Measuring the longevity risk in mortality projections

B. Wissenschaftliche Mitteilungen

N. Brouhns, M. Denuit, J. K. Vermunt, Louvain-la-Neuve, Tilburg

Measuring the Longevity Risk in Mortality Projections

1 Introduction and Motivation

As demonstrated in Benjamin & Soliman (1993), McDonald (1997) and McDonald et al. (1998), mortality at adult and old ages reveal decreasing annual death probabilities. These changes clearly affect pricing and reserving for life annuities, as stressed e.g. by Marocco & Pitacco (1998) and Olivieri (2001). The calculation of expected present values requires thus an appropriate mortality projection to avoid underestimation of future costs.

Projections are extensions of recent trends as far as they can be perceived from mortality statistics. Lee & Carter (1992) proposed a simple model for describing the secular change in mortality as a function of a single time index. The method describes the log of a time series of age-specific death rates as the sum of an age-specific component that is independent of time and another component that is the product of a time-varying parameter reflecting the general level of mortality, and an age-specific component that represents how rapidly or slowly mortality at each age varies when the general level of mortality changes. This model is fit to historical data. The resulting estimate of the time-varying parameter is then modeled and forecast as a stochastic time series using standard Box-Jenkins methods. From this forecast of the general level of mortality, the actual age-specific rates are derived using the estimated age effects. Recently, Brouhns, Denuit & Vermunt (2002) resorted to a Poisson log-bilinear regression model to build projected lifetables. Their approach, inspired from a comment made by Alho (2000) on Lee (2000), purposed to avoid some drawbacks inherent to the Lee & Carter (1992) original methodology.

The main statistical tool of Lee & Carter (1992) is least-squares estimation via singular value decomposition of the matrix of the log age-specific observed forces of mortality. This implicitly means that the errors are assumed to be homoskedastic, which is quite unrealistic: the logarithm of the observed force of mortality is much more variable at older ages than at younger ages because of the much smaller absolute number of deaths at older ages. Moreover, the required data have to fill a rectangular matrix because of singular value decomposition; this may pose a problem when the format of the available data has been modified in the past (the actuary has then first to complete the data using different techniques

which may bias the results). As we will see in Section 3, the method used by BROUHNS ET AL. (2002) avoid these drawbacks.

Of course, the projection of the mortality itself is affected by uncertainty. The effects of uncertainty coming from projections, in terms of the risk borne by the insurer, are investigated. Such an analysis is particularly important to decide upon the reinsurance needed. In BROUHNS ET AL. (2002), confidence intervals (for annuities and life expectancies) were obtained by ignoring all the errors except those in forecasting the mortality index. According to Appendix B of LEE & CARTER (1992), these errors dominate the others for annuities and expected remaining lifetimes. Because of the importance of appropriate measures of uncertainty in an actuarial context, the present paper aims to derive confidence intervals taking into account all the sources of variability. The nonlinear nature of the quantities of interest makes an analytical approach not tractable and we therefore resort to Monte-Carlo simulation (or parametric bootstrap).

Let us now describe the content of this paper. Section 2 describes the notation and assumptions adopted throughout this paper. The data used for the numerical illustrations are also presented there. Section 3 recalls the basic features of the projection model proposed by BROUHNS ET AL. (2002). The simulation method to derive confidence intervals for the quantities of interest is described there. Section 4 illustrates the approach on the mortality statistics presented in Section 2. Section 5 examines the distribution of the estimator of the net single life annuity premium and purposes to determine the safety loading with the help of a quantile of this distribution. Section 6 aims to evaluate the ruin probability relating to a portfolio of life annuities. The final Section 7 concludes. Appendices gather technical aspects.

2   Notation, assumption and data

2.1   Notation

We analyze the changes in mortality as a function of both age x and time t. Although age and time are theoretically free to vary in the half-positive real line, we work here with integer x and t. Henceforth, $\mu_x(t)$ will denote the force of mortality at age x during calendar year t. Similarly, $q_x(t)$ is the one-year death probability at age x in year t and the corresponding survival probability is $p_x(t) = 1 - q_x(t)$. We denote as D_{xt} the number of deaths recorded at age x during year t, from an exposure-to-risk ETR_{xt} (that is, ETR_{xt} is the number of person years from which D_{xt} occurred).

2.2   Piecewise constant forces of mortality

In this paper, we assume that the age-specific mortality rates are constant within bands of time and age, but allowed to vary from one band to the next. Specifically, given any integer age x and calendar year t, it is supposed that

$$\mu_{x+\xi}(t+\tau) = \mu_x(t) \quad \text{for any } 0 \leq \xi, \tau < 1. \tag{2.1}$$

This is best illustrated with the aid of a coordinate system that has calendar time as abscissa and age as coordinate. Such a representation is called a Lexis diagram after the German demographer who introduced it. Both time scales are divided into yearly bands, which partition the Lexis plane into square segments. Model (2.1) assumes that the mortality rate is constant within each square, but allows it to vary between squares.

2.3   Data

The model presented in this paper is fitted to the matrix of Dutch death rates, from year 1950 to 2000 and for ages 60 to 98. The data relate to the entire Dutch population and have been gathered by the *Centraal Bureau voor de Statistiek* (CBS). They are available from the CBS StatLine system (http://statline.cbs.nl). The observed number of deaths d_{xt}, is given by age and year, where age is year of death minus year of birth.

The raw estimate $\widehat{\mu}_x(t)$ of the force of mortality $\mu_x(t)$ is given by the ratio of the observed number of deaths d_{xt} for age x and year t to the "central exposed to risk"(henceforth denoted as ETR_{xt}), that is

$$\hat{\mu}_x(t) = \frac{d_{xt}}{ETR_{xt}} \quad \text{where } ETR_{xt} = \frac{l_{x,t} + l_{x+1,t+1}}{2}.$$

The one-year death probabilities are then estimated under (2.1) as

$$\hat{q}_x(t) = 1 - \hat{p}_x(t) = 1 - \exp\{-\hat{\mu}_x(t)\}.$$

2.4   Quantities of interest

Life expectancies are often used by demographers to measure the evolution of mortality. Specifically, $e_x(t)$ is the average number of years an x-aged individual

in year t will survive. We thus expect that this person will die in year $t + e_x(t)$ at age $x + e_x(t)$. The formula giving $e_x(t)$ under (2.1) is

$$
\begin{aligned}
e_x(t) &= \int_{\xi=0}^{1} \exp\left(-\mu_x(t)\xi\right)d\xi \\
&+ \sum_{k\geq 1}\left\{\prod_{j=0}^{k-1}\exp\left(-\mu_{x+j}(t+j)\right)\right\} \\
&\cdot \int_{\xi=k}^{k+1} \exp\left(-\mu_{x+k}(t+k)(\xi-k)\right)d\xi \\
&= \frac{1 - \exp\left(-\mu_x(t)\right)}{\mu_x(t)} \\
&+ \sum_{k\geq 1}\left\{\prod_{j=0}^{k-1}\exp\left(-\mu_{x+j}(t+j)\right)\right\} \\
&\cdot \frac{1 - \exp\left(-\mu_{x+k}(t+k)\right)}{\mu_{x+k}(t+k)}.
\end{aligned}
\tag{2.2}
$$

The actual computation of $e_x(t)$ requires the knowledge of $\mu_\xi(\tau)$ for ξ ranging from x until the ultimate age (ω, say) and for τ ranging from t to $t + \omega - x$. Of course, these mortality rates cannot be estimated at time t (since we do not have data at our disposal) and thus must be extrapolated from the past. We describe in Section 3 how this can be done in practice.

As actuaries, we are more interested in the price of an immediate life annuity sold to an individual aged x in year t, given by

$$
a_x(t) = \sum_{k\geq 0}\left\{\prod_{j=0}^{k} p_{\dot{x}+j}(t+j)\right\}v^{k+1}
\tag{2.3}
$$

where v is the yearly discount factor. We will see that mortality projections are particularly important to compute the premiums relating to such a contract.

3   Mortality projection method

3.1   Poisson modelling and piecewise constant mortality rates

The assumption (2.1) is compatible with Poisson modelling for death numbers. Indeed, let us focus on a particular couple age x - year t. We observe D_{xt} deaths

among N_{xt} individuals aged x on January 1. We assume that the remaining lifetimes of these individuals are independent and identically distributed.

