



City Research Online

City St George's, University of London

Citation: Butt, Z. & Haberman, S. (2009). llc: a collection of R functions for fitting a class of Lee-Carter mortality models using iterative fitting algorithms (Actuarial Research Paper No. 190). London, UK: Faculty of Actuarial Science & Insurance, City University London.

This is the unspecified version of the paper.

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

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/2321/>

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



Cass Business School
CITY UNIVERSITY LONDON

**Faculty of Actuarial
Science and Insurance**

**Actuarial Research Paper
No. 190**

**Ilc: A Collection of R Functions for Fitting a
Class of Lee-Carter Mortality Models using
Iterative Fitting Algorithms***

Zoltan Butt and Steve Haberman

September 2009

ISBN 978-1-905752-22-5

Cass Business School
106 Bunhill Row
London EC1Y 8TZ
Tel +44 (0)20 7040 8470
www.cass.city.ac.uk

"Any opinions expressed in this paper are my/our own and not necessarily those of my/our employer or anyone else I/we have discussed them with. You must not copy this paper or quote it without my/our permission".

ilc: A Collection of R Functions for Fitting a Class of Lee-Carter Mortality Models using Iterative Fitting Algorithms*

Version 1.0

Zoltan Butt[‡] and Steven Haberman

April 26, 2010

Abstract

We implement a specialised iterative regression methodology in R for the analysis of age-period mortality data based on a class of generalised Lee-Carter (LC) type modelling structures. The LC-based modelling framework is viewed in the current literature as among the most efficient and transparent methods of modelling and projecting mortality improvements. Thus, we make use of the modelling approach discussed in [Renshaw and Haberman \(2006\)](#), which extends the basic LC model and proposes to make use of a tailored iterative process to generate parameter estimates based on Poisson likelihood. Furthermore, building on this methodology we develop and implement a stratified LC model for the measurement of the additive effect on the log scale of an explanatory factor (other than age and time). This modelling methodology is implemented in a publically available collection of programming functions that facilitate both the preparation of mortality data and the fitting and analysis of the given log-linear modelling structures. Also, the package incorporates methods to produce forecasts of future mortality rates and to compute the corresponding future life expectancy.

Keywords : *generalised/extended Lee-Carter models, age-period-cohort models, iterative estimation approach, statistical programming in R.*

*Financial support for the development of the current statistical program in R from the CMI and from Lucida Plc is gratefully acknowledged. A special thanks goes to our colleague Arthur Renshaw for providing the original source code in GLIM of the iterative fitting technique for the generalised LC models. Also the authors are grateful to the CMI for releasing the mortality data.

[‡]Corresponding author at Sir John Cass Business School, London (UK).

Email: z.butt@city.ac.uk Tel: (+44) (0)20 7040 4104 Fax: (+44) (0)20 7040 8572

1 Introduction

The Iterative Lee-Carter ([ilc](#)) package is a freely available collection of [R](#)¹ (programming) functions for the analysis of age-period mortality data that implements specific regression and descriptive methods to fit a generalised class of LC type modelling structures. The purpose of the mortality modelling package described here is to apply an improved modelling framework, which extends the standard LC method based on the Normal error structure that was originally proposed in [Lee and Carter \(1992\)](#). Consequently, we depart from the traditional Singular Value Decomposition (SVD) fitting method, that assumes Gaussian residuals, and instead implement a regression tool based on Poisson likelihood maximization process. In particular, we make use of the approach proposed and illustrated in [Renshaw and Haberman \(2006\)](#), which generalises the basic LC modelling framework and extends the work of [Brouhns et al. \(2002\)](#), to develop a tailored iterative process for updating the parameter estimates. Furthermore, building on this methodology, we develop and implement a new modelling approach, referred to as the stratified (or extended) LC model, that can be applied to measure the overall effect of an explanatory factor (other than age and time) on the log mortality rates across all ages and periods.

This generalised modelling methodology is implemented within the R statistical software in the form of a specialised set of command functions that apply the above mentioned iterative fitting method. The package contains methods for the analysis of a class of six different types of log-linear models in the GLM framework with Poisson errors that includes the basic LC model too. In addition, the *ilc* package also include tools for the fitting and analysis of the stratified LC model. In order to assess the goodness of fit of the regression, the estimation routines support a range of residual analyses with corresponding target fitted values, which can be visualised by specialised diagnostic plots. The package allows preliminary data corrections, primarily in order to replace missing data-cells, but also to eliminate potential outliers that might result from data inaccuracies. Further, the package includes two simple methods of 'closing-out' procedures to correct the original data at very old ages before the application of the models. Finally, the functionality of this software is currently being enhanced with the inclusion of a number of control parameters and flexible plotting methods.

The remaining sections of this working paper are organised as follows. Section [2](#) presents in detail the variants of the adopted modelling framework and discusses the main features of mortality forecasting within the *ilc* application. Further, section [3](#) provides a brief description of the iterative fitting approach used for the estimation of the model parameters. Following on, section [4](#) gives instructions for installing and using the *ilc* package in R, including how to prepare mortality data and how to fit the models and to run the regression diagnostics. Some numerical illustrations are provided using the CMI pensioners mortality data.

¹A gentle introduction for beginners about methods of statistical analysis and graphical illustration in the R programming environment is provided in [Venables et al. \(2005\)](#).

2 Modelling Framework

The application and extension of the LC modelling approach has dominated the recent literature in the field of mortality forecasting (see [Brouhns et al. 2002](#), [Renshaw and Haberman 2003a,b](#), [Booth 2006](#) and further references therein). According to [Booth \(2006\)](#), the LC-based approach is widely considered in the current literature to be among the most efficient and transparent methods to date that produces fairly realistic life expectancy forecasts, which are used as reference values for other modelling methods. For instance, the accepted framework of modelling and projecting mortality improvements in the USA for the last decade or so has been the LC-based age-period (AP) model (see [Lee and Carter 1992](#), [Lee 2000](#)). Similarly, the model has been applied successfully to Canadian ([Lee and Nault 1993](#)), as well to Japanese (see [Wilmoth 1993](#)) mortality data and formed part of official projections. While the model has gained acceptance in the UK too, the persistent cohort effects observed for generations born between 1925 and 1945 has led to a special adaptation of the LC method by [Renshaw and Haberman \(2006\)](#), developing the so-called age-period-cohort (APC) log-bilinear generalised linear models (GLM) with Poisson error structures.

In terms of forecasting, the LC family of models are part of the extrapolative stochastic methods that assume that the observed historical trends of human mortality improvements will persist into the future. Many authors consider that the relative stability of the past trends provide a sufficiently reliable basis for future projections. While the validity of these assumptions have been debated (see [Guterman and Vanderhoof 2000](#)), the view of the majority is that these methods still offer the most effective and dependable alternative to date. Given the inherent complexity of the factors affecting human mortality and the lack of our understanding of the intricate mechanisms governing our aging process ([Brouhns et al. 2002](#)), econometric or structural models based on causality and interactions of biological and/or demographic factors have so far failed to give rise to plausible theory-informed forecasting methods (see [Booth 2006](#)).

In the LC type modelling approach, the age effects are assumed to be constant in time and the time-variant period and/or cohort effects are projected forward using autoregressive time-series models. Thus, the period and/or cohort factors are extrapolated in time by a stochastic ARIMA process (e.g. random walk with drift) in order to make forecasts of the future force of mortality and, implicitly, future (period- and cohort-based) life expectancy.

In the modelling framework described here, we aim to provide a common platform for fitting LC type models and making future forecasts of mortality and of life expectancy. Thus, we model the force of mortality based on GLM regression methods using a log-link with a class of parameterised predictors that contain bilinear terms. However, the presence of the bilinear predictors prevents the application of the standard estimation methods normally used within the GLM approach. Instead, the parameters are estimated through an iterative minimisation

technique applied to the deviance of the non-linear model structure that is dependent on the choice of error distribution. In the following we describe in more detail the particular modelling structures implemented in the `ilc` package.

2.1 Mortality Data

Consider a mortality experience observed at individual ages (x) and calendar years (t), giving rise to a total of $(k \times n)$ available data cells, so that we can estimate the central mortality rate (m_{xt}) and the corresponding force of mortality (μ_{xt}) by

$$(\hat{\mu}_{xt} =) \hat{m}_{xt} = \frac{y_{xt}}{e_{xt}},$$

where y_{xt} and e_{xt} represent the number of deaths and the matching central exposure for any given subgroup, respectively. In addition, for each combination of age x and period t , we define the cohort year $z = t - x$ representing the year of birth of each subgroup in the data.

2.2 Basic Age-Period (AP) Lee-Carter Model

The basic AP LC model was first proposed by [Lee and Carter \(1992\)](#) and it was introduced as a type of principal components model of the mortality rate (m_{xt}) dependent only on factors related to age and period. The model is expressed as

$$\text{LC : } \log(m_{xt}) = \alpha_x + \beta_x \kappa_t + \varepsilon_{xt}, \quad (1)$$

where the parameters are interpreted as follows:

α_x represents a constant age-specific pattern of mortality;

κ_t measures the trend in mortality over time;

β_x measures the age-specific deviations of mortality change from the overall trend;

ε_{xt} are Gaussian distributed $N(0, \sigma^2)$ random effects by age and time.

Due to the bilinear multiplicative construct ($\beta_x \kappa_t$) present in equation (1), there is a clear identifiability problem that is traditionally resolved by ensuring that these parameters satisfy a pair of specified constraints, given by

$$\sum_x \beta_x = 1, \quad \sum_{t=t_1}^{t_n} \kappa_t = 0. \quad (2)$$

Then, the standard LC model can be estimated using the singular value decomposition (SVD) method that leads to the following estimator of the age-specific effects:

$$\hat{\alpha}_x = \frac{1}{n} \sum_{t=t_1}^{t_n} \log(\hat{m}_{xt}) , \quad (3)$$

which minimises the sum of squares of the error term ($S = \sum_{xt} \varepsilon_{xt}^2$). Lee and Carter also advocates a set of adjustments to the $\hat{\kappa}_t$ estimates in order to ensure that in each year, the total deaths predicted by the model equals the total of the observed deaths $\sum_x y_{xt}$.

