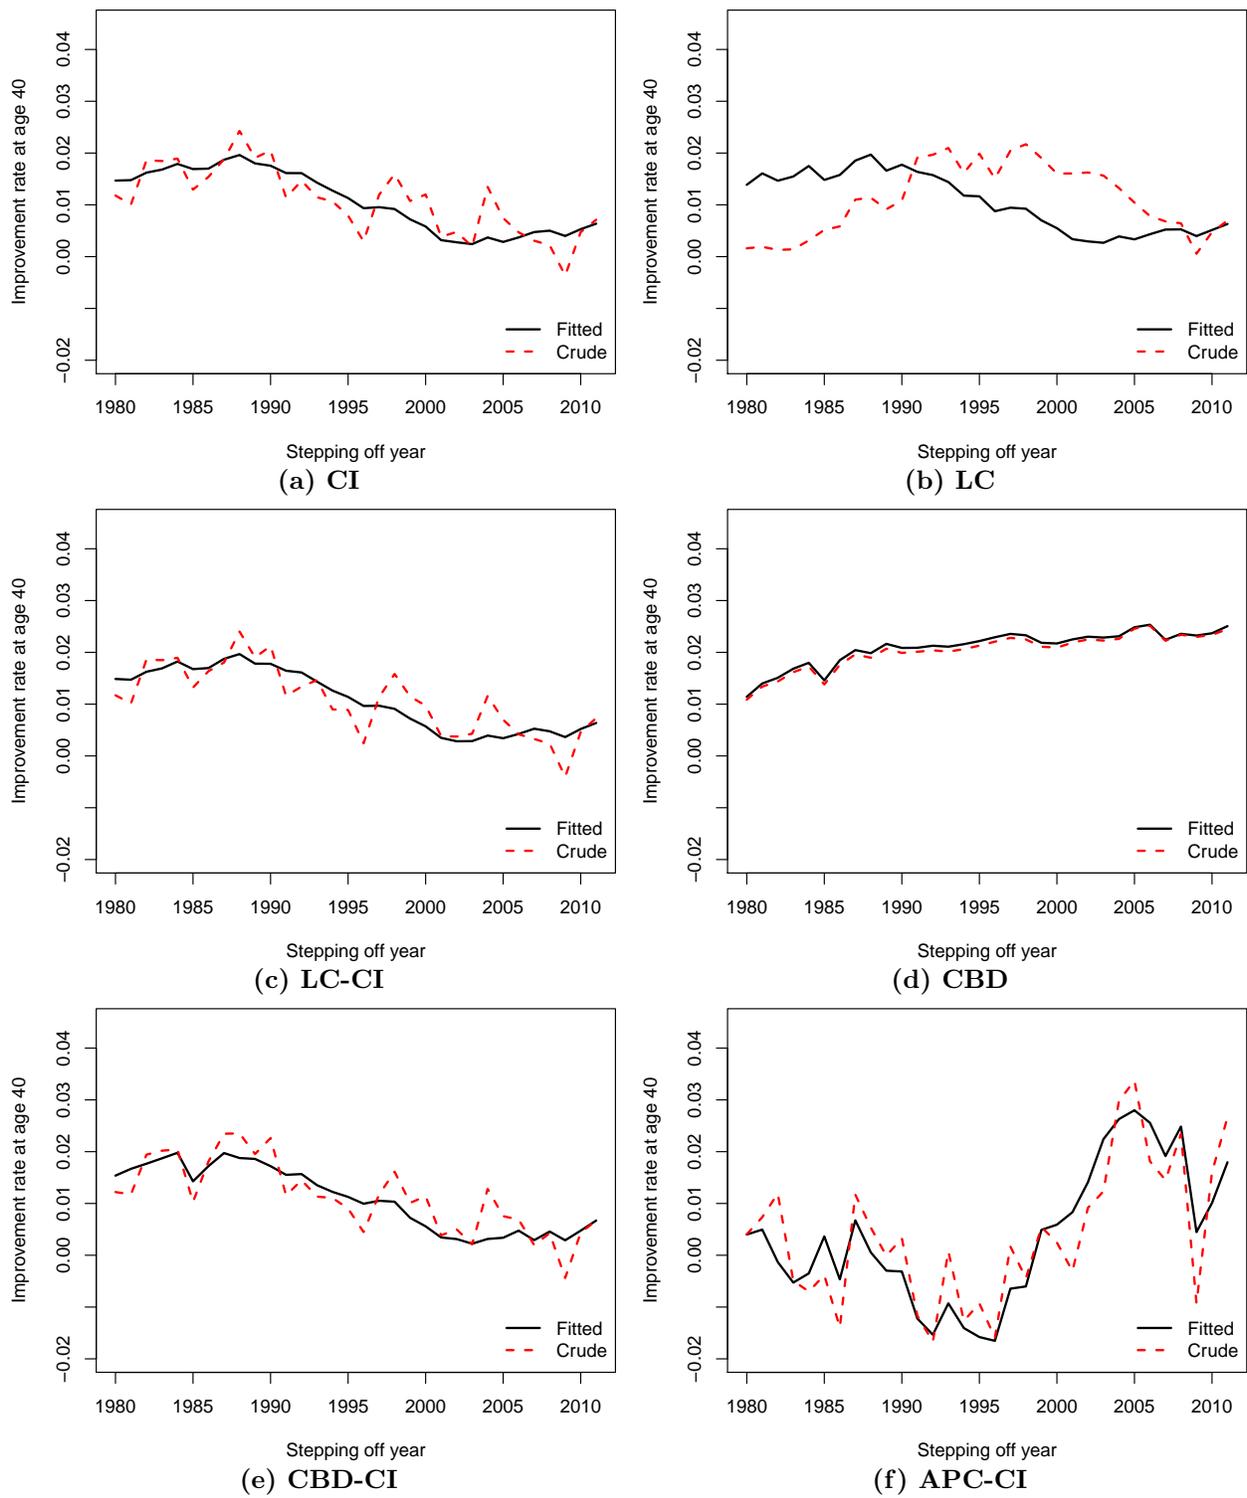


Fig. 13. Average 10-Year-Ahead Projected Improvement Rate at Age 40 With Different Stepping-Off Year: 20-Year Rolling Window

rapidly. However, this was offset by sharp declines in $\kappa_t^{(2)}$, indicating that this faster rate of improvement affected mainly those at younger ages.

Figure 15 shows the projected mortality rates for U.S. men, applying the CI, LC and CBD models using the VAR(1) model. We see that these are all broadly biologically reasonable, although there are similar differences between the “crude” and “fitted” approaches as were discussed for England and Wales in Section 4.3.¹²

6. 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 on 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 previous studies. Furthermore, we investigate the parameter estimates obtained under previous studies and find that they not only are subject to considerable parameter uncertainty but also give significantly different best-estimate forecasts of future mortality rates, which may be biased and less robust in comparison with our approach and with models of mortality rates.

However, we also note that the improvement rate approach makes any changes in the rate of improvement—i.e., systematic accelerations and decelerations in the rate of mortality improvement—more obvious and harder to ignore in the historical data. These were more obvious in the data for England and Wales, but there is no reason to suspect they are not a feature of other data sets that has been largely overlooked in studies that only project log-mortality rates linearly. However, we find that these are difficult to allow for in future projections without running the risk of biologically unreasonable predictions such as mortality crossovers. These trade-offs between the plausibility of future projections and the need for consistency with the trends observed in the recent past are features of many approaches to modeling mortality. However, we leave further examination of this topic for future work.

In summary, we find that the “fitted” 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.

¹²Figures comparable to those shown in Figures 11, 12 and 13 for England and Wales are not shown for reasons of space but are available on request from the authors.