To each of the N_{xt} individuals, we associate a variable δ_i indicating whether the person dies or not, i.e.

$$\delta_i = \begin{cases} 1 & \text{if person } i \text{ dies at age } x \\ 0 & \text{otherwise,} \end{cases}$$

$i = 1, 2, \ldots, N_{xt}$. We also define τ_i as the time lived by individual i. So $\tau_i = 1$ if individual i reaches age $x + 1$ (and $\delta_i = 0$) and $\tau_i < 1$ if individual i dies at age x (and $\delta_i = 1$). We assume that we have at our disposal observations (δ_i, τ_i) for each of the N_{xt} individuals.

Under the assumption (2.1), the contribution of individual i to the likelihood writes

- if he survives

$$p_x(t) = \exp(-\mu_x(t)) \, ;$$

- if he dies

$$\tau_i p_x(t) \mu_{x+\tau_i}(t + \tau_i) = \exp(-\tau_i \mu_x(t)) \mu_x(t) \, .$$

Therefore, the contribution of individual i can be cast into

$$\exp(-\tau_i \mu_x(t)) \{\mu_x(t)\}^{\delta_i} \, .$$

Invoking the mutual independence of the remaining lifetimes yields the likelihood

$$L(\mu_x(t)) = \prod_{i=1}^{N_{xt}} \exp(-\tau_i \mu_x(t)) \{\mu_x(t)\}^{\delta_i} \tag{3.1}$$

$$= \exp(-\mu_x(t) \tau_\bullet) \{\mu_x(t)\}^{\delta_\bullet}$$

where

$$\tau_\bullet = \sum_{i=1}^{N_{xt}} \tau_i \quad \text{and} \quad \delta_\bullet = \sum_{i=1}^{N_{xt}} \delta_i \, .$$

Clearly, $\delta_\bullet$ is the total number of deaths recorded for age x during year t, i.e. $\delta_\bullet = D_{xt}$, and $\tau_\bullet$ is the total exposure-to-risk (in person-years), i.e. $\tau_\bullet = \text{ETR}_{xt}$. The likelihood (3.1) is proportional to the Poisson likelihood, i.e. the one obtained under the assumption $D_{xt} \sim \text{Poisson}(\text{ETR}_{xt} \mu_x(t))$. Therefore, provided we resort

to the maximum likelihood estimation procedure, it is equivalent to work on the basis of the "true" likelihood (3.1) or on the Poisson likelihood, once the assumption (2.1) has been made.

BRILLINGER (1986) showed that under reasonable assumptions about the processes governing births and deaths, the Poisson distribution is a good candidate to model the numbers of deaths at different ages.

3.2 Poisson log-bilinear model

Under the assumption (2.1), we have seen in the latter subsection that the Poisson assumption appears to be plausible. Following BROUHNS ET AL. (2002), we thus consider that

$$D_{xt} \sim \text{Poisson}\Big(ETR_{xt}\ \mu_x(t)\Big) \quad \text{with } \mu_x(t) = \exp\left(\alpha_x + \beta_x \kappa_t\right) \qquad (3.2)$$

where the parameters α_x, β_x and κ_t are constrained by

$$\sum_t \kappa_t = 0 \quad \text{and} \quad \sum_x \beta_x = 1 \qquad (3.3)$$

ensuring model identification.

The force of mortality is thus assumed to have the log-bilinear form $\ln \mu_x(t) = \alpha_x + \beta_x \kappa_t$. Moreover, the expected number of deaths is given by $F_{xt} = ETR_{xt} \exp(\alpha_x + \beta_x \kappa_t)$. The meaning of the α_x, β_x, and κ_t parameters is essentially the same as in the classical Lee-Carter model, that is,

$\exp \alpha_x$: is the general shape across age of the mortality schedule or, more precisely, the geometric mean of $\mu_x(t)$ in the observation period;

κ_t: represents the time trend;

β_x: indicates the sensitivity of the logarithm of the force of mortality at age x to variations in the parameter κ_t. The shape of the β_x profile tells which rates decline rapidly and which slowly over time in response of change in κ_t.

3.3 Estimation of the parameters

We estimate the parameters α_x, β_x and κ_t by maximizing the log-likelihood based on model (3.2), which is given by

$$L(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\kappa}) = \sum_{x,t} \left\{ D_{xt}(\alpha_x + \beta_x \kappa_t) - ETR_{xt} \exp(\alpha_x + \beta_x \kappa_t) \right\} + \text{const.}$$

Because of the presence of the bilinear term $\beta_x \kappa_t$, it is not possible to estimate the proposed model with commercial statistical packages that implement Poisson regression. However, the LEM program (VERMUNT, 1997a, 1997b) can be used for this purpose. In Appendix A, we give the quite simple LEM input files that we used for our analyses.

The algorithm implemented in LEM to solve the likelihood equations is a unidimensional or elementary Newton method. GOODMAN (1979) was the first who proposed this iterative method for estimating log-linear models with bilinear terms. In iteration step $\nu + 1$, a single set of parameters is updated fixing the other parameters at their current estimates using the following updating scheme

$$\widehat{\xi}^{(\nu+1)} = \widehat{\xi}^{(\nu)} - \frac{\partial L^{(\nu)}/\partial \xi}{\partial^2 L^{(\nu)}/\partial \xi^2}$$

where $L^{(\nu)} = L(\widehat{\xi}^{(\nu)})$. In our application, there are three sets of parameters; that is, the α_x, the β_x, and the κ_t terms.

The updating scheme is as follows: starting with $\widehat{\alpha}_x^{(0)} = 0$, $\widehat{\beta}_x^{(0)} = 1$, and $\widehat{\kappa}_t^{(0)} = 0$ (random values can also be used)

$$\widehat{\alpha}_x^{(\nu+1)} = \widehat{\alpha}_x^{(\nu)} - \frac{\sum_t \left(d_{xt} - \widehat{F}_{xt}^{(\nu)} \right)}{-\sum_t \widehat{F}_{xt}^{(\nu)}}, \qquad \begin{aligned} \widehat{\beta}_x^{(\nu+1)} &= \widehat{\beta}_x^{(\nu)}, \\ \widehat{\kappa}_t^{(\nu+1)} &= \widehat{\kappa}_t^{(\nu)}, \end{aligned}$$

$$\widehat{\kappa}_t^{(\nu+2)} = \widehat{\kappa}_t^{(\nu+1)} - \frac{\sum_x \left(d_{xt} - \widehat{F}_{xt}^{(\nu+1)} \right) \widehat{\beta}_x^{(\nu+1)}}{-\sum_x \widehat{F}_{xt}^{(\nu)} \left(\widehat{\beta}_x^{(\nu+1)} \right)^2}, \qquad \begin{aligned} \widehat{\alpha}_x^{(\nu+2)} &= \widehat{\alpha}_x^{(\nu+1)}, \\ \widehat{\beta}_x^{(\nu+2)} &= \widehat{\beta}_x^{(\nu+1)}, \end{aligned}$$

$$\widehat{\beta}_x^{(\nu+3)} = \widehat{\beta}_x^{(\nu+2)} - \frac{\sum_t \left(d_{xt} - \widehat{F}_{xt}^{(\nu+2)} \right) \widehat{\kappa}_t^{(\nu+2)}}{-\sum_t \widehat{F}_{xt}^{(\nu+2)} \left(\widehat{\kappa}_t^{(\nu+2)} \right)^2}, \qquad \begin{aligned} \widehat{\alpha}_x^{(\nu+3)} &= \widehat{\alpha}_x^{(\nu+2)}, \\ \widehat{\kappa}_t^{(\nu+3)} &= \widehat{\kappa}_t^{(\nu+2)}, \end{aligned}$$

where $\widehat{F}_{xt}^{(\nu)} = ETR_{xt} \exp(\widehat{\alpha}_x^{(\nu)} + \widehat{\beta}_x^{(\nu)} \widehat{\kappa}_t^{(\nu)})$ is the estimated number of deaths after iteration step ν. The criterion used to stop the procedure is a very small

increase of the log-likelihood function (the default value of LEM is 10^{-6}, but it can be recommended to set the criterion a little bit sharper, so to 10^{-10}).

After updating the κ_t parameters, we have to impose a location constraint. LEM uses the centering constraint $\sum_t \widehat{\kappa}_t = 0$, which is the same constraint as in (3.3). This constraint is specified with a design matrix, namely the *spe(...)* statement in the code given in Appendix A. After updating the β_x parameters, a scaling constraint has to be imposed. The scaling constraint used by LEM is $\widehat{\beta}_1 = 1$, which is different from (3.3). In order to obtain the parameterization in which $\sum_x \widehat{\beta}_x = 1$, one has to divide the LEM estimates for β_x by $\sum_x \widehat{\beta}_x$ and multiply the LEM estimates for κ_t by the same number.