Subsequently, the LC model was re-evaluated in the mortality forecasting literature (see [Tableau 2001](#), [Brouhns et al. 2002](#), [Renshaw and Haberman 2003a](#)) and it was proposed that the model can also be formulated within a GLM framework with a generalised error distribution. In this setting, the LC model parameters can be estimated by maximum likelihood (ML) methods based on the choice of error distribution. Thus, in line with traditional actuarial practice, this approach assumes that the age- and period-specific number of deaths are independent realizations from a Poisson distribution with parameters

$$E[Y_{xt}] = e_{xt} \mu_{xt} , \quad \text{Var}[Y_{xt}] = \phi E[Y_{xt}] , \quad (4)$$

where ϕ is a measure of over-dispersion to allow for heterogeneity (e.g. from duplicate policies in the case of insurance data). Making use of the LC type parameterization (1), now in terms of the force of mortality (μ_{xt}), equations (4) correspond to a GLM model of the response variable Y_{xt} with log-link and non-linear parameterized predictor:

$$LC : \quad \eta_{xt} = \log(\hat{y}_{xt}) = \log(e_{xt}) + \alpha_x + \beta_x \kappa_t . \quad (5)$$

In order to obtain unique parameter values, the above model is formulated in line with the same constraints (2), while $\log(e_{xt})$ is treated as an offset value during fitting.²

It is important to emphasise that model (5) is conceptually different from the original LC framework (1), because the modelling errors have a generalised class of distribution that are determined by the direct fitting of the number of deaths instead of the logarithmic transform of the rates. That is, the GLM regression is based on ML methods with theory-based distributional assumptions in contrast to the SVD fitting, which relies on empirical measures (i.e. least squares). Indeed, the parameter estimates under the original framework (1) can also be obtained within the GLM approach by adjusting the target variable to $Y_{xt} = \log(m_{xt})$ and applying the identity link function with a Normal error structure.

A measure of the overall goodness of fit in the GLM settings is the scaled deviance between the observed and the fitted target variable values, which depends on the chosen distributional

²The interpretation and treatment of model (5) in terms of a mortality reduction factor $F(x, t)$ is beyond the scope of the current paper (see [Renshaw and Haberman 2006](#)).

assumption. Thus, ML point estimates under the GLM approach are obtained at the minimum value of the total deviance of model (5) with Poisson errors, which is given by

$$D(y_{xt}, \hat{y}_{xt}) = \sum_{x,t} dev(x, t) = \sum_{x,t} 2\omega_{xt} \left\{ y_{xt} \log \frac{y_{xt}}{\hat{y}_{xt}} - (y_{xt} - \hat{y}_{xt}) \right\}, \quad (6)$$

where $dev(x, t)$ are the deviance residuals that depend on a set of prior weights ω_{xt} where $\omega_{xt} = 1$ is assigned to each non-empty data cell, with $\omega_{xt} = 0$ for empty cells.³ However, standard minimisation techniques cannot be applied due to the presence of the bilinear interaction term ($b_x \kappa_t$). Thus, we resort to an alternative fitting strategy, as described in [Renshaw and Haberman \(2006\)](#), which is based on an iterative Newton-Raphson method applied to the deviance function (6). In section 3, we offer a brief description of the core algorithmic rule that governs the fitting process of this approach, with specific application to the LC model summarized in section 3.1.

Model diagnostics of goodness of fit can be carried out by visual inspection and by formal testing of the following types of residuals, that are listed below in an increasing order of their relevance in the current modelling framework:

- a) **log-rates:** $r_{xt} = \log(\mu_{xt}) - \log \hat{\mu}_{xt}$;
- b) **rates:** $r_{xt} = (\mu_{xt}) - (\hat{\mu}_{xt})$;
- c) **deaths:** $r_{xt} = y_{xt} - \hat{y}_{xt} = e_{xt} \mu_{xt} - e_{xt} \hat{\mu}_{xt}$;
- d) **deviance:** $r_{xt} = \text{sign}(y_{xt} - \hat{y}_{xt}) \sqrt{\frac{dev_{xt}}{\hat{\phi}}}$, $\hat{\phi} = \frac{D(y_{xt}, \hat{y}_{xt})}{\nu}$,

where $\hat{\phi}$ is an empirical scaling factor and ν represents the degrees of freedom, dependent on the particular model structure.

Thus, in the *ilc* package, we make available plotting methods that can produce residual plots of the above residuals with respect to age, period and year of birth. The latter can be used also to check for cohort effects in case these are not directly measured in the model. As an additional model diagnostic, the program can also make plots of the fitted values (i.e. either mortality rates or number of deaths) against age and period.

2.3 Generalized Family of Lee-Carter Models

In a more recent development, the basic setting has been further extended to include an additional bilinear term, containing a second period effect (as in [Renshaw and Haberman 2003b](#)) or

³In contrast to the GLM approach, in the SVD fitting the application of data matrix containing empty cells is not possible. Nonetheless, as mentioned before, the *ilc* program can optionally correct missing data cells by 'closing-out' methods in order to improve fitting.

a cohort effect (as in [Renshaw and Haberman 2006](#)). In particular, the latter approach sheds new light on the early 20th century England and Wales mortality patterns. Thus, the basic LC model can be transformed into a more general framework in order to analyse the relationship between age and time and their joint impact on the mortality rates. In the current application, we follow the APC modelling framework and fitting methodology proposed by [Renshaw and Haberman \(2006\)](#) that specifies the force of mortality by a generalized structure written as

$$M : \quad \mu_{xt} = \exp \left(\alpha_x + \beta_x^{(0)} \iota_{t-x} + \beta_x^{(1)} \kappa_t \right) , \quad (7)$$

where α_x maps the main age profile of mortality, ι_{t-x} and κ_t represent the cohort and period effects, respectively, whereas $\beta_x^{(0)}$ and $\beta_x^{(1)}$ parameters measure the corresponding interactions with age.

We note that model (7) represents a family of six generalized non-linear models of the LC type structure with log-link function. The sub-categories of the overall model can be defined by independently setting the interaction parameters $\left(\beta_x^{(0,1)} \right)$ to one of the following:

- a) unknown (to be estimated);
- b) =1 (fixed);
- c) =0 (void).

Thus, the basic LC type structure results by defining the age-specific parameters as

$$LC : \quad \beta_x^{(0)} = 0 \ (\forall x) \quad \text{and} \quad \beta_x^{(1)} = \beta_x .$$

Alternative formulation can result by cancelling out the period effect altogether and maintaining only the age and the cohort effects, as follows:

$$AC : \quad \beta_x^{(0)} = \beta_x \quad \text{and} \quad \beta_x^{(1)} = 0 \ (\forall x) .$$

Following the same approach, other 3 substructures can be defined, namely (using the notations introduced by [Renshaw and Haberman 2006](#)):

$$H_0 : \beta_x^{(0,1)} = 1 ; \quad H_1 : \beta_x^{(0)} = 1 ; \quad H_2 : \beta_x^{(1)} = 1 .$$

We note that the main regression function of the `ilc` package implements all six substructures of model (7) making use of either the Gaussian or the Poisson error distribution. The overall estimation of this class of model structures proceeds along the same iterative minimisation techniques, which are described in section 3. However, in order to obtain unique parameter estimates, we need to make slight modifications to the parameter updating cycle depending on the particulars of the sub-structure. Given its overwhelming importance, we illustrate the algorithmic rule of the most general APC framework (i.e. model *M*) in part 3.2.

2.4 Stratified (or Extended) Lee-Carter Model

The purpose of the methodology described here is to quantify the differences in the mortality experience of population subgroups differentiated by an additional measurable covariate (other than age and period). This is a new modelling approach that assumes a direct additive effect of an observable factor on the log mortality rates across all ages and calendar time periods. Clearly, the usefulness of an all-encompassing additional factor strongly depends on the size and nature of the mortality experience. Examples where additional effects might exist that could act constantly across age and time in human mortality experience include factors related to geographical, socio-economic or race differences. The modelling framework and estimation methodology proposed here builds on the previous LC type structure with Poisson errors presented in the previous section.

Consider a cross-classified mortality experience observed over age (x), period (t) and an extra variate (g), made up of $(k \times n \times l)$ data cells, such that we can estimate the central mortality rates (m_{xtg}) and the force of mortality (μ_{xtg}) for any given subgroup by the ratio of the number of deaths and the corresponding central exposure (see section 2.1).

As in the previous approaches, our aim is to model the number of deaths (y_{xtg}) within a generalized LC framework with a Poisson error structure, shaped by the following parameterized (non-linear) predictor:

$$SLC : \quad \eta_{xtg} = \log(\hat{y}_{xtg}) = \log(e_{xtg}) + \alpha_x + \alpha_g + \beta_x \kappa_t, \quad (8)$$

where $\log(e_{xtg})$ is treated as an offset value during fitting and the model parameters are subject to the usual constraints defined in equations (2).

We note that relationship (8) can be viewed as an adjusted LC model, whereas the overall trend of mortality change (κ_t) over time and its interaction (β_x) with age is the same for the entire population, while the main effect is now *stratified* in order to capture both the effect of age and an additional variate (g), namely:

$$\hat{\mu}_{xtg} = \exp(\alpha_{xg} + \beta_x \kappa_t),$$

where $\alpha_{xg} = \alpha_x + \alpha_g$. We note that, in this formulation, the parameter α_g measures the relative differences between the age-specific mortality profiles on the log scale of the population subgroups defined by the extra variate (g). It is interesting to observe that this modelling structure corresponds to the “common factor” model of Li and Lee (2005). The estimation method of this modelling framework is presented in more detail in section 3.3.

This is the simplest extension of the LC model to allow for stratification. More complex models involving β_{xg} and κ_{tg} could also be introduced — these are left for future development.

2.5 Forecasting Approach

The forecasting of mortality rates in the case of the LC family of models (7) is based on time series prediction of the calendar time dependent parameters (ι_{t-x} , κ_t). This can be written as follows:

$$\dot{\mu}_{x,t_n+s} = \exp\left(\hat{\alpha}_x + \hat{\beta}_x^{(0)} \dot{i}_{t_n+s-x} + \hat{\beta}_x^{(1)} \dot{\kappa}_{t_n+s}\right), \quad s > 0, \quad (9)$$

where $\dot{i}_{t_n+s-x}$ and $\dot{\kappa}_{t_n+s}$ represent the forecasted cohort and period effects, respectively. Observe that, in the case of cohort effects, the forecasted values revert to the fitted parameters (*i.e.* $\dot{i}_{t_n+s-x} = \hat{i}_{t_n+s-x}$) whenever the forecasting horizon falls within the available data range (*i.e.* $\forall s \leq x - x_1$). This forecasting method allows us to generate future average values and to evaluate the future variability of the central mortality rates. In turn, the variability of the predictions can be applied to measure the uncertainty in the longevity risk.⁴

The most common type of time series extrapolation methods applied in the LC framework are the univariate ARIMA (Auto-Regressive Integrated Moving Average) processes, which are characterised by three parameters (p, d, q). The type of ARIMA model used depends on the fitted parameter profile within the available data range (e.g. the size of deviations from the mean, extent of stationarity etc.). In the majority of applications of the LC framework the random walk with drift (0,1,0) is the usual choice for the period effects (κ_t), which can be expressed as:

$$\kappa_t = \kappa_{t-1} + d + e_t, \quad (10)$$

where d measures the drift in the form of average annual deviations and e_t represents the white noise in the stochastic process.