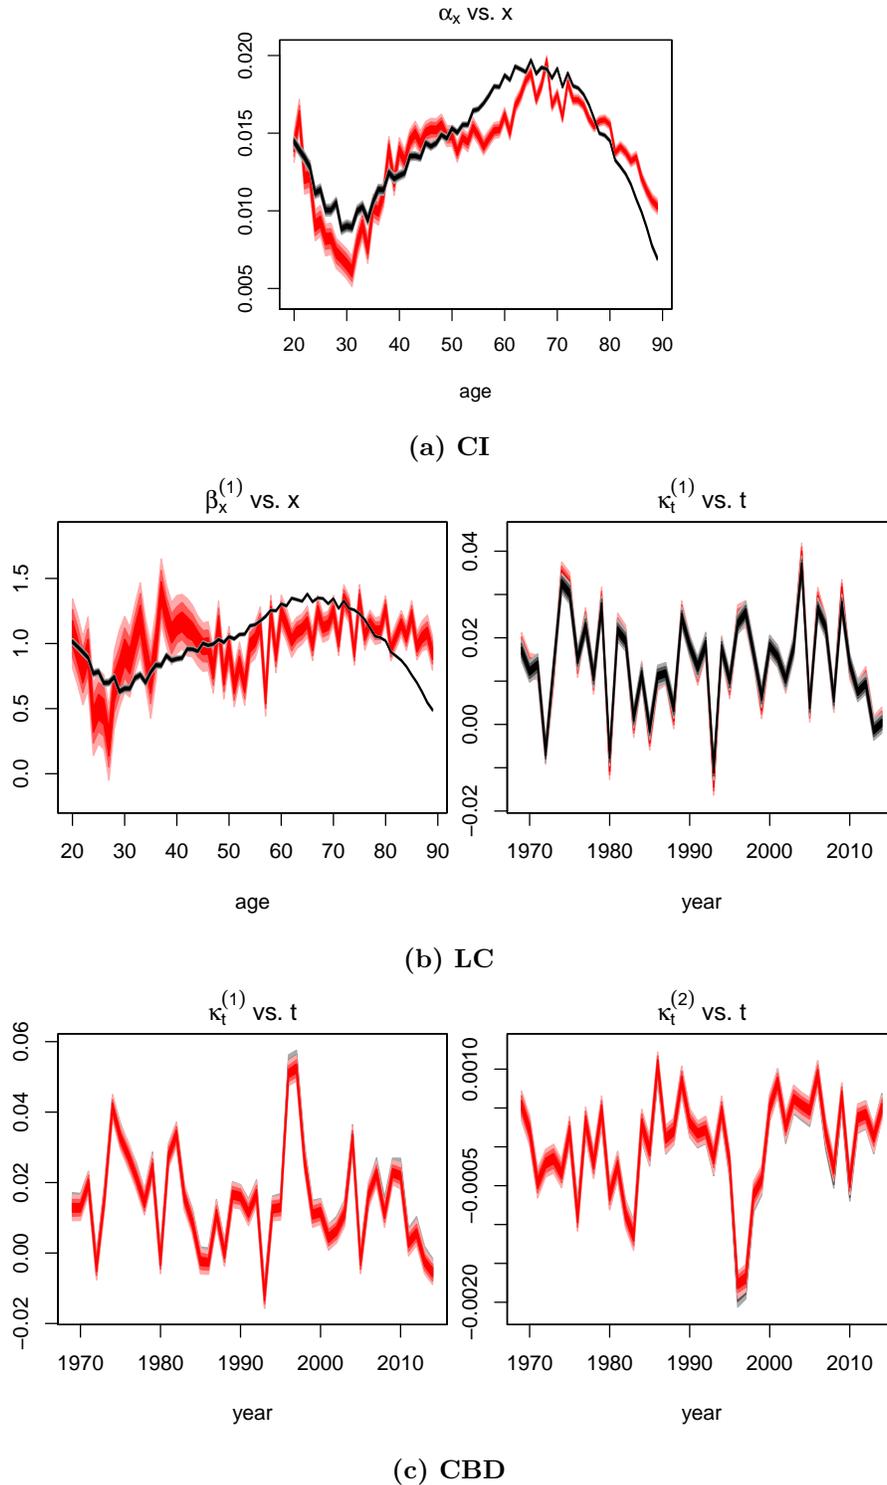


Fig. 14. Parameters for Model CI With Parameter Uncertainty: U.S. Males, Ages 20–89, 1968–2014.