3.4   Modelling the time-factor

As in the Lee-Carter methodology the time factor κ_t is intrinsically viewed as a stochastic process. Box-Jenkins techniques are therefore used to estimate and forecast κ_t within an ARIMA(p, d, q) times series model, which takes the general form

$$(1 - B)^d \kappa_t = \mu + \frac{\Theta_q(B)\epsilon_t}{\Phi_p(B)}$$

where

B is the delay operator, $B(\kappa_t) = \kappa_{t-1}$, $B^2(\kappa_t) = \kappa_{t-2}$, ... ;

$1 - B$ is the difference operator, $(1 - B)\kappa_t = \kappa_t - \kappa_{t-1}$, $(1 - B)^2 \kappa_t = \kappa_t - 2\kappa_{t-1} + \kappa_{t-2}$, ... ;

$\Theta_q(B)$ is the Moving Average polynomial, with coefficients $\boldsymbol{\theta} = (\theta_1, \theta_2 \ldots \theta_q)$;

$\Phi_p(B)$ is the Autoregressive polynomial, with coefficients $\boldsymbol{\phi} = (\phi_1, \phi_2 \ldots \phi_p)$;

ϵ_t is white noise with variance σ_ϵ^2.

The parameters of the models are μ, $\boldsymbol{\theta}$, $\boldsymbol{\phi}$ and σ_ϵ. The method we use to obtain estimates for the ARIMA parameters is conditional least squares. Forecasted values of time parameters will be denoted by κ_t^*.

As is discussed in the next sections, the parameter estimates of the Poisson model and the forecasts κ_t^* can be used to obtain projected age-specific mortality rates, life expectancies, and annuities single premiums. We also present a simulation-based method that can be used to take the various error sources into account.

3.5 Confidence intervals for the parameters

In forecasting, it is important to provide information on the uncertainty of the forecasted quantities. In that respect, confidence intervals are particularly useful. However, in the current application it is impossible to derive the relevant confidence intervals analytically. The reason for this is that two very different sources of uncertainty have to be combined: sampling errors in the parameters of the Poisson model and forecast errors in the projected ARIMA parameters. An additional complication is that the measures of interest – mortality rates, life expectancies, and annuities single premiums – are complicated non-linear functions of the Poisson parameters α_x, β_x, and κ_t and the ARIMA parameters μ, $\boldsymbol{\theta}$, ϕ, and σ_ε.

Because of the problems associated with analytic methods, we propose estimating confidence intervals by Monte-Carlo simulation. Our simulation procedure yields M samples of α_x, β_x, and κ_t parameters and future values of the time parameters, denoted by κ_t^*. Let the mth simulated set of these basic parameters be denoted by $\boldsymbol{\xi}^m$ and the measures of interest by ψ. Since the ψ parameters are (non-linear) functions $f(\boldsymbol{\xi})$ of the basic parameters $\boldsymbol{\xi}$, the mth set of ψ parameter can be obtained by $\psi^m = f(\boldsymbol{\xi}^m)$. In other words, our procedure yields M samples of ψ parameters which can be used to compute their confidence intervals.

The two sources of uncertainty that have to be combined are the sampling fluctuation in the α_x, β_x, and κ_t parameters and the forecast error in the κ_t^* parameters. Since we resorted to maximum likelihood to estimate the parameters of the Poisson model, we know that $(\widehat{\alpha}, \widehat{\beta}, \widehat{\kappa})$ is asymptotically multivariate normally (MVN) distributed, with mean (α, β, κ) and covariance matrix given by the inverse of the Fisher information matrix $\mathcal{I}$, whose elements equal minus the expected value of the second derivatives of the log-likelihood with respect to the parameters of interest. Appendix B shows how to obtain the information matrix and how to sample values from the MVN distribution of interest. The second source of uncertainty is captured by the estimated ARIMA standard error $\widehat{\sigma}_\varepsilon$.

Once we estimated the parameters α_x, β_x and κ_t of the Poisson model (3.2) as described in Section 3.2 as well as their variance-covariance matrix $\mathcal{I}^{-1}$ as described in Appendix B, the mth sample in the Monte Carlo simulation is obtained by the following 4 steps:

1. Generate α_x^m, β_x^m, and κ_t^m from the appropriate MVN distribution (see Appendix B for details).

2. Estimate the ARIMA model using the κ_t^m as data points. This yields a new set μ^m, $\boldsymbol{\theta}^m$, ϕ^m, and σ_ε^m of the parameters μ, $\boldsymbol{\theta}$, ϕ, and σ_ε.

3. Generate a projection of κ_t^{*m} using the ARIMA parameters. The future errors ε_t^{*m} are sampled from a univariate normal distribution with a mean of 0 and a standard deviation of σ_ε^m.

4. Compute the measures of interest ψ^m.

The first step is meant to take into account the insecurity about the Poisson parameters. The second step deals with the fact that the insecurity about the ARIMA parameters depends on the insecurity about the Poisson parameters. The third makes that the insecurity about the forecasted κ_t^* not only depends on the ARIMA standard error, but also on the insecurity of the ARIMA parameters themselves. Finally, in the computation of the relevant measures in step four, all sources of insecurity are taken into account.

4   An application to Dutch population mortality statistics

4.1   Estimation of the parameters

We apply the Poisson modelling to the Dutch data presented in the introductory section. The Poisson parameters α_x, β_x and κ_t involved in (3.2) are estimated by the procedure described in Section 3.2. Figure 4.1 plots the estimated α_x, β_x and κ_t. We can see that the $\widehat{\alpha_x}$'s represent the average mortality across time: the $\widehat{\alpha_x}$'s clearly increase in x, reflecting higher mortality at older ages, as expected. The $\widehat{\beta_x}$'s decrease with age but remain positive. The $\widehat{\kappa_t}$'s for women exhibit regular behavior decreasing from 10 to -10. This reveals the improvements of mortality at ages 60 to 98 for Dutch women during the observation period. The $\widehat{\kappa_t}$'s for men behave quite irregularly, beginning to decrease only in the seventies.

4.2   Modelling the time factor

Following the early work of LEE & CARTER (1992), we use the Box-Jenkins methodology (identification - estimation - diagnosis) to generate the appropriate ARIMA time series model for the male and female mortality indexes.
A good model for the women is ARIMA(0,1,0), which is a random walk with drift:

$$(1 - B)\kappa_t = \kappa_t - \kappa_{t-1} = \mu + \epsilon_t. \tag{4.1}$$

For the men, the situation is a bit more complicated. Looking at the data (see Figure 4.1) gives the feeling that there is a break in the series: data before year
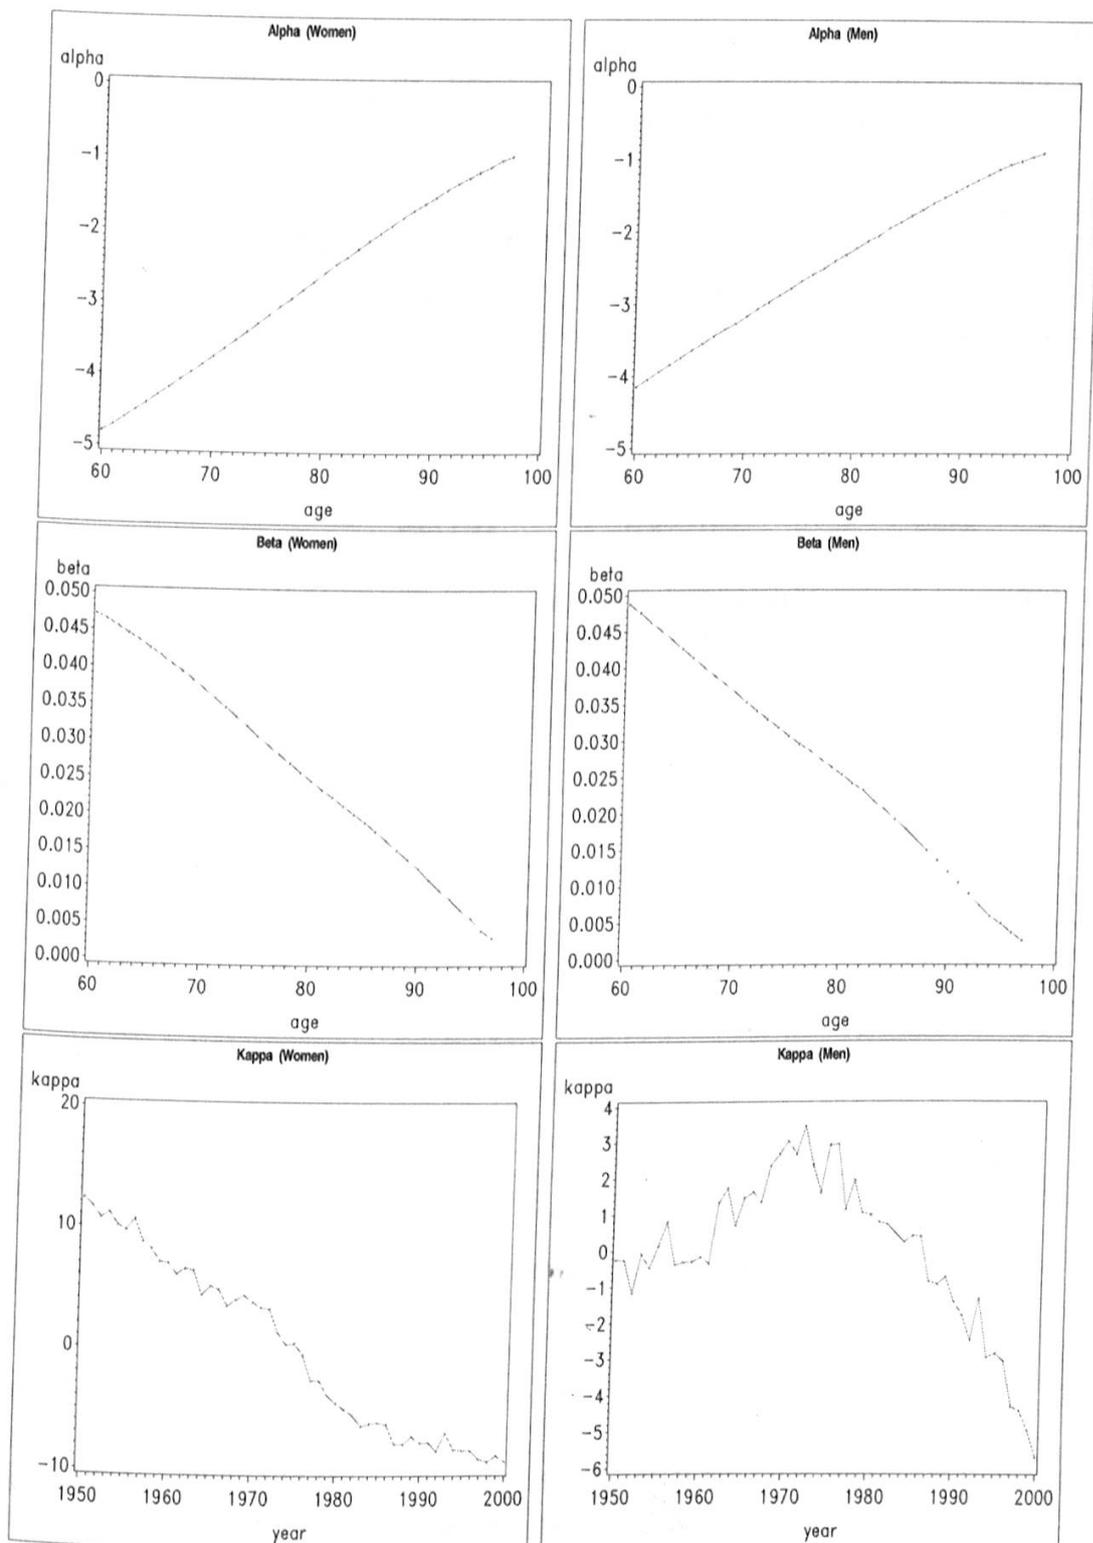
Figure 4.1: Estimations of α_x, β_x and κ_t for women (left panels) and men (right panels).

1970 behave differently from data after 1970. We thus split the series into two parts, each having its own stochastic behaviour. In the following we will use data from 1970 for projecting the κ_t for the male population. Moreover, the ARIMA(0,1,0) model (4.1) appears to be a good choice in this case as well, bringing us close to the work of LEE & CARTER (1992).

The estimated parameters for the ARIMA(0,1,0) models (4.1) are given in Table 4.1 for men and women. The sex-specific estimated values of κ_t together with the projected κ_t^* are shown with their 95% interval forecasts in Figure 4.2.

Sex	$\widehat{\mu}$	$\widehat{\sigma_\varepsilon}$
Women	−0.4293	0.8698
Men	−0.2503	0.3749

Table 4.1: Estimation of the parameters μ and σ_ϵ of the ARIMA(0,1,0) models.
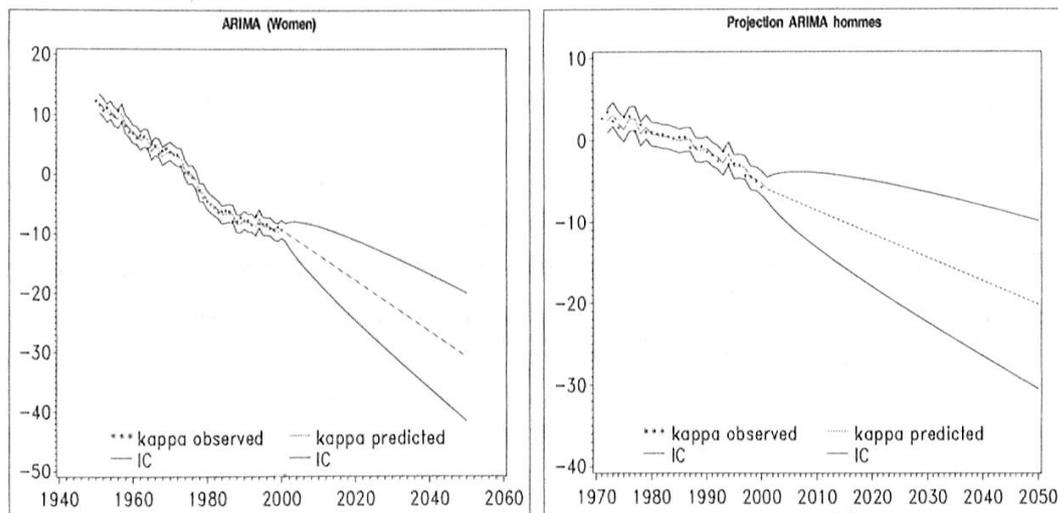

Figure 4.2: Estimated and projected values of κ_t for males and females.

4.3   Confidence intervals

A Monte Carlo simulation is then used to generate $10,000$ samples of the original parameters α_x, β_x and κ_t. On the basis of each of these 10,000 realizations, we estimated the ARIMA(0,1,0) parameters μ and σ_ϵ. The simulation thus also gives a sample of size $10,000$ of the ARIMA parameters, whose significant percentile values are given in Table 4.2. The interval $[\widehat{q_{0.05}}, \widehat{q_{0.95}}]$ is best regarded as an approximate 90% confidence interval for the quantities of interest.

	$\overline{\mu}$	$\widehat{q_{0.05}}$	$\widehat{q_{0.95}}$	$\overline{\sigma_\epsilon}$	$\widehat{q_{0.05}}$	$\widehat{q_{0.95}}$
Women	−0.4286	−0.4374	−0.4200	0.8995	0.8371	0.9631
Men	−0.2501	−0.2611	−0.2391	0.3947	0.3506	0.4398

Table 4.2: Simulation outcomes for the parameters of the ARIMA(0,1,0) models: $\overline{\mu}$ is the average over the 10,000 samples of the estimations for μ, and $\overline{\sigma_\epsilon}$ is the analogue for σ_ϵ. These values are supplemented with the 5% and 95% percentiles of the outcomes.