According to Booth et al. (2006), ARIMA(0,1,0) is a reasonable choice in the cases where there is a stable linear tendency in the annual mortality improvements, but would be inappropriate for the cases characterised by regular dynamic changes in slope (*i.e.* non-linear). Nevertheless, the authors have found that this model has performed well in many large data applications, even when a more complex model might have been indicated by the shape of the period effects. Similarly, on inspection of the output results of our own empirical trials, we are satisfied that this method is appropriate for many human mortality data sets. In the current version of the *ilc* package, there are methods only for the time dependent parameter to be projected forward, although with slight adjustments it is possible to extrapolate (indirectly) the cohort dependent parameter values too. In future versions, we plan to implement complete and automated forecasting methods using a wider range of ARIMA models for all six modelling structures considered in this application. Note there are several choices for the forecast formula

⁴We recognise that a method based purely on the extrapolated time dependent parameters might fail to capture all of the variability in future predicted values because it does not allow for the uncertainty in the other model parameters. However, as noted by Lee and Carter (1992), this simplified approach should still provide a good approximation for the calculation of the prediction intervals. This has recently been explored in extensive bootstrapping investigations, as evidenced, for example by Renshaw and Haberman (2008, 2009).

– thus, equation (9) uses the model μ_{x,t_n} as the jump off value for forecasting. We could also use last observed data point $\hat{\mu}_{x,t_n}$ (see Lee 2000) or an average value.

3 Fitting Methodology

As mentioned before, the fitting methodology implemented in this application is based on an iterative algorithm that minimises the deviance function. That is, we make use of a cyclical updating process of the parameter estimates until the minimum difference between the likelihood of the fitted model and the likelihood of the saturated model (i.e. one parameter for each observation) is achieved. Thus, the updating mechanism for a given parameter θ is provided by the Newton-Raphson minimisation method applied to the deviance function, which can be expressed as

$$u(\hat{\theta}) = \hat{\theta} - \frac{\frac{\partial D}{\partial \theta}}{\frac{\partial^2 D}{\partial \theta^2}}. \quad (11)$$

Looking at the deviance function (6) with Poisson error structure, we can observe that

$$\begin{aligned} \frac{\partial D}{\partial \theta} &= \sum \frac{\partial dev}{\partial \theta} = \sum 2\omega \left\{ -y \frac{\hat{y}'}{\hat{y}} + \hat{y}' \right\} \\ &= \sum 2\omega \frac{\hat{y}'}{\hat{y}} (\hat{y} - y) = \sum 2\omega a (\hat{y} - y), \end{aligned} \quad (12)$$

where

$$\hat{y}' = \frac{\partial \hat{y}}{\partial \theta} \Rightarrow \begin{cases} \frac{\partial \hat{y}}{\partial \alpha_x} = \hat{y} \\ \frac{\partial \hat{y}}{\partial b_x} = \kappa_t \hat{y} \\ \frac{\partial \hat{y}}{\partial \kappa_t} = b_x \hat{y} \end{cases} = a \hat{y} \quad \text{such that} \quad \begin{cases} a = 1 \\ a = \kappa_t \\ a = b_x \end{cases}.$$

Making use of the above simplified notations, we can express the second partial derivative of the deviance function as follows:

$$\frac{\partial^2 D}{\partial \theta^2} = \sum 2\omega a \hat{y}' = \sum 2\omega a^2 \hat{y}. \quad (13)$$

Substituting the expressions (12) and (13) into (11) yields the following general fitting routine:

$$u(\hat{\theta}) = \hat{\theta} - \frac{\sum 2\omega a (\hat{y} - y)}{\sum 2\omega a^2 \hat{y}} = \hat{\theta} + \frac{\sum 2\omega a (y - \hat{y})}{\sum 2\omega a^2 \hat{y}}. \quad (14)$$

We note that similar updating rule can be determined in the case of the model with Gaussian distributed errors (see Renshaw and Haberman 2006). Without going into further details, we note that the *ilc* package implements the updating algorithms corresponding to the models with both Gaussian and Poisson error structures. For the purpose of the current paper, in the

following parts we focus on the detailed estimation methodology of the latter with respect to the base LC, the APC and the SLC modelling frameworks.

3.1 Updating cycle of the base LC fitting

1. Get appropriate initial values:

$$\hat{\alpha}_x = \frac{1}{n} \sum_t \log \hat{m}_{xt} \text{ (i.e. make use of the SVD estimate (3));}$$

$$\hat{b}_x = \frac{1}{k}; \quad \hat{\kappa}_t = 0.$$

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x, \hat{\kappa}_t) \rightarrow$ calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

2. Update parameter $\hat{\alpha}_x$:

$$\hat{\alpha}_x = \hat{\alpha}_x + \frac{\sum_t 2\omega(y - \hat{y})}{\sum_t 2\omega \hat{y}}$$

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x, \hat{\kappa}_t) \rightarrow$ calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

3. Update parameter $\hat{\kappa}_t$:

$$\hat{\kappa}_t = \hat{\kappa}_t + \frac{\sum_x 2\omega(y - \hat{y})}{\sum_x 2\omega \hat{b}_x^2 \hat{y}}$$

– adjust the updated parameter such that $\hat{\kappa}_t = \hat{\kappa}_t - \overline{\hat{\kappa}_t}$;

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x, \hat{\kappa}_t) \rightarrow$ calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

4. Update parameter $\hat{b}_x$:

$$\hat{b}_x = \hat{b}_x + \frac{\sum_t 2\omega(y - \hat{y})}{\sum_t 2\omega \hat{\kappa}_t^2 \hat{y}}$$

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x, \hat{\kappa}_t) \rightarrow$ calculate deviance $D_u(y_{xt}, \hat{y}_{xt})$.

5. Check deviance convergence:

$$\Delta D = D - D_u$$

were D_u is the updated deviance at step 4.

– if $\Delta D > 1 \times 10^{-6} \Rightarrow$ goto step 2.

– Stop iterative process once $\Delta D \approx 0$ and take the fitted parameters as the ML estimates to the observed data.

– Alternatively, stop if $\Delta D < 0$ for a consecutive 5 updating cycles and consider using other starting values or declare the iterations non-convergent.

6. Once convergence is achieved, re-scale the interaction parameters: $\hat{b}_x$ and $\hat{\kappa}_t$:

$$\hat{b}_x = \frac{\hat{b}_x}{\sum_x \hat{b}_x}; \quad \hat{\kappa}_t = \hat{\kappa}_t \times \left(\sum_x \hat{b}_x \right),$$

in order to satisfy the usual LC model constraints $\sum_t \kappa_t = 0$ and $\sum_x b_x = 1$.

3.2 Updating cycle of APC fitting

In the full age-period-cohort GLM model (7), the sum $\log(e_{xt}) + \alpha_x$ is treated as an offset value. Consequently, the α_x parameter is not adjusted during the iterative process when both the year and the cohort effects are included in the model structure.

1. Estimate the (fix) age effects:

$$\hat{\alpha}_x = \frac{1}{n} \sum_t \log \hat{m}_{xt} \text{ (i.e. make use of the SVD estimate (3));}$$

2. Get appropriate initial values:

$$\hat{b}_x^{(0)} = \hat{b}_x^{(1)} = \frac{1}{k};$$

Estimate the simplified period-cohort predictor (i.e. model H_0 , see section 2.3):

$$\eta_{xt} = (\log(e_{xt}) + \alpha_x) + \iota_z + \kappa_t;$$

in order to get initial values for ι_z and κ_t .

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x^{(0)}, b_x^{(1)}, \hat{\iota}_z, \hat{\kappa}_t) \rightarrow$ calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

3. Update parameter $\hat{\iota}_z$:

$$\hat{\iota}_z = \hat{\iota}_z + \frac{\sum_x 2\omega(y - \hat{y})}{\sum_x 2\omega(\hat{b}_x^{(0)})^2 \hat{y}}$$

– shift the updated parameter such that $\hat{\iota}_z = \hat{\iota}_z - \hat{\iota}_1$;

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x^{(0)}, b_x^{(1)}, \hat{\iota}_z, \hat{\kappa}_t) \rightarrow$ calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

4. Update parameter $\hat{b}_x^{(0)}$:

$$\hat{b}_x^{(0)} = \hat{b}_x^{(0)} + \frac{\sum_t 2\omega(y - \hat{y})}{\sum_t 2\omega \hat{\iota}_z^2 \hat{y}}$$

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x^{(0)}, b_x^{(1)}, \hat{\iota}_z, \hat{\kappa}_t) \rightarrow$ calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

5. Update parameter $\hat{\kappa}_t$:

$$\hat{\kappa}_t = \hat{\kappa}_t + \frac{\sum_x 2\omega(y - \hat{y})}{\sum_x 2\omega(\hat{b}_x^{(1)})^2 \hat{y}}$$

- shift the updated parameter such that $\hat{\kappa}_t = \hat{\kappa}_t - \hat{\kappa}_1$;
- calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x^{(0)}, b_x^{(1)}, \hat{\iota}_z, \hat{\kappa}_t)$ → calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

6. Update parameter $\hat{b}_x^{(1)}$:

$$\hat{b}_x^{(1)} = \hat{b}_x^{(0)} + \frac{\sum_t 2\omega(y - \hat{y})}{\sum_t 2\omega \hat{\kappa}_t^2 \hat{y}}$$

- calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{b}_x^{(0)}, b_x^{(1)}, \hat{\iota}_z, \hat{\kappa}_t)$ → calculate deviance $D(y_{xt}, \hat{y}_{xt})$.

7. Check deviance convergence:

$$\Delta D = D - D_u$$

were D_u is the updated deviance at step 6.

- if $\Delta D > 1 \times 10^{-6} \Rightarrow$ goto step 3.
- Stop iterative process once $\Delta D \approx 0$ and take the fitted parameters as the ML estimates to the observed data.
- Alternatively, stop if $\Delta D < 0$ for a consecutive 5 updating cycles and consider using other starting values or declare the iterations non-convergent.

8. Once convergence is achieved, re-scale the interaction parameters: $\hat{b}_x^{(0)}$, $\hat{b}_x^{(1)}$, $\hat{\iota}_z$ and $\hat{\kappa}_t$:

$$\hat{b}_x^{(0)} = \frac{\hat{b}_x^{(0)}}{\sum_x \hat{b}_x^{(0)}} , \quad \hat{b}_x^{(1)} = \frac{\hat{b}_x^{(1)}}{\sum_x \hat{b}_x^{(1)}} ; \quad \hat{\kappa}_t = \hat{\kappa}_t \times \left(\sum_x \hat{b}_x^{(1)} \right) ,$$

in order to satisfy the APC model constraints $\sum_x b_x^{(0)} = \sum_x b_x^{(1)} = 1$ and $\sum_t \kappa_t = 0$.

3.3 Updating cycle of SLC fitting

Due to the stratified nature of the main effect variable (α_{xg}) and the target Poisson error structure, the parameters of model (8) cannot be fitted by the SVD method used in the traditional LC approach. Therefore, in order to estimate the above SLC model (8) we make use of the iterative methodology given in section 3 by making a few necessary adjustments to allow for the extra explanatory variable (α_g). Thus, the extended deviance function of model (8) with Poisson errors is given by the sum of the deviance residuals in all of the available data cells, and this can be written as:

$$D(y_{xtg}, \hat{y}_{xtg}) = \sum_{x,t,g} dev(x, t, g) = \sum 2\omega \left\{ y \log \frac{y}{\hat{y}} - (y - \hat{y}) \right\} , \quad (15)$$

where in the last sum notation we drop the subscripts for the sake of simplicity.