Note: Shades in the fan represent confidence intervals at the 50%, 80% and 95% level. Black fans correspond to the “fitted” estimation approach, and red fans to the “crude” estimation approach.

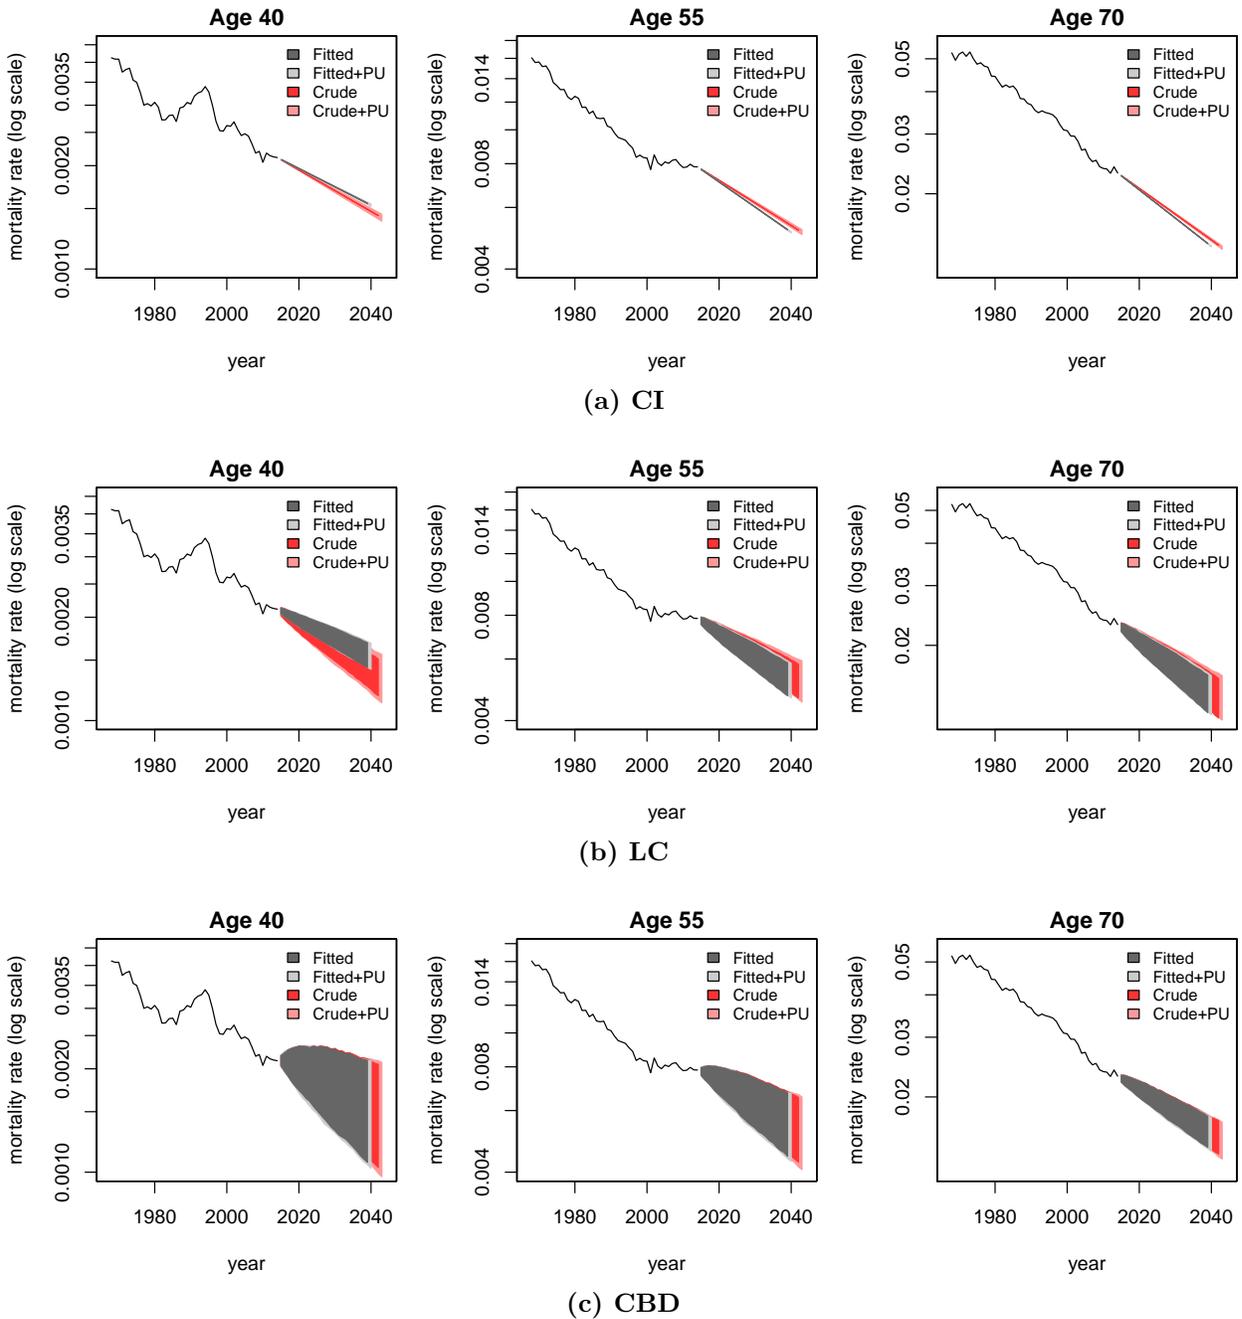


Fig. 15. Fan Charts for Mortality Rates $m_{x,t}$ at Ages $x = 40, 55, 70$ From the CI, LC and CBD Models Applied to U.S. Males, Ages 20–89, 1968–2014.

Note: The solid lines show historical mortality rates for the period 1968–2014. Shades in the fan represent prediction intervals at the 95% level.

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 “crude” approach to the fitting of a mortality improvement rate model, while the constraints in the third column are applied when using the “fitted” 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 sets of constraints make sense intuitively, since the $\alpha_x t$ term in the 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,

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

Table A.2. Parameter Constraints for the Structures Considered in This Paper

Model	Improvement Model (η_{xt})	Equivalent Mortality Model ($\tilde{\eta}_{xt}$)
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	$\sum_t \kappa_t^{(1)} = 0, \sum_y \gamma_y = 0,$ $\sum_y (y - \bar{y}) \gamma_y = 0$	$K_0^{(1)} = 0, \sum_t (t - \bar{t}) K_t^{(1)} = 0,$ $\Gamma_{-X} = 0, \sum_y (y - \bar{y}) \Gamma_y = 0,$ $\sum_y (y - \bar{y})^2 \Gamma_y = 0$

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.