The last step is then to compute values for the quantities of interest. As explained in Section 3.4, each set $\alpha_x^m, \beta_x^m, \kappa_t^m$ and κ_t^{*m} of simulated $\alpha_x, \beta_x, \kappa_t$ and κ_t^* gives a realization of this quantity, so that the procedure also provides the actuary with a sample of size 10,000 on the basis of which standard errors and quantiles can be estimated. If we consider for example the evolution of the mortality rates at age 65 through years, we obtain 10,000 realizations from

$$\mu_{65}^m(t) = \exp(\alpha_{65}^m + \beta_{65}^m \kappa_t^m)$$

for $t \le 2,000$ and

$$\mu_{65}^m(t) = \exp(\alpha_{65}^m + \beta_{65}^m \kappa_t^{*m})$$

for $t \ge 2,001$. This is represented on Figure 4.3. Similarly, the evolution of the mortality rates in 2005 through ages from 60 is depicted in Figure 4.4. For each situation, the point estimates (given by the average over the 10,000 samples) are supplemented with 90% confidence bands $[\widehat{q_{0.05}}, \widehat{q_{0.95}}]$.
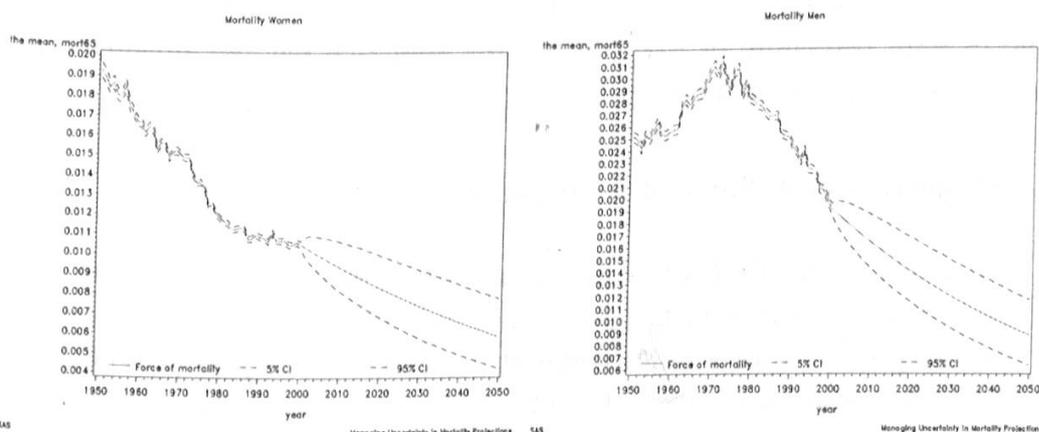

Figure 4.3: Mortality rates $\mu_{65}(t)$, $t \ge 1950$, with 90% confidence intervals, for women (left) and for men (right).
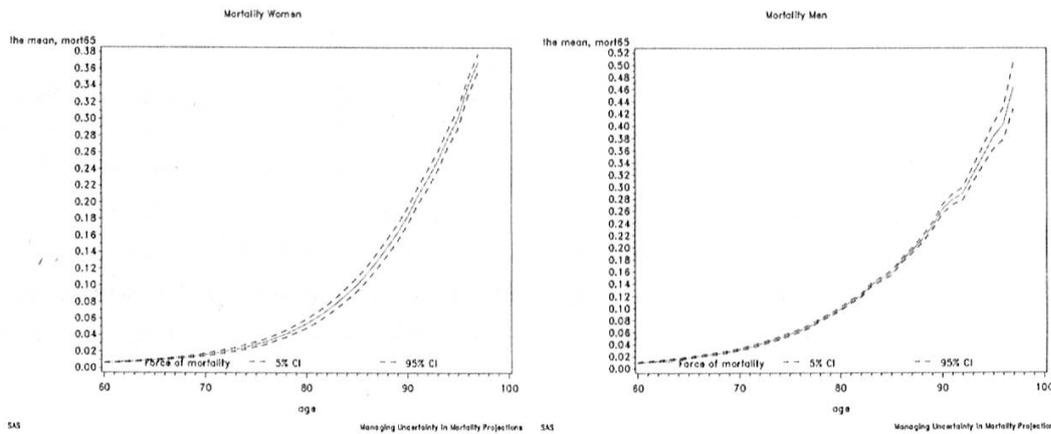
Figure 4.4: Mortality rates $\mu_x(2005)$, $x \geq 60$, with 90% confidence intervals, for women (left) and for men (right).

5 Distribution of the estimator of the life annuity net single premium

Using the projected lifetables generated in the preceding section, we deduce confidence intervals on life expectancies $e_x(t)$ and on the net present values $a_x(t)$. These quantities indeed depend on the future evolution of mortality. Specifically, having generated mortality rates $\mu_x^m(t)$, $m = 1, \ldots, 10{,}000$, we get one-year survival probabilities $p_x^m(t)$ and we can thus compute $e_x^m(t)$ and $a_x^m(t)$ according to formulas (2.2) and (2.3). Figure 5.5 displays an estimation of the density function of $\widehat{e_{65}(2000)}$ and $\widehat{a_{65}(2000)}$ for women and Figure 5.6 is the analogue for men, with $v = 1.04^{-1}$.

Once the estimation of the density of $\widehat{a_{65}(2000)}$ is available, the actuary can decide about the height of the safety loading. Indeed, the company could charge the 90 or 95th percentile of $\widehat{a_{65}(2000)}$. This approach has the advantage to offer a clear understanding of the way the safety margin is computed.

6 Projecting cash flows of a life annuity portfolio

Let us consider a portfolio of n immediate life annuities ($n = 10{,}000$ in all the numerical illustrations), sold to 65-year-old individuals at January 1 of year 2000 and providing them with a unit capital at the end of each year provided they are still alive. The random number of contracts at time t (calendar year $2000 + t$) is N_t. Having generated a sequence of $\mu_{65+t}^m(2000+t)$ as explained in the preceding sections we compute sequences of