Then, we make the corresponding adjustments regarding the extra dimension in the model, so that the Newton-Raphson minimising routine of the adjusted deviance function (15) can proceed along similar lines to those described earlier. Thus, we can make use of equations (12) and (13) in order to find the first and the second order differentials, as follows:

$$\frac{\partial \hat{y}}{\partial \alpha_g} = \hat{y} \quad (= a \hat{y}) .$$

Hence, one needs to substitute $a = 1$ value in the updating rule (14) corresponding to parameter α_g . The iterative calculations need to take into account the higher dimension in the cross-classified data by age, period and factor g . In the following, we demonstrate the adjusted updating cycle that allows for this extra dimension in the observed mortality experience.

1. Get appropriate initial values:

$$\begin{aligned} \hat{\alpha}_x &= \frac{1}{n \times l} \sum_{t,g} \log \hat{m}_{xtg} \text{ (i.e. the average logrates across all } t, g \text{ indexed cells);} \\ \hat{\alpha}_g &= 0; \quad \hat{b}_x = \frac{1}{k}; \quad \hat{\kappa}_t = 0 . \\ &\rightarrow \text{ calculate fitted values } \hat{y}(\hat{\alpha}_x, \hat{\alpha}_g, \hat{b}_x, \hat{\kappa}_t) \rightarrow \text{ calculate deviance } D(y_{xtg}, \hat{y}_{xtg}) . \end{aligned}$$

2. Update parameter $\hat{\alpha}_x$:

$$\begin{aligned} \hat{\alpha}_x &= \hat{\alpha}_x + \frac{\sum_{t,g} 2\omega(y - \hat{y})}{\sum_{t,g} 2\omega \hat{y}} \\ &\rightarrow \text{ calculate fitted values } \hat{y}(\hat{\alpha}_x, \hat{\alpha}_g, \hat{b}_x, \hat{\kappa}_t) \rightarrow \text{ calculate deviance } D(y_{xtg}, \hat{y}_{xtg}) . \end{aligned}$$

3. Update parameter $\hat{\alpha}_g$:

$$\begin{aligned} \hat{\alpha}_g &= \hat{\alpha}_g + \frac{\sum_{x,t} 2\omega(y - \hat{y})}{\sum_{x,t} 2\omega \hat{y}} \\ &\text{-- adjust the updated parameter such that } \hat{\alpha}_g = \hat{\alpha}_g - \hat{\alpha}_{g_1}, \text{ where } g_1 \text{ is the first} \\ &\text{level/group of the extra variate } g \text{ (i.e. set the first level as a base value);} \\ &\rightarrow \text{ calculate fitted values } \hat{y}(\hat{\alpha}_x, \hat{\alpha}_g, \hat{b}_x, \hat{\kappa}_t) \rightarrow \text{ calculate deviance } D(y_{xtg}, \hat{y}_{xtg}) . \end{aligned}$$

4. Update parameter $\hat{\kappa}_t$:

$$\begin{aligned} \hat{\kappa}_t &= \hat{\kappa}_t + \frac{\sum_{x,g} 2\omega(y - \hat{y})}{\sum_{x,g} 2\omega \hat{b}_x^2 \hat{y}} \\ &\text{-- adjust the updated parameter such that } \hat{\kappa}_t = \hat{\kappa}_t - \overline{\hat{\kappa}_t} ; \\ &\rightarrow \text{ calculate fitted values } \hat{y}(\hat{\alpha}_x, \hat{\alpha}_g, \hat{b}_x, \hat{\kappa}_t) \rightarrow \text{ calculate deviance } D(y_{xtg}, \hat{y}_{xtg}) . \end{aligned}$$

5. Update parameter $\hat{b}_x$:

$$\hat{b}_x = \hat{b}_x + \frac{\sum_{t,g} 2\omega(y - \hat{y})}{\sum_{t,g} 2\omega \hat{\kappa}_t^2 \hat{y}}$$

→ calculate fitted values $\hat{y}(\hat{\alpha}_x, \hat{\alpha}_g, \hat{b}_x, \hat{\kappa}_t)$ → calculate deviance $D_u(y_{xtg}, \hat{y}_{xtg})$.

6. Check deviance convergence:

$$\Delta D = D - D_u$$

were D_u is the updated deviance at step 5.

- if $\Delta D > 1 \times 10^{-6}$ ⇒ goto step 2.
- Stop iterative process once $\Delta D \approx 0$ and take the fitted parameters as the ML estimates to the observed data.
- Alternatively, stop if $\Delta D < 0$ for a consecutive 5 updating cycles and consider using other starting values or declare the iterations non-convergent.

7. Once convergence is achieved, re-scale the interaction parameters: $\hat{b}_x$ and $\hat{\kappa}_t$:

$$\hat{b}_x = \frac{\hat{b}_x}{\sum_x \hat{b}_x} \quad ; \quad \hat{\kappa}_t = \hat{\kappa}_t \times \left(\sum_x \hat{b}_x \right),$$

in order to satisfy the usual LC model constraints $\sum_t \kappa_t = 0$ and $\sum_x b_x = 1$.

4 Application of the Generalized LC Models in R with ilc

In the following, we present the most important features of using the *ilc* package to fit and analyse age and time dependent mortality models. The data manipulation and regression methods are illustrated in context of the CMI (lives) data containing the mortality experience of male life office pensioners retiring at or after normal retirement age. The data is made up of observed central exposure and deaths for ages 50-108, all durations combined, investigation years 1983-2003 (Source: Continuous Mortality Investigation). The main regression and diagnostic methods used in the *ilc* package are adequate to run independently, however most data formatting and life expectancy forecasting features are built such that to integrate with the **demography** and **forecast** packages of R, written by Rob J Hyndman.⁵ However, the *ilc* package accommodates many specific methods which allow improved inspection and graphical visualisation of both the mortality data and the regression outputs.

⁵Detailed reference manuals of the **demography** and other complementary packages are available at URL: www.robhyndman.info/Rlibrary/demography.

The `ilc` package has been developed and tested in the R statistical software version 2.8 and the following packages are required for its trouble-free use:⁶

- `demography` (version 0.98);
- `forecast` (version 1.11);
- `tseries` (version 0.10-14);
- `adddb` (version 3.221);
- `mgcv` (version 1.3-29);
- `zoo` (version 1.5-5).
- `survival` (version 2.34-1).

4.1 Package Installation

The `ilc` package is still in an early development phase and it is not fully prepared yet for unaided installation in R (e.g. issuing: `> install.packages("ilc")` command). Nevertheless, the program command functions and the mortality data used for illustration purposes are provided in a binary pre-compiled form and they can be made easily available in R. Thus, the compressed package archive (`ilc-v1.0.zip`) can be downloaded [here](#) (e.g. open the link in a web browser and save the zip file into a local directory.), which contains detailed instructions, examples for demonstration and the source code itself, which is provided for ease of reference and also to encourage the users to contribute fixes and new features. Thus, the installation of the `ilc` package content into R can be carried out in the following way:

1. Extract the `ilc-v1.0.zip` archive into a chosen working directory:
e.g. `"c:\Program Files\R"`;
(Note that this creates by default a subfolder `ilc`.)
2. From the R console set the working directory to the newly created folder:
`> setwd("c:/Program Files/R/ilc")`
3. Attach the pre-compiled program functions (as provided in the zip archive):
`> attach("ilc.rdata")`

⁶The latest versions of these packages can be downloaded and installed in R from the CRAN archive (or one of its mirrors) at: cran.r-project.org using the `> install.packages` function. We note that newer versions of the `demography` and `forecast` package than the ones illustrated here seem to run into errors when computing life expectancy forecasts.

4. Optionally, attach the pre-compiled CMI male pensioners' mortality experience data set – designated as: `dd.cmi.pens` – (as provided in the `demography` format in the zip archive for illustration purposes):


```
> attach("cmi.rdata")
```

4.2 Preparing the Mortality Data for Analysis

In order to fit the generalised LC type family models the mortality data need to be arranged in a `demogdata` class format of the `demography` package. For instance, assuming that the above mentioned CMI mortality experience is made up by the cross-tabulated mortality rates (`mu`) and the central exposures (`e`) by individual ages (`x`) and calendar years (`t`) sequences, we can create an R data object (`dd.cmi.pens`) for the generalised LC analysis by making use of the following purpose-built function:

```
> dd.cmi.pens <- demogdata(data=mu, pop=e, ages=x, years=t, type="mortality",
  label="CMI", name="male")
```

where the arguments `data` and `pop` must be matrices (or data-frames) of equal dimensions. Also, the arguments `label` and `name` are additional (string) qualifiers that specify the origin and the series (e.g. gender) of the data, respectively. Such data objects can contain more than one set of mortality experiences that can be identified by the `name` argument. For further details and examples of using the `demogdata` format/function, the reader is referred to the `demography` package manual. Following on, a summary description can be printed out by typing the data object's name:

```
> dd.cmi.pens
Mortality data for CMI
Series: male
Years: 1983 - 2003
Ages: 50 - 108
```

Alternatively, more detailed data inspections and/or graphical illustrations may be produced using the following type of commands:⁷

- print a query table of mortality rates:


```
> insp.dd(dd.cmi.pens, age=50:80, year=1985:1990)
```
- print a query table of central exposures:


```
> insp.dd(dd.cmi.pens, what='pop', age=70:100, year=1988:1993)
```

⁷Observe here that only the available segments of data are used whenever the ages and/or years sequences mismatch the given data array.