References

- Brouhns, N., Denuit, M. M., Van Keilegom, I., 2005. Bootstrapping the Poisson Log-Bilinear Model for Mortality Forecasting. *Scandinavian Actuarial Journal* (3): 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 (3): 373–393.
- Cairns, A. J., Blake, D., Dowd, K., 2006a. 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., Blake, D., Dowd, K., 2006b. Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk. *ASTIN Bulletin* 36 (1): 79–120.
- Callot, L., Haldrup, N., Lamb, M. K., 2014. Deterministic and Stochastic Trends in the Lee-Carter Mortality Model. Technical Report, Aarhus University.
- Continuous Mortality Investigation, 2002. Working Paper 1 - An interim Basis for Adjusting the “92” Series Mortality Projections for Cohort Effects.
Available at <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper-1>
- Continuous Mortality Investigation, 2009. Working Paper 38 - A Prototype Mortality Projection Model: Part One - an Outline of the Proposed Approach.
Available at <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-papers-38-and-39>
- Continuous Mortality Investigation, 2016a. Working Paper 90 – CMI Mortality Projections Model consultation.
- Continuous Mortality Investigation, 2016b. Working Paper 91 – CMI Mortality Projections Model Consultation - Technical 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 (5): 385–401.
- Haberman, S., Kaishev, V. K., Millossovich, P., Villegas, A. M., Baxter, S. D., Gaches, A. T., Gunnlaugsson, S., Sison, M., 2014. Longevity Basis Risk: A Methodology for Assessing Basis Risk. 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 (3): 309–333.
- Haberman, S., Renshaw, A., 2013. Modelling and Projecting Mortality Improvement Rates Using a Cohort Perspective. *Insurance: Mathematics and Economics* 53 (1): 150–168.
- Human Mortality Database, 2014. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany).
Available at www.mortality.org
- Hunt, A., Blake, D., 2015a. Identifiability in Age/Period Mortality Models. Technical Report, Pensions Institute PI 15-08, Cass Business School, City University of London.
Available at <http://www.pensions-institute.org/workingpapers/wp1508.pdf>
- Hunt, A., Blake, D., 2015b. Identifiability in Age/Period/Cohort Mortality Models. Technical Report, Pensions Institute PI 15-09, Cass Business School, City University of London.
Available at <http://www.pensions-institute.org/workingpapers/wp1509.pdf>
- Hunt, A., Blake, D., 2015c. On the Structure and Classification of Mortality Models. Technical Report, Pensions Institute PI 15-06, Cass Business School, City University of London.
Available at <http://www.pensions-institute.org/workingpapers/wp1506.pdf>
- Koissi, M., Shapiro, A., Hognas, G., 2006. Evaluating and Extending the Lee-Carter Model for Mortality Forecasting: Bootstrap Confidence Interval. *Insurance: Mathematics and Economics* 38 (1): 1–20.
- Lee, R. D., Carter, L. R., 1992. Modeling and Forecasting U.S. Mortality. *Journal of the American Statistical Association* 87 (419): 659–671.
- Li, J. S.-H., Hardy, M. R., Tan, K. S., 2009. Uncertainty in Mortality Forecasting: An Extension to the Classical Lee-Carter Approach. *ASTIN Bulletin* 39 (1): 137–164.
- Mitchell, D., Brockett, P. L., Mendoza-Arriaga, R., Muthuraman, K., 2013. Modeling and Forecasting Mortality Rates. *Insurance: Mathematics and Economics* 52 (2): 275–285.
- Murphy, M., 2009. The “Golden Generations” in Historical Context. *British Actuarial Journal* 15 (S1): 151–184.
- Pfaff, B., 2008. Var, Svar and Svec models: Implementation Within R Package vars. *Journal of Statistical Software* 27 (4).
Available at <http://www.jstatsoft.org/v27/i04/>
- Renshaw, A., Haberman, S., 2003. On the Forecasting of Mortality Reduction Factors. *Insurance: Mathematics and Economics* 32 (3): 379–401.
- 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., Millossovich, P., Kaishev, V. K., 2016. StMoMo: An R Package for Stochastic Mortality Modelling. R package version 0.4.0.
Available at <https://CRAN.R-project.org/package=StMoMo>

Willeits, R., 2004. The Cohort Effect: Insights and Explanations. *British Actuarial Journal* 10 (4): 833–877.