$$p_{65+t}^m(2000 + t) = \exp\{-\mu_{65+t}^m(2000 + t)\}$$

Life expectancies Women

Managing Uncertainty in Mortality Projections

Annuities Women

Managing Uncertainty in Mortality Projections

Figure 5.5: Life expectancies $e_{65}\widehat{(2000)}$ (top) and annuities $a_{65}\widehat{(2000)}$ (bottom) distributions for women.
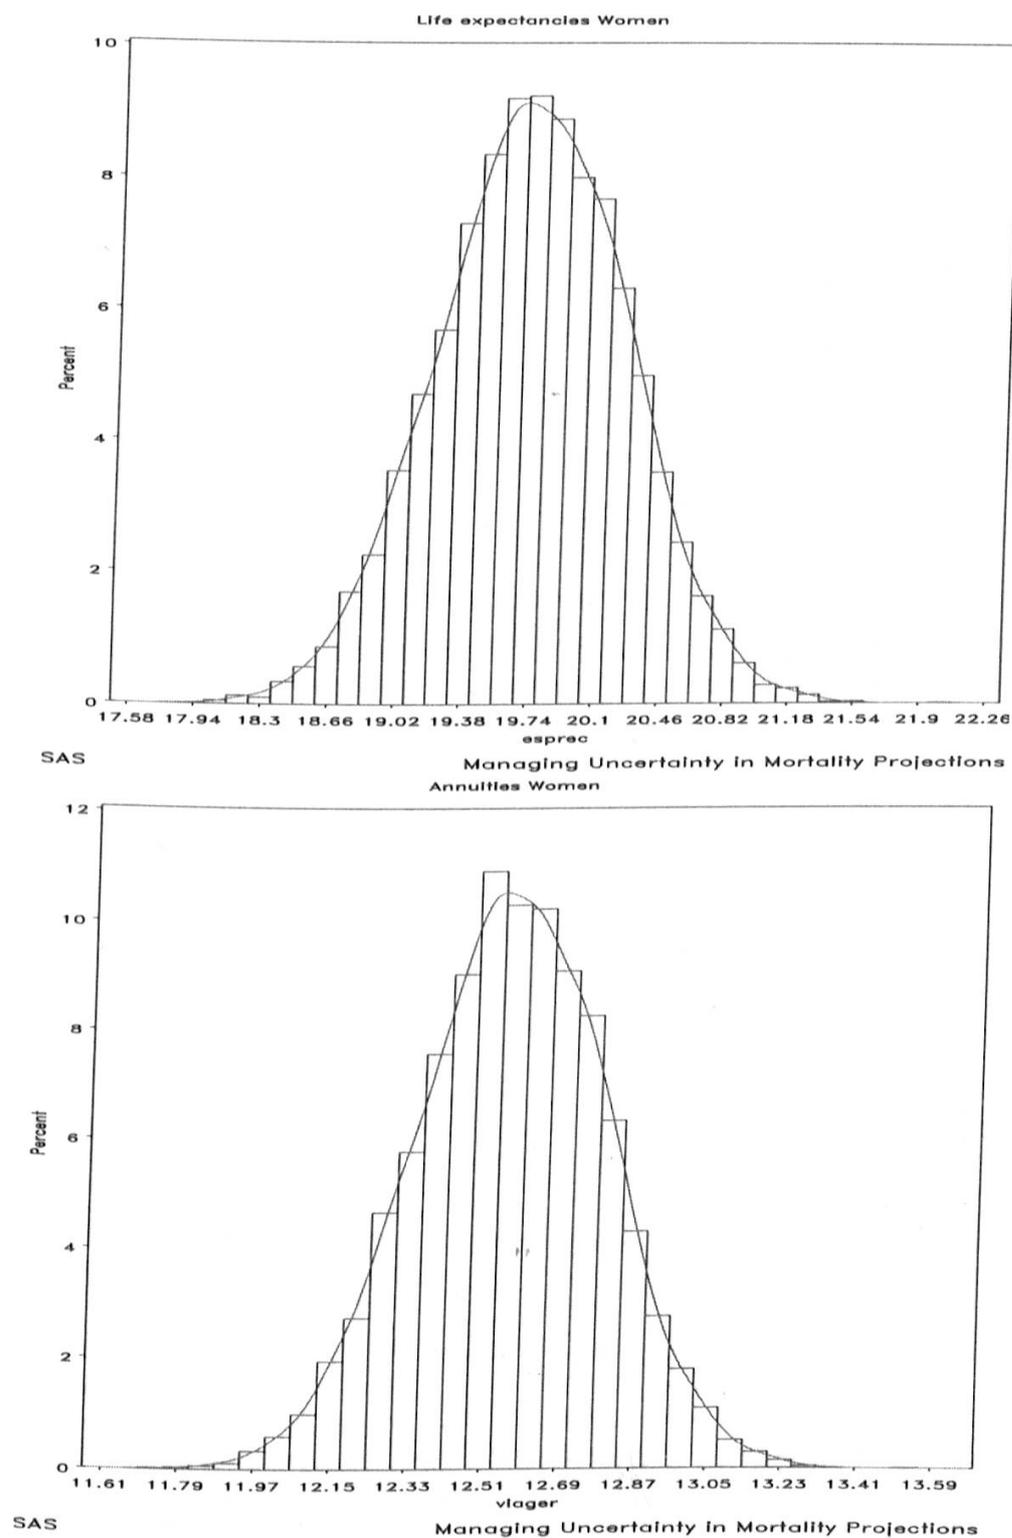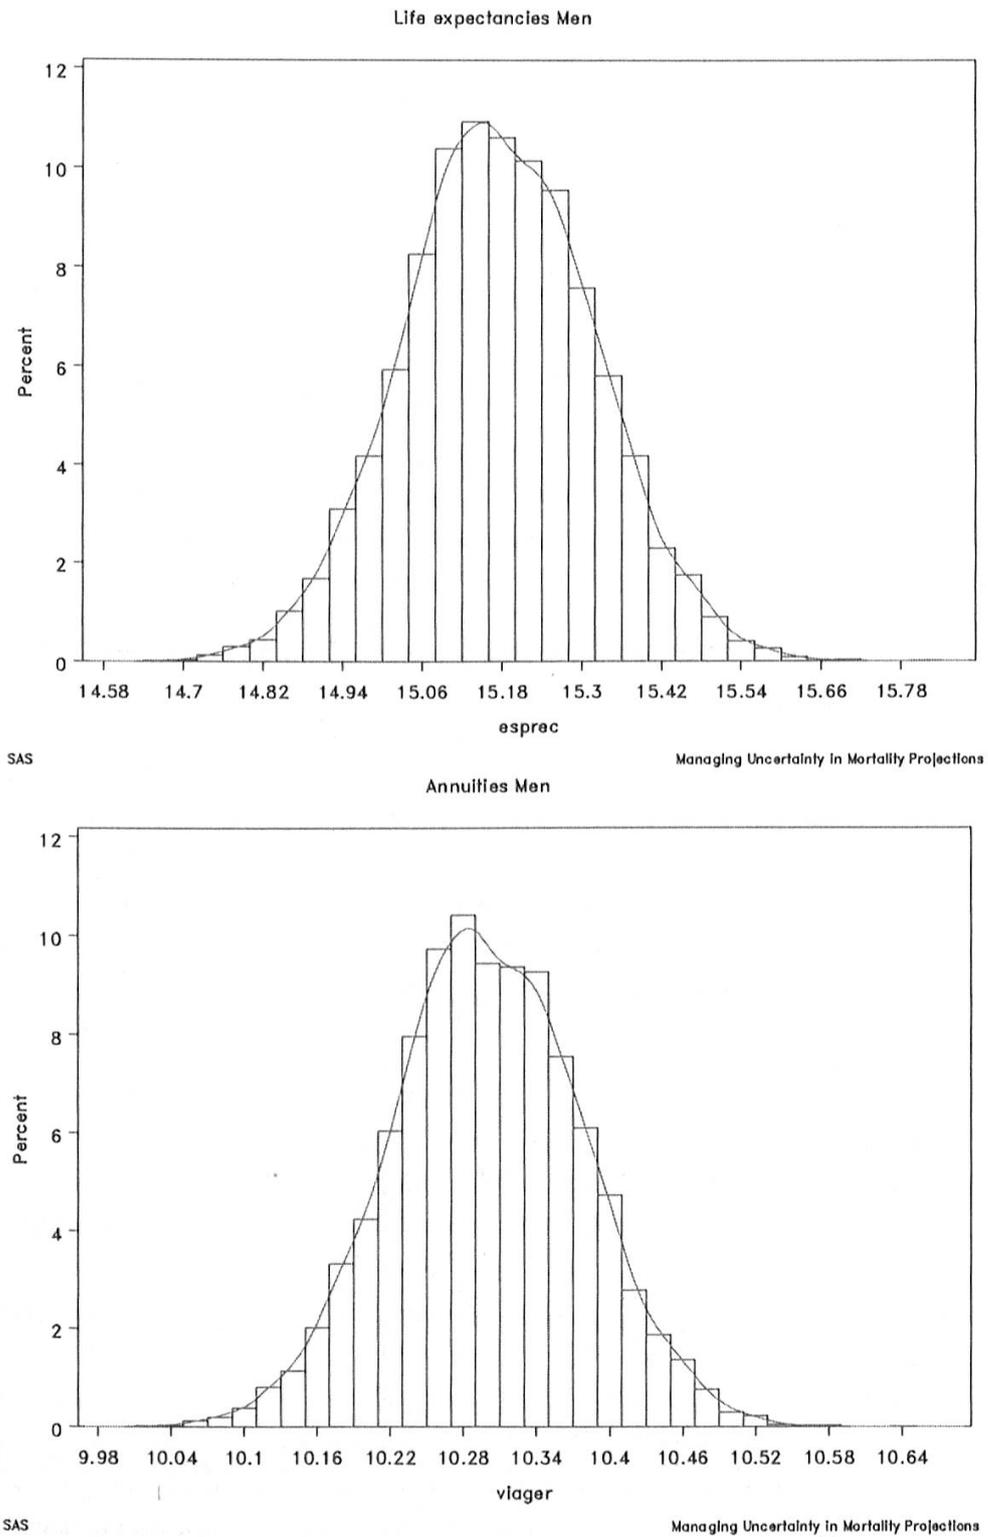
Life expectancies Men

Annuities Men

Figure 5.6: Life expectancies $\widehat{e_{65}(2000)}$ (top) and annuities $\widehat{a_{65}(2000)}$ (bottom) distributions for men.

and

$$q_{65+t}^m(2000+t) = 1 - \exp\{-\mu_{65+t}^m(2000+t)\}, \quad m = 1, \ldots, 10,000.$$

There are different ways to simulate the number of survivors at the different ages:

1. A first possibility consists in generating at once all the years of death for the 10,000 annuitants. Denoting as $T_{65}(2000)$ the remaining lifetime of an insured aged 65 in year 2000, the associated survival function is

$$\Pr[T_{65}(2000) > \xi] = \begin{cases} \exp\left(-\mu_{65}(2000)\right) & \text{if } \xi \leq 1 \\ \exp\left(-(\xi - [\xi])\mu_{65+[\xi]}(2000 + [\xi])\right) \\ \qquad \prod_{k=0}^{[\xi]-1} \exp\left(-\mu_{65+k}(2000+k)\right) & \text{if } \xi > 1 \end{cases}$$

where $[\xi]$ is the integer part of the positive real number ξ. To simulate a realization of $T_{65}(2000)$, it then suffices to invert the latter survival function and to evaluate it in a random unit uniform number. More precisely, having simulated u from the unit uniform distribution, we look for the index j such that

$$\prod_{k=0}^{j} \exp\left(-\mu_{65+k}(2000+k)\right) \geq u \geq \prod_{k=0}^{j+1} \exp\left(-\mu_{65+k}(2000+k)\right).$$

The annuitant then dies at age $65 + j + 1$. If

$$u \geq \exp\left(-\mu_{65}(2000)\right)$$

then the annuitant dies at age 65 and the insurer does not have to pay anything.