- print a query table of number of deaths:


```
> insp.dd(dd.cmi.pens, what='deaths', age=seq(100), year=1980:2010)
```
- produce simple plots (i.e. without legend) of log- or untransformed rates:


```
> plot(dd.cmi.pens)
> plot(dd.cmi.pens, transf=F)
```
- produce annotated plots (i.e. with legend) of log- or untransformed rates:


```
> plot.dd(dd.cmi.pens, xlim=c(40, 110), lpar=list(x.int=-0.2, y.int=0.9, cex=0.85))
```

 where the optional `lpar` list controls the legend layout (see pane *a*) of Figure 1)


```
> plot.dd(dd.cmi.pens, year=1985:1995, transf=F)
> plot.dd(dd.cmi.pens, year=1995:1997, transf=F, lty=1:3, col=1:3)
```
- produce annotated plots of number of deaths:


```
> tmp.d <- extract.deaths(dd.cmi.pens, ages=55:100)
      # without correction of empty cells, or
> tmp.d <- extract.deaths(dd.cmi.pens, ages=55:100, fill='perks')
      # This makes use of fill.demogdata() function to replace all
      # empty cells using the 'Perks' model (see ilc source code).
      # Other correction methods available are: 'interpolate' and 'mspline'
      # (see demography package manual).
> tmp.d$type <- 'mortality'
> plot.dd(tmp.d, year=1995:2003, transf=F, lty=1:8)
```

 (see pane *b*) of Figure 1)

Since in the case of the SLC model, the mortality experience is cross-classified by an additional covariate, the data set is best represented by a three dimensional matrix (i.e. `array`). For this purpose, the `ilc` package introduces a special class of data object (`rhdata`) that holds the necessary information about the grouping factors and the aggregate data of number of deaths, central exposures and the corresponding mortality rates. For example, consider a raw data set (`tab`) that comes in the form of individual observations of survival times and additional covariate(s), such as:

```
> tab[1:5, ] # show first 5 observations only
```

refno	dob	dev	event	cov1	cov2	(dob)	(dev)
1	-14485	15177	1	k	1	05/05/1920	21/07/2001
2	-13993	15177	1	j	1	09/09/1921	21/07/2001
3	-15800	15177	0	a	3	28/09/1916	21/07/2001
4	-15973	15177	1	c	2	08/04/1916	21/07/2001
5	-12776	15177	1	j	1	08/01/1925	21/07/2001

where the columns headed `dob` and `dev` represent the date of birth and of the date of event (i.e. 1=death, 0=survive), respectively, of individual cases with reference `refno`, that must be




Figure 1: Illustration of CMI (lives) pensioners mortality experience:
a) log central mortality rates and b) observed number of deaths.

entered in a format of class `date` (i.e. Julian dates – number of days since 1/1/1960, see `survival` package manual). Further, the last columns, headed `cov1` and `cov2`, represent some additional grouping factors (other than age and time) with observable levels `a–m` and `1–3`, respectively. Then the `rhdata` function of `ilc` can extract the aggregate data matrices for individual ages 60–95 over the period 2000–2005 by, say, `cov1` and place them in the appropriate format:⁸

```
> mtab <- rhdata(dat=tab, covar='cov1', xbreaks=60:96, xlabels=60:95,
  ybreaks=mdy.date(1,1,2000:2006), ylabels=2000:2005, name='M', label='MDat')
```

A short synopsis about the data source and the cross-tabulation parameters can be printed out by typing the newly created `rhdata` object's name:

```
> mtab
Multidimensional Mortality data for: MDat [M]
Across covariates:
years: 2000 - 2005
```

⁸We note that the column names `dob`, `dev` and `event` of the source data set (`tab`) cannot be changed.

```
ages: 60 - 95
cov1: a, b, c, d, e, f, g, h, i, j, k, l, m
```

Here, we note that the sub-grouping of the data set can be carried out by more than one additional covariate at once by specifying the argument `covar=c('cov1', 'cov2')`.

Currently, there are no user-friendly methods to extract specific parts of the `rhdata` class data object (e.g. covariate-specific tables by given ages and years). However, we can run the following commands to show the components of `mtab` data set:

- print a query table of mortality rates by individual ages 70-75 and by the first level ('a') of `cov1`:
`> mtab$mu[60:95%in%70:75, ,1]`
- print a query table of central exposures by individual ages 70-75 and level 'e' of `rscov1`:
`> mtab$pop[60:95%in%70:75, ,mtab$covariates$cov1%in%c('e')]`
- print a query table of number of deaths for all ages by levels 'k-m' in the first 3 years:
`> mtab$deaths[,1:3, 11:13]`

Due to the extensive data segmentation, we are likely to get a considerable number of undetermined mortality rates corresponding to zero exposures. Thus, it can be useful, before fitting the SLC model, to make use of a suitable 'closing-out' procedure to replace these data cells. This can be carried out with the aid of `fill.rhdata` function, as follows:

```
> mtab <- fill.rhdata(mtab, method='mspline')
# multidimensional wrapper of the fill.demogdata() function;
```

The above routine makes use of the `smooth.demogdata` function wherever it is needed in order to fit monotonic regression splines (see `demography` package manual) to the age-specific mortality rates and replaces all zero or missing values. Similarly, it is possible to make use of the 'interpolate' method from the `demography` package that interpolates between the values corresponding to the available nearby years of the same age group. An alternative smoothing method implemented in the `ilc` package is 'perks', which attempts to fit a generalised Perks model ($\mu_x = \frac{a}{1+\exp(b-px)}$) to the age-specific mortality rates (see [Thatcher 1999](#)).

For demonstration and/or testing purposes, it may be helpful to create an artificially stratified mortality experience with a Poisson error structure from a `demogdata` class object. The function `dd.rfp` can take a `demogdata` class object of 'mortality' type and adjust the observed log mortality rates by a vector of Poisson distributed additive effects (i.e. reduction factors) with predetermined means (for further details see `ilc` source code). For instance, taking the CMI experience as the base data set, we can produce a randomly stratified mortality data of `rhdata` format, as follows:

```
> rfp.cmi <- dd.rfp(dd.cmi.pens, rfp=c(0.5, 1.2, -0.7, 2.5))
```

with a data summary shown as

```
> rfp.cmi
Multidimensional Mortality data for: CMI [male]
Across covariates:
  years: 1983 - 2003
  ages: 50 - 108
  X: base, a, b, c, d
```

Plots of the central exposures and log mortality rates held in the `rfp.cmi` by the additional covariate (X) can be produced in the following way (see Figure 2):

```
> matplot(rfp.cmi$age, rfp.cmi$pop[,1]), type='l', xlab='Age',
  ylab='Ec', main='Base Level') # base level
> matplot(rfp.cmi$age, rfp.cmi$pop[,2]), type='l', xlab='Age',
  ylab='Ec', main='Level 1') # first level (a)
  :
> matplot(rfp.cmi$age, log(rfp.cmi$mu[,1]), type='l', xlab='Age',
  ylab='log(mu)', main='Base Level') # base level
> matplot(rfp.cmi$age, log(rfp.cmi$mu[,2]), type='l', xlab='Age',
  ylab='log(mu)', main='Level 1') # first level (a)
  :
```

The plots illustrated in Figure 2 of the randomised data (`rfp.cmi`) with respect to the (artificial) additional effect (X) show entirely indistinguishable central exposures and log mortality profiles. However, as will be demonstrated further on, the SLC fitting method can successfully identify the base mortality experience and estimate accurately the means of the additive effects.

4.3 Fitting the Mortality Models and Making Forecasts

In order to explore the fitted model objects and to run diagnostic checks, the `ilc` package caters for specialised methods of generic functions (like `coef`, `plot`, `fitted` and `residuals`) and also contains model specific utility functions (like `deviance.lca`, `residual.plots`, and `fitted.plots`). In the following, we illustrate the use of these tools and we give a brief interpretation of the outputs.




Figure 2: Illustration of randomised CMI (lives) pensioners mortality experience: central exposures and log central mortality rates by additional covariate (X).

4.3.1 Analysis of the Generalised LC Model Structures

The `lca.rh` is a universal routine of the `ilc` package developed to fit any of the six variants of the LC model structures (i.e. including the base LC model) using the iterative fitting method (see sections 2 and 3). The function arguments are defined as:⁹

```
> args(lca.rh)
function (dat, year = dat$year, age = dat$age, series = 1, max.age = 100, dec.conv = 6,
  clip = 3, error = c("poisson", "gaussian"), model = c("m", "h0", "h1", "h2", "ac", "lc"),
  restype = c("logrates", "rates", "deaths", "deviance"), scale = F, interpolate = F,
  verbose = T, spar = NULL)
```

The functionality of the arguments are aimed to be self-explanatory and user-friendly. In

⁹We acknowledge that `lca.rh` is designed to mimic some of the features and functionality of the `lca` function of the `demography` package. Also, as mentioned before, it makes use of the 'interpolate' correction method to replace missing data cells. However, the modelling and fitting methodology implemented in `lca.rh` are based entirely on the iterative approach presented in this paper.

the following we clarify further the main features:

`dat` : source data object of `demogdata` class;

`series` : target series to be used from the source data;

`dec.conv` : number of decimal places used to achieve convergence;

`clip` : number of marginal cohorts to remove from the rectangular data array (i.e. give 0 weights – it's only applicable to the first 5 models);

`error` : type of error structure of the model choice;

`model` : model choice (see section 2.3) – it can be a character or a numeric value (1-6) corresponding to the described models;

`restype` : type of residuals, which controls the type of the fitted value too;

Thus, in the cases of 'logrates' and 'rates' the function returns as fitted values the log and untransformed mortality rates, respectively. Likewise, the choices of 'deaths' and 'deviance' correspond to the fitted number of deaths.

`scale` : based on lca of `demography` package to re-scale the interaction parameters so that the κ_t has drift parameter equal to 1;

`spar` : numerical smoothing spline parameter (see `smooth.spline` function);

If not NULL (i.e. ranging from 0 to 1, with a recommended value of 0.6) the interaction effects $(\beta_x^{(0,1)})$ are smoothed out after fitting. As a consequence, the period/cohort effects are adjusted accordingly.

`verbose` : logical parameter to control the output amount of process information;

If set to TRUE the program prints out the updated deviance values along with the starting and final parameter estimates.

In the following two examples, we aim to give a general feel of how to make use of the above iterative fitting routine and then we discuss briefly the program outputs:

1) Estimate the base LC model (with Poisson errors)

In this application, we make use of the CMI (lives) data up to the age of 100 to avoid any data irregularities at very old ages and any remaining 0/NA values we can replace by interpolation:

```
> mod6 <- lca.rh(dd.cmi.pens, mod='lc', interpolate=T, verbose=F)
```

Original sample: Mortality data for CMI

Series: male

Years: 1983 - 2003

Ages: 50 - 108

Applied sample: Mortality data for CMI (Corrected: interpolate)

Series: male

Years: 1983 - 2003

Ages: 50 - 100

Fitting model: [$LC = a(x) + b1(x) * k(t)$]

- with Poisson error structure and with deaths as weights -

Iterations finished in: 14 steps

Warning messages:

- 1: In lca.set(dat, year, age, series, max.age, interpolate) :
⇒ data above age 100 are grouped.
- 2: A total of 62 0/NA central mortality rates are re-estimated by the "interpolate" method.
- 3: In lca.set(dat, year, age, series, max.age, interpolate) :
There are 45 cells with 0/NA exposures, which are ignored in the current analysis.
Try reducing the maximum age or choosing a different age range.
Alternatively, fit LC model with error= "gaussian" .




Figure 3: LC regression parameters for CMI male pensioners (lives) for age range 50 – 100 over the observation period of 1983 – 2003.

We note here that the same call to `lca.rh` function with `error="gaussian"` setting, computes the standard LC model of Lee and Carter (1992), however, using the iterative fitting method instead of the traditional SVD. Alternatively, the `lca` function of the `demography` package can fit the standard LC model with SVD approach by issuing a call like:

```
> modlc <- lca(dd.cmi.pens, interpolate=T, adjust='none')
```

that yields the same parameter estimates (for further details of using the `lca` function, the reader is referred to the `demography` manual).

A short printout of the model summary is produced by:

```
> mod6
```

Iterative Lee-Carter Family Regression:
Fitted Model: $LC = a(x) + b1(x) * k(t)$

Call: `lca.rh(dat = dd.cmi.pens, model = "lc", interpolate = T, verbose = F)`

Error Structure: poisson

Data Source: CMI [male] over

calendar years: (1983 - 2003) and ages: (50 - 100)

Deviance convergence in: 14 iterations

		dev	dev.c	df	df.c
1	Mean deviance base	1.386		df base	905
2	Mean deviance total	1.733		df tot	969

The estimated model parameters can be printed out using the `coef` function:

```
> coef(mod6)
```

	ax	ax.c	bx1	bx1.c	kt	kt.c
1	50	-3.665	50	0.110	1983	13.735
2	51	-4.199	51	0.048	1984	11.988
3	52	-4.633	52	0.037	1985	12.331
4	53	-4.812	53	0.017	1986	9.747
5	54	-4.664	54	0.013	1987	10.772
:						

where the columns headed with `.c` extension give the estimated coefficients and the other columns indicate the corresponding parameter labels. Alternatively, we can illustrate graphically the fitted parameters (see Figure 3) by the simple command:

```
> plot(mod6)
```




Figure 4: LC cross-classified fitted values for CMI male pensioners (lives) for age range 50 – 100 over the observation period of 1983 – 2003.
a) by age versus year and b) by year versus age

Further graphical illustrations of the regression outcome can be produced with the following command:

```
> fitted.plots(mod6)
```

that plots the cross-classified fitted values by age against calendar year; and also by year against age (see Figure 4 panes a) and b), respectively).

According to [Renshaw and Haberman \(2006\)](#), the preferred type of residuals to conduct diagnostic checks on the model are the standardised deviance residuals. Thus, we should change the current LC fitted object's residual values from 'logrates' type, which was only needed in order to produce the corresponding fitted values. In order to compute the 'deviance' residuals from a fitted object with different type of residuals, we can make use of the function `lca.dev.res`, though this utility also needs the central exposures matrix used in the LC fitting (see source code for further details and examples). In cases where deviance convergence is achieved fairly quickly, it is also possible to simply re-fit the original model, as follows:

```
> mod6d <- lca.rh(dd.cmi.pens, mod='lc', restype='deviance', interpolate=T, verbose=F)
```

Then, we can run the residuals plotting method on the new output object (see Figure 5):

```
> residual.plots(mod6d)
```

although, we note that the above function works on any type of residuals of the LC class family models.

Finally, we can produce forecasts of future mortality improvements and the corresponding future life expectancy based on the fitted LC model. The `ilc` application makes use of the `forecast` package to predict future values of the trend parameter (κ_t) using a traditional ARIMA(0, 1, 0) model over a given time horizon. This is accomplished by running the `forecast` method on the fitted model object. For instance, in order to produce a forecast over a 20 years period, we can issue the following type of command:

```
> forc6 <- forecast(mod6, h=20, jump='fit', level=90, shift=F)
```

which returns a “`fmforecast`” class object that contains the predicted mean trend parameter and the corresponding predicted mean mortality rates, alongside with their lower and upper limits of a 90 % confidence interval (CI).

We can visualise the forecasted log-mortality rates with the `demogdata` plotting method:

```
plot.dd(forc6, xlim=c(45, 100), lpar=list(x.int=-0.2, y.int=0.9, cex=0.95))
```

Figure 6 shows the above and we can note that the overly low rates at age 50 are the results of the corresponding peaked interaction effect (β_{50}), as it can be seen in Figure 3.

Further, the forecast object `forc6` also contains the predicted mean life expectancy and its 90 % CI, which can be extracted by:

```
> forc6$e0
Time Series:
Start = 2004
End = 2023
Frequency = 1
```

	e0	e0.lo	e0.hi
2004	34.18014	33.63142	34.73744
2005	34.46238	33.66483	35.28137
2006	34.74726	33.74325	35.79047
2007	35.03522	33.84356	36.28940
2008	35.32670	33.95714	36.78839

However, we can also compute life expectancy forecasts at other ages too by making use of the




Figure 5: LC standardised deviance residuals for CMI male pensioners (lives) for age range 50 – 100 over the observation period of 1983 – 2003.




Figure 6: LC future log mortality rates values for CMI male pensioners (lives) for age range 50 – 100 over a 20 years prediction horizon.

following `demography` package command, say, at target age of 60:¹⁰

```
> le6 <- life.expectancy(forc6, age=60)
```

The `ilc` package contains two specialised functions: `fle.plot` and `fle.plot` that can make forecasts and produce the corresponding plots directly from the LC model object. The former creates plots only of the predicted (period) life expectancy at any age with the chosen prediction interval (PI), whereas the latter produces the plots of both the predicted trend parameter and the predicted life expectancy at any age alongside the estimated PIs. For example, Figure 7 illustrates the plotting output of the following command:

```
> fle.plot(mod6, at=60, h=30, level=90)
```

with the same parameter settings as in the previous examples.

2) Estimate the APC model (with Poisson errors)

In this application we make use of the CMI data using a restricted age range (e.g. to avoid data correction) and 'deviance' residuals. It is possible to choose a reduced convergence precision to achieve faster processing, although for proper fit it is recommended to use the default value

¹⁰Further details about the application of this command are available in the `demography` package help files – e.g. by typing at the R console `> ?life.expectancy`.




Figure 7: Illustration of LC forecast over a 20 years prediction horizon with 90% CI for CMI male pensioners (lives)
 a) trend parameter κ_t and b) future life expectancy at age 60.

(it can lead to a slower convergence cycle for this model):

```
> mod1 <- lca.rh(dd.cmi.pens, age=60:95, res='dev', dec=3, verb=F)
```

Original sample: Mortality data for CMI

Series: male

Years: 1983 - 2003

Ages: 50 - 108

Applied sample: Mortality data for CMI

Series: male

Years: 1983 - 2003

Ages: 60 - 95

Fitting model: [$M = a(x) + b_0(x) \cdot i(t-x) + b_1(x) \cdot k(t)$]

- with Poisson error structure and with deaths as weights -

Iterations finished in: 445 steps

Warning messages:

1: In lca.set(dat, year, age, series, max.age, interpolate) :

There are 1 cells with 0/NA mu, which are ignored in the current analysis.

Try reducing the maximum age or setting interpolate=TRUE.

```
2: In lca.rh(dd.cmi.pens, age = 60:95, int = F, res = "dev", dec = 3, :
```

```
    The cohorts outside [1891, 1940] were zero weighted (clipped).
```

The corresponding model summary can be printed out by writing:

```
> mod1
```

```
Iterative Lee-Carter Family Regression:
```

```
Fitted Model:  $M = a(x) + b_0(x) \cdot i(t-x) + b_1(x) \cdot k(t)$ 
```

```
Call: lca.rh(dat = dd.cmi.pens, age = 60:95, dec.conv = 3, restype = "dev",
  interpolate = F, verbose = F)
```

```
Error Structure: poisson
```

```
Data Source: CMI [male] over
```

```
    calendar years: (1983 - 2003) and ages: (60 - 95)
```

```
Deviance convergence in: 445 iterations
```

	dev	dev.c	df	df.c
1 Mean deviance base	1.386		df base	597
2 Mean deviance total	1.648		df tot	684

Similarly, in the case of the APC fitted model, we can repeat the above procedures to investigate the regression, that gives the following outputs:

```
> coef(mod1)
```

	itx	itx.c	ax	ax.c	bx0.c	bx1.c	kt	kt.c
1	1888	0.000	60	-3.923	-0.049	0.051	1983	0
2	1889	0.000	61	-4.15	-0.021	0.03	1984	-1.267
3	1890	0.000	62	-4.307	0.04	0.039	1985	-1.001
4	1891	3.381	63	-4.394	0.027	0.018	1986	-2.901
5	1892	3.649	64	-4.117	0.06	0.027	1987	-2.357
:								

where we can note that both trend parameters (κ_t , ι_{t-x}) are re-scaled during fitting to start from 0 (see section 3.2). The regression plot in Figure 8 reveals a strong cohort effect for the pensioners born between 1910–1920:

```
> plot(mod1)
```

The other 4 model constructs can be estimated in a similar way by entering the corresponding `model` argument value in the main function call. Usually, in the case of large data sets, the fitting cycle is fast and produces stable parameter estimates. We have not yet implemented any object oriented methods in the `ilc` package to produce forecasts for the models that allow




Figure 8: APC regression parameters for CMI male pensioners (lives) for age range 60 – 95 over the observation period of 1983 – 2003.

for the cohort effect. This feature is going to be developed in future versions of the software.

4.3.2 Analysis of the Stratified LC Model

The `ilc` package provides the purpose-built `elca.rh` program to fit the extended (i.e. stratified) LC model structure using the iterative fitting method (see sections 2.4 and 3.3). This function follows closely the structure of the `lca.rh` regression routine and offers the same choice of argument settings. In addition, fixed base age effect (α_x) can be imputed through the optional `ax.fix` argument, which are then not modified during the fitting process. The function is specified in the following way:¹¹

```
> args(elca.rh)
function(dat, year=dat$year, age=dat$age, dec.conv = 6, error = c("poisson", "gaussian"),
  restype = c("logrates", "rates", "deaths", "deviance"), scale = F, interpolate = F,
  verbose = T, spar=NULL, ax.fix = NULL)
```

¹¹Observe that the `max.age` feature is not implemented in the current version of the `elca.rh` function.

where the arguments listed below have an updated functionality from the previous description:

dat : source data object of **rhdata** class with only one additional grouping factor (i.e. covariate other than age and time);

ax.fix : vector of predetermined parameter estimates of the main (base) age effect, which must be of the same length as the **age** argument.

Therefore, if it is not NULL the parameter (α_x) is ignored during the updating cycle.

The only multidimensional data sets available to us, that were used to develop this part of the program, are currently commercially sensitive and thus are restricted for publication. Nevertheless, we can still demonstrate the use of the program on the randomly stratified CMI mortality data set (**rfp.cmi**) presented in section 4.2, as follows:

```
> mod6e <- elca.rh(rfp.cmi, age=50:100, interp=T, dec=3, verb=F)
Original sample: Multidimensional Mortality data for: CMI [male]
Across covariates:
  years: 1983 - 2003
  ages: 50 - 108
  X: base, a, b, c, d
Applied sample: Multidimensional Mortality data for: CMI [male]
Across covariates:
  years: 1983 - 2003
  ages: 50 - 100
  X: base, a, b, c, d
Fitting model: [ LC(g) = a(x)+a(g)+b(x)*k(t) ]
               - with Poisson error structure and with deaths as weights -
Iterations finished in: 38 steps
Warning messages:
1: A total of 1160 0/NA central mortality rates are re-estimated by the "interpolate" method.
2: In elca.rh(rfp.cmi, age = 50:100, int = T, dec = 3, verb = F) :
   There are 152 cells with 0/NA exposures, which are ignored in the current analysis.
   Try reducing the fitted age range.
   Alternatively, fit ELC model with error= "gaussian" .
```

The corresponding model summary output is provided by writing:

```
> mod6e
```

```
Extended Lee-Carter Regression:
Fitted Model: LC(g) = a(x)+a(g)+b(x)*k(t)
```

Call: `elca.rh(dat = rfp.cmi, age = 50:100, dec.conv = 3, interpolate = T, verbose = F)`

Error Structure: poisson

Data Source: CMI : male over

calendar years: (1983 - 2003) , ages: (50 - 100)

and groups: base a b c d

Deviance convergence in: 38 iterations

	dev	dev.c	df	df.c
1 Mean deviance base	264.316		df base	3648
2 Mean deviance total	202.249		df tot	4845




Figure 9: SPC regression parameters for artificially stratified CMI male pensioners (lives) for age range 50 – 100 over the observation period of 1983 – 2003.