2. A second approach consists in generating the annual number of deaths from the binomial distribution with parameters N_t and $q_{65+t}(2000+t)$.

3. A third possibility is to resort to the Poisson approximation for the binomial distribution.

In practice, these three approaches provide very similar results. Here, we continue with the Poisson modelling and proceed as follows. We simulate the future evolution of this portfolio as follows. Starting from $N_0 = n$, we first calculate the exposure as

$$ETR_{65,2000} = -N_0 \frac{q_{65}(2000)}{\ln p_{65}(2000)}.$$

Then we simulate the number of deaths at age 65 in year 2000 as

$$D_{65}(2000) \sim \text{Poisson}(ETR_{65,2000}\mu_{65}(2000))$$

and the number of remaining policies in year 2001 as

$$N_1 = N_0 - D_{65}(2000)$$

The company pays an amount N_1 and gets returns on the reserve.
Proceeding iteratively for $t = 1, 2, \ldots$, we simulate until the cohort totally vanished according to the following equations:

$$ETR_{65+t,2000+t} = -N_t \frac{q_{65+t}(2000 + t)}{\ln p_{65+t}(2000 + t)}$$

$$D_{65+t}(2000 + t) \sim \text{Poisson}(ETR_{65+t,2000+t}\mu_{65+t}(2000 + t))$$

and

$$N_{t+1} = N_t - D_{65+t}(2000 + t).$$

is the amount to be paid by the company at the end of year $2000 + t$.
It is worth mentioning that the simulating procedure does not guarantee that the number of deaths $D_{65+t}(2000 + t)$ is smaller than the number of remaining policies N_t.
Let us now project the future cash flows corresponding to this portfolio. At time 0 the company gets $na_{65}(2000)$. Then we observe the extinction of the cohort and compute the cashflows and the evolution of the reserves each year. For the reserves, we have taken the same yearly interest rate as the one used in the calculation of the annuity, namely $i = 4\%$, which corresponds to a quite pessimistic view. In Table 6.1, four methods for computing the net single premiums have been compared, namely

1. the transversal vision, based on observed data from 1998 to 2000: in this case, the mortality rates are estimated on the basis of the statistics related to the years 1998–2000 and are used directly (without projection) to price the life annuity contracts;

2. the longitudinal vision, pure premium: the mortality rates are projected according to the method described in Section 3 but no safety loading is added to the pure premium so obtained;

3. The longitudinal vision, 90-th percentile value: the mortality rates are projected and a safety loading is added by charging the 90th percentile of $\widehat{a_{65}(2000)}$, as discussed in Section 5;

4.   The longitudinal vision, 95-th percentile value: the mortality rates are projected and a safety loading is added by charging the 95th percentile of $\widehat{a_{65}(2000)}$, as discussed in Section 5.

The different columns of Table 6.1 give the following results:

1.   the net single premium of the life annuity, computed according to the 4 strategies described above;

2.   the global probability of ruin (in %) at total extinction of the cohort, that is, the probability that the premium income got in 2000 does not suffice to fund all the promised payments, if the interest rate obtained on the reserves is equal to 4% (which is a quite pessimistic scenario);

3.   The mean time to ruin (in years), that is, the average number of years elapsed before ruin, given that ruin occurs;

4.   The mean severity of ruin (the year the ruin occurs), which is the deficit the year the company runs out of funds;

5.   the mean number of remaining contracts when ruin occurs;

6.   the interest rate on the reserves needed to ensure that the global probability of ruin is below 1%.

Let us comment the results of Table 6.1. First and foremost, they enlighten the importance of mortality projections: ratemaking on the basis of transversal data results in negative cashflows in almost 100% of the cases simulated (99.84% for women and 97.94% for men). Moreover, it appears to be necessary to include a safety loading since charging the pure premium leads to ruin in about 50% of the cases (as expected). It is interesting to notice the different results obtained according to the gender of the policyholders. When the percentile premium calculation principle is used, the ruin probability is higher for men than for women, ruin occurs on average more rapidly for men than for women (provided ruin indeed occurs) but the deficit is much higher for women than for men. Similarly, the number of pending policies when ruin occurs (that is, the victims of the bankruptcy) is higher for women. The last column shows that, as it is well known by practitioners, it is possible to counteract the longevity risk by sufficiently high financial returns on the reserves.

Women						
Premium Principle	Annuity	Global ruin probability (in %)	Mean time to ruin (in years)	Mean severity of ruin	Mean nb of remaining contracts	i (in %)
Transv.	11.82	99.84	22.8	−194	407	5.2
Long.	12.57	55.58	28.7	−90	190	4.6
Long. 90%	12.85	17.37	30.5	−65	146	4.3
Long. 95%	12.93	11.00	31.3	−62	131	4.3
Men						
Premium Principle	Annuity	Global ruin probability (in %)	Mean time to ruin (in years)	Mean severity of ruin	Mean nb of remaining contracts	i (in %)
Transv.	9.97	97.94	22.0	−100	214	4.9
Long.	10.30	50.66	25.8	−50	109	4.5
Long. 90%	10.40	26.90	26.7	−41	90	4.4
Long. 95%	10.43	22.12	27.0	−38	85	4.3

Table 6.1: Risk measures for the annuity portfolio.

7   Conclusion

To the knowledge of the authors, the present paper offers the first attempt to quantify the longevity risk, that is, the variability of the life annuity premiums computed on the basis of projected mortality rates. Since in the log-bilinear Poisson regression approach, this amounts to combine different sources of sampling fluctuations (namely, the variability of the estimations $\widehat{\alpha_x}$, $\widehat{\beta_x}$ and $\widehat{\kappa_t}$ together with the prediction errors of the κ_t^*), an analytical approach turns out to be virtually impossible. Therefore, we have opted for a Monte-Carlo approach. The simulation strategy adopted in this paper is fully parametric (in the sense that the confidence intervals are obtained under the hypothesis that the model (3.2) is correct) and based on large sample properties of the ML estimators. Specifically, we have generated M samples from the multivariate normal distribution with mean vector $(\widehat{\alpha}, \widehat{\beta}, \widehat{\kappa})^t$ and covariance matrix $\widehat{\mathcal{I}}^{-1}$.

The parametric Monte-Carlo method used in this paper is an interesting first attempt to quantify the uncertainty of the future mortality. Nevertheless, its major drawback is that it is based on a normal approximation which may not be very accurate. Since the second step of the simulation consists in estimating an ARIMA model using generated data that are multivariate normal with known mean and covariance matrix, the whole simulation is fully determined by the estimated α_x's, β_x's and κ_t's together with the information matrix. Hence, the real structure of the data cannot be fully captured.

To avoid these drawbacks, we can contemplate other possibilities for dealing with the insecurity of the parameters of the Poisson model. Two of these

are semiparametric and nonparametric bootstrapping. Both procedures involve generating M new tables of observed numbers of deaths and reestimating the Poisson model with each of these generated data matrices. This yields the M sets of α_x, β_x, and κ_t parameters that are needed in the subsequent steps. The two bootstrapping methods differ in the manner in which the M new data sets are generated. A straightforward manner to implement the semiparametric bootstrap is to generate observed numbers of deaths from the Poisson distribution defined by estimates of the Poisson parameters and the observed exposures times. Another, more complicated, implementation involves generating cohort survival tables using the estimated Poisson rates, where the risk population is adapted depending on the numbers of deaths in the previous year. In the nonparametric bootstrap, the M new data matrices are obtained by means of sampling with replacement from the original data matrix.

In a forthcoming paper, we will compare the fully parametric approach worked out in the present article to semi- and nonparametric bootstrap, to check whether the confidence intervals on the life annuitiy premiums derived in this paper are not artificially too small.

To end with, let us mention that the study of the variability of the amounts of premium, and of the corresponding ruin probabilities, are of prime importance for deciding upon the level of reinsurance needed.

Appendices

A   LEM input files

This is the LEM input file that estimates the Poisson parameters α_x, β_x and κ_t involved in (3.2):

```
man 2
dim 39 51
lab X T
mod {wei(XT), X, spe(T,1a,X,b)}
dat deaths.dat
sta wei(XT) exposures.dat
```

The command "man" indicates the number of (manifest) variables, in this case 2 (age and calendar time). With "dim", one specifies the number of levels of the variables. For females, we had 39 age groups and 51 time points. The command "lab" is used to specify variable labels. The "mod" statement is used to specify the three relevant model terms: the exposures [wei(XT)], the age effect [X], and the bilinear term [spe(T,1a,X,b)]. It is assumed that the files "deaths.dat" and "exposures.dat" contain the tables with observed counts d_{xt} and exposure ETR_{xt}. The commands "dat" and "sta" are used to specify these data files.