Also, we can print out the fitted parameter values of the additive effect as:

```
> coef(mod6e)
```

	ax	ax.c	bx.c	kt	kt.c	ag	ag.c
1	50	-4.033	0.164	1983	16.162	base	0
2	51	-4.319	0.044	1984	11.457	a	0.496
3	52	-4.801	0.022	1985	13.075	b	1.17

4	53	-4.896	0.030	1986	10.106	c	-0.735
5	54	-4.874	0.015	1987	8.772	d	2.518
⋮							

We note that the fitting algorithm converges fairly quickly, in just 38 iterations, when using the precision of `dec.conv=3` and estimates the parameters of the additional effect (see values in column `ag.c`) close to the simulated Poisson means (i.e. $\text{rfp}=\text{c}(0.5, 1.2, -0.7, 2.5)$). Also, considering the extent of noise imposed on the base CMI data (see Figure 2), the remaining parameter estimates are overall similar to the coefficients of the standard LC (`lca.rh`) fit with Poisson error structure (`mod6`), as it can be seen in the corresponding interaction and period effects shown in the plots of Figures 3 and 9. Based on empirical trials carried out on actual mortality data, we can report that the parameters of the bilinear term, practically, remain the same after adding an observed additional effect (α_g) to the model.




Figure 10: Illustration of forecast result over a 20 years period in the SLC modelling framework:

a) future trend parameter and b) future life expectancy at age 60.

Once we allow for the stratification of the main effect parameter, forecasting in the SLC modelling framework can proceed along the same method applied in the traditional LC ap-

proach (see section 2.5). In the current `ilc` package, there are no specialized methods to produce predictions directly, but we can still make use of the `demography` package `forecast.lca` functions to produce forecasted trend parameter (κ_t). Then, we can make use of an adapted version of the `fle.plot` method to illustrate the corresponding future life expectancy differentiated by the additional effect using the following commands:

```
> mod6ef <- forecast.lca(mod6e, h=20, level=90, jump='fit', shift=F)
> plot(mod6ef$kt, ylab='kt', xlab='Year')
> matfle.plot(mod6e$lca, mod6, at=60, label='RFP CMI', h=20)
```

Thus, Figure 10 illustrates the resulting plots of predicted trend parameter (panel *a*) and the future life expectancy at age 60 over a 20 year period (panel *b*).

References

- Booth, H. (2006), “Demographic forecasting: 1980 to 2005 in review”, *International Journal of Forecasting* **22**(3), 547–581.
- Booth, H., Hyndman, R. J., Tickle, L. and de Jong, P. (2006), “Lee-Carter mortality forecasting: a multi-country comparison of variants and extensions”, *Demographic Research* **15**(Article 9), 289–310. (<http://www.demographic-research.org/Volumes/Vol15/9/>).
- Brouhns, N., Denuit, M. and Vermunt, J. (2002), “A Poisson log-bilinear regression approach to the construction of projected life tables”, *Insurance: Mathematics and Economics* **31**, 373–393.
- Guterman, S. and Vanderhoof, I. (2000), “Forecasting changes in mortality: a search for a law of causes and effects”, *North American Actuarial Journal* **2**, 135–138.
- Lee, L. (2000), “The Lee-Carter method for forecasting mortality, with various extensions and applications”, *North American Actuarial Journal* **4**, 80–93.
- Lee, R. and Carter, L. (1992), “Modelling and forecasting U.S. mortality”, *Journal of the American Statistical Association* **87**, 659–671.
- Lee, R. and Nault, F. (1993), *Modeling and forecasting provincial mortality in Canada*. Paper presented at the World Congress of International Union for the Scientific Study of Population, Montreal. Unpublished.
- Li, N. and Lee, R. D. (2005), “Coherent mortality forecasts for a group of populations: an extension of the Lee-Carter method”, *Demography* **42**(3), 575–594.
- Renshaw, A. E. and Haberman, S. (2003a), “Lee-Carter mortality forecasting: a parallel generalized linear modelling approach for England and Wales mortality projections”, *Journal of the Royal Statistical Society, Series C* **52**(1), 119–137.

- Renshaw, A. E. and Haberman, S. (2003*b*), “Lee-Carter mortality forecasting with age specific enhancement”, *Insurance: Mathematics and Economics* **33**, 255–272.
- Renshaw, A. E. and Haberman, S. (2006), “A cohort-based extension to the Lee-Carter model for mortality reduction factors”, *Insurance: Mathematics and Economics* **38**, 556–570.
- Renshaw, A. E. and Haberman, S. (2008), “On simulation-based approaches to risk measurement in mortality with specific reference to Poisson Lee-Carter modelling”, *Insurance: Mathematics and Economics* **42**(2), 797–816.
- Renshaw, A. E. and Haberman, S. (2009), “On age-period-cohort parametric mortality rate projections”, *Insurance: Mathematics and Economics* . To appear.
- Tabeau, E. (2001), A review of demographic forecasting models for mortality, *in* E. Tabeau, A. Van Den Berg Jeths and C. Heathcote, eds, “Forecasting mortality in developed countries: insights from a statistical, demographic and epidemiological perspective”, Dordrecht: Kluwer Academic Publishers, pp. 1–32.
- Thatcher, A. R. (1999), “The long-term pattern of adult mortality and the highest attained age”, *Journal of the Royal Statistical Society, Series A* **162**, 5–30.
- Venables, W. N., Smith, D. M. and the R Development Core Team (2005), “An introduction to R. Notes on R: A programming environment for data analysis and graphics”. Version 2.1.1 (2005-06-20). ISBN 3-900051-12-7.
- Wilmoth, J. (1993), Computational methods for fitting and extrapolating the Lee-Carter model of mortality change, Technical report, Berkeley: Department of Demography, University of California.



FACULTY OF ACTUARIAL SCIENCE AND INSURANCE

Actuarial Research Papers since 2001

Report Number	Date	Publication Title	Author
135.	February 2001.	On the Forecasting of Mortality Reduction Factors. ISBN 1 901615 56 1	Steven Haberman Arthur E. Renshaw
136.	February 2001.	Multiple State Models, Simulation and Insurer Insolvency. ISBN 1 901615 57 X	Steve Haberman Zoltan Butt Ben Rickayzen
137.	September 2001	A Cash-Flow Approach to Pension Funding. ISBN 1 901615 58 8	M. Zaki Khorasanee
138.	November 2001	Addendum to "Analytic and Bootstrap Estimates of Prediction Errors in Claims Reserving". ISBN 1 901615 59 6	Peter D. England
139.	November 2001	A Bayesian Generalised Linear Model for the Bornhuetter- Ferguson Method of Claims Reserving. ISBN 1 901615 62 6	Richard J. Verrall
140.	January 2002	Lee-Carter Mortality Forecasting, a Parallel GLM Approach, England and Wales Mortality Projections. ISBN 1 901615 63 4	Arthur E. Renshaw Steven Haberman.
141.	January 2002	Valuation of Guaranteed Annuity Conversion Options. ISBN 1 901615 64 2	Laura Ballotta Steven Haberman
142.	April 2002	Application of Frailty-Based Mortality Models to Insurance Data. ISBN 1 901615 65 0	Zoltan Butt Steven Haberman
143.	Available 2003	Optimal Premium Pricing in Motor Insurance: A Discrete Approximation.	Russell J. Gerrard Celia Glass
144.	December 2002	The Neighbourhood Health Economy. A Systematic Approach to the Examination of Health and Social Risks at Neighbourhood Level. ISBN 1 901615 66 9	Les Mayhew
145.	January 2003	The Fair Valuation Problem of Guaranteed Annuity Options : The Stochastic Mortality Environment Case. ISBN 1 901615 67 7	Laura Ballotta Steven Haberman
146.	February 2003	Modelling and Valuation of Guarantees in With-Profit and Unitised With-Profit Life Insurance Contracts. ISBN 1 901615 68 5	Steven Haberman Laura Ballotta Nan Want
147.	March 2003.	Optimal Retention Levels, Given the Joint Survival of Cedent and Reinsurer. ISBN 1 901615 69 3	Z. G. Ignatov Z.G., V. Kaishev R.S. Krachunov
148.	March 2003.	Efficient Asset Valuation Methods for Pension Plans. ISBN 1 901615707	M. Iqbal Owadally
149.	March 2003	Pension Funding and the Actuarial Assumption Concerning Investment Returns. ISBN 1 901615 71 5	M. Iqbal Owadally

Report Number	Date	Publication Title	Author
150.	Available August 2004	Finite time Ruin Probabilities for Continuous Claims Severities	D. Dimitrova Z. Ignatov V. Kaishev
151.	August 2004	Application of Stochastic Methods in the Valuation of Social Security Pension Schemes. ISBN 1 901615 72 3	Subramaniam Iyer
152.	October 2003.	Guarantees in with-profit and Unitized with profit Life Insurance Contracts; Fair Valuation Problem in Presence of the Default Option ¹ . ISBN 1-901615-73-1	Laura Ballotta Steven Haberman Nan Wang
153.	December 2003	Lee-Carter Mortality Forecasting Incorporating Bivariate Time Series. ISBN 1-901615-75-8	Arthur E. Renshaw Steven Haberman
154.	March 2004.	Operational Risk with Bayesian Networks Modelling. ISBN 1-901615-76-6	Robert G. Cowell Yuen Y, Khuen Richard J. Verrall
155.	March 2004.	The Income Drawdown Option: Quadratic Loss. ISBN 1 901615 7 4	Russell Gerrard Steven Haberman Bjorn Hojgarrd Elena Vigna
156.	April 2004	An International Comparison of Long-Term Care Arrangements. An Investigation into the Equity, Efficiency and sustainability of the Long-Term Care Systems in Germany, Japan, Sweden, the United Kingdom and the United States. ISBN 1 901615 78 2	Martin Karlsson Les Mayhew Robert Plumb Ben D. Rickayzen
157.	June 2004	Alternative Framework for the Fair Valuation of Participating Life Insurance Contracts. ISBN 1 901615-79-0	Laura Ballotta
158.	July 2004.	An Asset Allocation Strategy for a Risk Reserve considering both Risk and Profit. ISBN 1 901615-80-4	Nan Wang
159.	December 2004	Upper and Lower Bounds of Present Value Distributions of Life Insurance Contracts with Disability Related Benefits. ISBN 1 901615-83-9	Jaap Spreeuw
160.	January 2005	Mortality Reduction Factors Incorporating Cohort Effects. ISBN 1 90161584 7	Arthur E. Renshaw Steven Haberman
161.	February 2005	The Management of De-Cumulation Risks in a Defined Contribution Environment. ISBN 1 901615 85 5.	Russell J. Gerrard Steven Haberman Elena Vigna
162.	May 2005	The IASB Insurance Project for Life Insurance Contracts: Impact on Reserving Methods and Solvency Requirements. ISBN 1-901615 86 3.	Laura Ballotta Giorgia Esposito Steven Haberman
163.	September 2005	Asymptotic and Numerical Analysis of the Optimal Investment Strategy for an Insurer. ISBN 1-901615-88-X	Paul Emms Steven Haberman
164.	October 2005.	Modelling the Joint Distribution of Competing Risks Survival Times using Copula Functions. ISBN 1-901615-89-8	Vladimir Kaishev Dimitrina S. Dimitrova Steven Haberman
165.	November 2005.	Excess of Loss Reinsurance Under Joint Survival Optimality. ISBN1-901615-90-1	Vladimir K. Kaishev Dimitrina S. Dimitrova

Report Number	Date	Publication Title	Author
166.	November 2005.	Lee-Carter Goes Risk-Neutral. An Application to the Italian Annuity Market. ISBN 1-901615-91-X	Enrico Biffis Michel Denuit
167.	November 2005	Lee-Carter Mortality Forecasting: Application to the Italian Population. ISBN 1-901615-93-6	Steven Haberman Maria Russolillo
168.	February 2006	The Probationary Period as a Screening Device: Competitive Markets. ISBN 1-901615-95-2	Jaap Spreeuw Martin Karlsson
169.	February 2006	Types of Dependence and Time-dependent Association between Two Lifetimes in Single Parameter Copula Models. ISBN 1-901615-96-0	Jaap Spreeuw
170.	April 2006	Modelling Stochastic Bivariate Mortality ISBN 1-901615-97-9	Elisa Luciano Jaap Spreeuw Elena Vigna.
171.	February 2006	Optimal Strategies for Pricing General Insurance. ISBN 1901615-98-7	Paul Emms Steve Haberman Irene Savoulli
172.	February 2006	Dynamic Pricing of General Insurance in a Competitive Market. ISBN1-901615-99-5	Paul Emms
173.	February 2006	Pricing General Insurance with Constraints. ISBN 1-905752-00-8	Paul Emms
174.	May 2006	Investigating the Market Potential for Customised Long Term Care Insurance Products. ISBN 1-905752-01-6	Martin Karlsson Les Mayhew Ben Rickayzen
175.	December 2006	Pricing and Capital Requirements for With Profit Contracts: Modelling Considerations. ISBN 1-905752-04-0	Laura Ballotta
176.	December 2006	Modelling the Fair Value of Annuities Contracts: The Impact of Interest Rate Risk and Mortality Risk. ISBN 1-905752-05-9	Laura Ballotta Giorgia Esposito Steven Haberman
177.	December 2006	Using Queuing Theory to Analyse Completion Times in Accident and Emergency Departments in the Light of the Government 4-hour Target. ISBN 978-1-905752-06-5	Les Mayhew David Smith
178.	April 2007	In Sickness and in Health? Dynamics of Health and Cohabitation in the United Kingdom. ISBN 978-1-905752-07-2	Martin Karlsson Les Mayhew Ben Rickayzen
179.	April 2007	GeD Spline Estimation of Multivariate Archimedean Copulas. ISBN 978-1-905752-08-9	Dimitrina Dimitrova Vladimir Kaishev Spiridon Penev
180.	May 2007	An Analysis of Disability-linked Annuities. ISBN 978-1-905752-09-6	Ben Rickayzen
181.	May 2007	On Simulation-based Approaches to Risk Measurement in Mortality with Specific Reference to Poisson lee-Carter Modelling. ISBN 978-1-905752-10-2	Arthur Renshaw Steven Haberman

Report Number	Date	Publication Title	Author
182.	July 2007	High Dimensional Modelling and Simulation with Asymmetric Normal Mixtures. ISBN 978-1-905752-11-9	Andreas Tsanakas Andrew Smith
183.	August 2007	Intertemporal Dynamic Asset Allocation for Defined Contribution Pension Schemes. ISBN 978-1-905752-12-6	David Blake Douglas Wright Yumeng Zhang
184.	October 2007	To split or not to split: Capital allocation with convex risk measures. ISBN 978-1-905752-13-3	Andreas Tsanakas
185.	April 2008	On Some Mixture Distribution and Their Extreme Behaviour. ISBN 978-1-905752-14-0	Vladimir Kaishev Jae Hoon Jho
186.	October 2008	Optimal Funding and Investment Strategies in Defined Contribution Pension Plans under Epstein-Zin Utility. ISBN 978-1-905752-15-7	David Blake Douglas Wright Yumeng Zhang
187.	May 2008	Mortality Risk and the Valuation of Annuities with Guaranteed Minimum Death Benefit Options: Application to the Italian Population. ISBN 978-1-905752-16-4	Steven Haberman Gabriella Piscopo
188.	January 2009	The Market Potential for Privately Financed Long Term Care Products in the UK. ISBN 978-1-905752-19-5	Leslie Mayhew
189.	June 2009	Whither Human Survival and Longevity or the Shape of things to Come. ISBN 978-1-905752-21-8	Leslie Mayhew David Smith
190	September 2009	ilc: A Collection of R Functions for Fitting a Class of Lee Carter Mortality Models using Iterative fitting Algorithms* ISBN 978-1-905752-22-5	Zoltan Butt Steven Haberman

Statistical Research Papers

1.	December 1995.	Some Results on the Derivatives of Matrix Functions. ISBN 1 874 770 83 2	P. Sebastiani
2.	March 1996	Coherent Criteria for Optimal Experimental Design. ISBN 1 874 770 86 7	A.P. Dawid P. Sebastiani
3.	March 1996	Maximum Entropy Sampling and Optimal Bayesian Experimental Design. ISBN 1 874 770 87 5	P. Sebastiani H.P. Wynn
4.	May 1996	A Note on D-optimal Designs for a Logistic Regression Model. ISBN 1 874 770 92 1	P. Sebastiani R. Settini
5.	August 1996	First-order Optimal Designs for Non Linear Models. ISBN 1 874 770 95 6	P. Sebastiani R. Settini
6.	September 1996	A Business Process Approach to Maintenance: Measurement, Decision and Control. ISBN 1 874 770 96 4	Martin J. Newby
7.	September 1996.	Moments and Generating Functions for the Absorption Distribution and its Negative Binomial Analogue. ISBN 1 874 770 97 2	Martin J. Newby
8.	November 1996.	Mixture Reduction via Predictive Scores. ISBN 1 874 770 98 0	Robert G. Cowell.
9.	March 1997.	Robust Parameter Learning in Bayesian Networks with Missing Data. ISBN 1 901615 00 6	P. Sebastiani M. Ramoni

Report Number	Date	Publication Title	Author
10.	March 1997.	Guidelines for Corrective Replacement Based on Low Stochastic Structure Assumptions. ISBN 1 901615 01 4.	M.J. Newby F.P.A. Coolen
11.	March 1997	Approximations for the Absorption Distribution and its Negative Binomial Analogue. ISBN 1 901615 02 2	Martin J. Newby
12.	June 1997	The Use of Exogenous Knowledge to Learn Bayesian Networks from Incomplete Databases. ISBN 1 901615 10 3	M. Ramoni P. Sebastiani
13.	June 1997	Learning Bayesian Networks from Incomplete Databases. ISBN 1 901615 11 1	M. Ramoni P. Sebastiani
14.	June 1997	Risk Based Optimal Designs. ISBN 1 901615 13 8	P. Sebastiani H.P. Wynn
15.	June 1997.	Sampling without Replacement in Junction Trees. ISBN 1 901615 14 6	Robert G. Cowell
16.	July 1997	Optimal Overhaul Intervals with Imperfect Inspection and Repair. ISBN 1 901615 15 4	Richard A. Dagg Martin J. Newby
17.	October 1997	Bayesian Experimental Design and Shannon Information. ISBN 1 901615 17 0	P. Sebastiani. H.P. Wynn
18.	November 1997.	A Characterisation of Phase Type Distributions. ISBN 1 901615 18 9	Linda C. Wolstenholme
19.	December 1997	A Comparison of Models for Probability of Detection (POD) Curves. ISBN 1 901615 21 9	Wolstenholme L.C
20.	February 1999.	Parameter Learning from Incomplete Data Using Maximum Entropy I: Principles. ISBN 1 901615 37 5	Robert G. Cowell
21.	November 1999	Parameter Learning from Incomplete Data Using Maximum Entropy II: Application to Bayesian Networks. ISBN 1 901615 40 5	Robert G. Cowell
22.	March 2001	FINEX : Forensic Identification by Network Expert Systems. ISBN 1 901615 60X	Robert G. Cowell
23.	March 2001.	Wren Learning Bayesian Networks from Data, using Conditional Independence Tests is Equivalent to a Scoring Metric ISBN 1 901615 61 8	Robert G Cowell
24.	August 2004	Automatic, Computer Aided Geometric Design of Free-Knot, Regression Splines. ISBN 1-901615-81-2	Vladimir K Kaishev, Dimitrina S. Dimitrova, Steven Haberman Richard J. Verrall
25.	December 2004	Identification and Separation of DNA Mixtures Using Peak Area Information. ISBN 1-901615-82-0	R.G. Cowell S.L. Lauritzen J Mortera,
26.	November 2005.	The Quest for a Donor : Probability Based Methods Offer Help. ISBN 1-90161592-8	P.F. Mostad T. Egeland., R.G. Cowell V. Bosnes Ø. Braaten

Report Number	Date	Publication Title	Author
27.	February 2006	Identification and Separation of DNA Mixtures Using Peak Area Information. (Updated Version of Research Report Number 25). ISBN 1-901615-94-4	R.G.Cowell S.L.Lauritzen J Mortera,
28.	October 2006	Geometrically Designed, Variable Knot Regression Splines : Asymptotics and Inference. ISBN 1-905752-02-4	Vladimir K Kaishev Dimitrina S.Dimitrova Steven Haberman Richard J. Verrall
29.	October 2006	Geometrically Designed, Variable Knot Regression Splines : Variation Diminishing Optimality of Knots. ISBN 1-905752-03-2	Vladimir K Kaishev Dimitrina S.Dimitrova Steven Haberman Richard J. Verrall
30.	November 2008	Scheduling Reentrant Jobs on Parallel Machines with a Remote Server. ISBN 978-1-905752-18-8	Konstantin Chakhlevitch Celia Glass

Papers can be downloaded from

<http://www.cass.city.ac.uk/arc/actuarialreports.html>

Faculty of Actuarial Science and Insurance

Cass Business School

The support of CMI Bureau is gratefully acknowledged.

Copyright 2009 © Faculty of Actuarial Science and Insurance,
Cass Business School
106 Bunhill Row,
London EC1Y 8TZ.

ISBN 978-1-905752-22-5