B   Fisher Information matrix

In order to simplify notation in the description of the elements of the Fisher information matrix, we write the expected number of deaths at age x in year t for an exposure to risk ETR_{xt} in the Poisson model in a slightly different form; that is,

$$F_{xt} = ETR_{xt} \exp\left[\left(\sum_{y=x_{\min}}^{x_{\max}} a_{xy}\alpha_y\right) + \left(\sum_{y=x_{\min}}^{x_{\max}} b_{xy}\beta_y\right)\left(\sum_{r=t_{\min}}^{t_{\max}} k_{tr}\kappa_r\right)\right]$$

where $x_{\min}$, $x_{\max}$, $t_{\min}$ and $t_{\max}$ have obvious meanings. Here, a_{xy}, b_{xy} and k_{tr} denote elements of three design matrices A, B, and K, whose columns are associated with the three sets of Poisson parameters. More precisely, a_{xy} and b_{xy} equal 1 if $x = y$, and 0 otherwise. Moreover, k_{tr} equals 1 if $t = r$, -1 if $t = t_{\max}$, and 0 otherwise. Note that setting $k_{t_{\max}r} = -1$ amounts to saying that $\kappa_{t_{\max}} = -\sum_{r=t_{\min}}^{t_{\max}-1} \kappa_r$, which is needed for identification. As a result, K

contains only $t_{\max} - t_{\min}$ instead of $t_{\max} - t_{\min} + 1$ columns. For identification, we also fix $\beta_{x_{\min}}$ to 1. As was explained in the text, it is straightforward to switch from this parameterization to the Lee-Carter parameterization in which $\sum_{y=x_{\min}}^{x_{\max}} \beta_y = 1$. With this parameterization, it is much easier to derive the Fisher information matrix.

Using L as a shorthand for $L(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\kappa})$, the elements of the Fisher information matrix for the free parameters $\alpha_{x_{\min}}$ to $\alpha_{x_{\max}}$, $\beta_{x_{\min}+1}$ to $\beta_{x_{\max}}$, and $\kappa_{t_{\min}}$ to $\kappa_{(t_{\max}-t_{\min})}$ can be obtained by

$$-E\left(\frac{\partial^2 L}{\partial \alpha_y \alpha_{y'}}\right) = \sum_x \sum_t F_{xt} a_{xy} a_{xy'}$$

$$-E\left(\frac{\partial^2 L}{\partial \beta_y \beta_{y'}}\right) = \sum_x \sum_t F_{xt} (\kappa_t b_{xy})(\kappa_t b_{xy'})$$

$$-E\left(\frac{\partial^2 L}{\partial \kappa_r \kappa_{r'}}\right) = \sum_x \sum_t F_{xt} (\beta_x k_{tr})(\beta_x k_{tr'})$$

$$-E\left(\frac{\partial^2 L}{\partial \alpha_y \beta_{y'}}\right) = \sum_x \sum_t F_{xt} a_{xy} (\kappa_t b_{xy'})$$

$$-E\left(\frac{\partial^2 L}{\partial \alpha_y \kappa_r}\right) = \sum_x \sum_t F_{xt} a_{xy} (\beta_x k_{tr})$$

$$-E\left(\frac{\partial^2 L}{\partial \beta_y \kappa_r}\right) = \sum_x \sum_t F_{xt} (\kappa_t b_{xy})(\beta_x k_{tr})$$

In the first step of the our Monte Carlo simulation procedure, we generate α_x^m, β_x^m, and κ_t^m from a MVN distribution with means equal to the maximum likelihood (ML) estimates $\widehat{\alpha}_x$, $\widehat{\beta}_x$, and $\widehat{\kappa}_t$ and covariance matrix equal to $\widehat{\mathcal{I}}^{-1}$. Note that the estimated Fisher information matrix is obtained by filling in the ML estimates in the above formulas.

In practice, simulation from a MVN distribution is done as follows:

$$\boldsymbol{\xi}^m = \widehat{\boldsymbol{\xi}} + \widehat{\boldsymbol{C}} \boldsymbol{u}^m.$$

Here, $\widehat{\boldsymbol{\xi}}$ denotes the vector with ML estimates, $\boldsymbol{u}^m$ is a vector of independent standard normal deviates, and $\widehat{\boldsymbol{C}}$ is the Choleski decomposition of $\widehat{\mathcal{I}}^{-1}$.

Before going to the second step in which the ARIMA model is estimated using κ_t^m as data points, we rescale the β_x^m and κ_t^m terms so that they are in agreement with the Lee-Carter parameterization.

Acknowledgements

References

Alho, J.M. (2000). Discussion of *Lee* (2000). *North American Actuarial Journal* **4**, 91-93.

Benjamin, B., & Soliman, A.S. (1993). *Mortality on the Move.* Institute of Actuaries, Oxford.

Brillinger, D.R. (1986). The natural variability of vital rates and associated statistics. *Biometrics* **42**, 693-734.

Brouhns, N., Denuit, M. & Vermunt, J.K. (2002). A Poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics & Economics*, to appear.

Goodman, L.A. (1979). Simple models for the analysis of association in cross-classifications having ordered categories. *Journal of the American Statistical Association* **74**, 537-552.

Lee, R.D. (2000). The Lee-Carter method of forecasting mortality, with various extensions and applications. *North American Actuarial Journal* **4**, 80-93.

Lee, R.D., & Carter, L. (1992). Modelling and forecasting the time series of US mortality. *Journal of the American Statistical Association* **87**, 659-671.

McDonald, A.S. (1997) Editor. *The Second Actuarial Study of Mortality in Europe.* Groupe Consultatif des Associations d'Actuaires des Pays des Communautés Européennes, Oxford.

McDonald, A.S., Cairns, A.J.C., Gwilt, P.L., & Miller, K.A. (1998). An international comparison of recent trends in mortality. *British Actuarial Journal* **4**, 3-141.

Marocco, P., & Pitacco, E. (1998). Longevity risk and life annuity reinsurance. *Transactions of the 26th International Congress of Actuaries*, Birmingham, 453-479.

Olivieri, A. (2001). Uncertainty in mortality projections: an actuarial perspective. *Insurance: Mathematics & Economics* **29**, 231-245.

Vermunt, J.K. (1997a). *Log-linear models for event histories.* Thousand Oakes: Sage Publications.

Vermunt, J.K. (1997b). *LEM: A general program for the analysis of categorical data.* Users' manual. Tilburg, The Netherlands: Tilburg University (see "www.kub.nl/mto" under software).

Natacha Brouhns
Michel Denuit
Université Catholique de Louvain
Louvain-la-Neuve
Belgium

Jeroen K. Vermunt
Department of Methodology and Statistics
Tilburg University
LE Tilburg
The Netherlands

Abstract

Projected lifetables are used to price life annuities because they include a forecast of the future trends in mortality. However, such tables may not properly represent future mortality, originating the so-called longevity risk. The present work purposes to quantify the uncertainty inherent to mortality projections in the framework of the log-bilinear Poisson regression model of BROUHNS, DENUIT & VERMUNT (2002).

Résumé

Les tables de mortalité prospectives sont utilisées pour la tarification des rentes viagères car elles incorporent une anticipation de la mortalité future. Cependant, ces tables pourraient ne pas décrire adéquatement la mortalité future, donnant ainsi naissance au risque de longévité. Le présent travail a pour but de quantifier l'incertitude entachant les projections de mortalité dans le cadre du modèle de Poisson log-bilinéaire de BROUHNS, DENUIT & VERMUNT (2002).

Zusammenfassung

Projizierte Sterbetafeln werden eingesetzt, um den Preis von Leibrenten zu bestimmen, weil in ihnen eine Voraussage der Entwicklung der zukünftigen Sterblichkeit enthalten ist. Solche Tafeln werden in der Regel aber den Verlauf der zukünftigen Sterblichkeit nicht genau darstellen woraus das sogenannte Langlebigkeitsrisiko entsteht. Ziel des vorliegenden Artikels ist es, die Ungewissheit die solchen Sterblichkeitsprojektionen innewohnt zu quantifizieren. Dies geschieht im Rahmen des log-bilinear Poisson Regressionsmodells von BROUHNS, DENUIT & VERMUNT (2002).