Insurance: Mathematics and Economics 69 (2016) 70–81

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

journal homepage: www.elsevier.com/locate/ime

A multivariate evolutionary credibility model for mortality improvement rates

Edo Schinzinger [a,*], Michel M. Denuit [b], Marcus C. Christiansen [c]

[a] DFG Research Training Group 1100, University of Ulm, Ulm, Germany

[b] Institut de Statistique, Biostatistique et Sciences Actuarielles, Université Catholique de Louvain, Louvain-la-Neuve, Belgium

[c] Department of Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, United Kingdom

ARTICLE INFO

Article history:
Received June 2015
Received in revised form
April 2016
Accepted 13 April 2016
Available online 27 April 2016

Keywords:
Mortality projection
Predictive distribution
Multi-population modeling
ARMA process
Lee-Carter model

ABSTRACT

The present paper proposes an evolutionary credibility model that describes the joint dynamics of mortality through time in several populations. Instead of modeling the mortality rate levels, the time series of population-specific mortality rate changes, or mortality improvement rates are considered and expressed in terms of correlated time factors, up to an error term. Dynamic random effects ensure the necessary smoothing across time, as well as the learning effect. They also serve to stabilize successive mortality projection outputs, avoiding dramatic changes from one year to the next. Statistical inference is based on maximum likelihood, properly recognizing the random, hidden nature of underlying time factors. Empirical illustrations demonstrate the practical interest of the approach proposed in the present paper.

1. Introduction

Mortality forecasts are used in a wide variety of fields. Let us mention health policy making, pharmaceutical research, social security, retirement fund planning and life insurance, to name just a few.

Following the elegant approach to mortality forecasting pioneered by Lee and Carter (1992) many projection models decompose the death rates (on the logarithmic scale) or the one-year death probabilities (on the logit scale) into a linear combination of a limited number of time factors. See, e.g., Hunt and Blake (2014). In a first step, regression techniques are used to extract the time factors from the available mortality data. In a second step, the time factors are intrinsically viewed as forming a time series to be projected to the future. The actual age-specific death rates are then derived from this forecast using the estimated age effects. This in turn yields projected life expectancies.

In the first step of the two-step model calibration procedure, the random nature of the unobservable time factor is disregarded, and this may bias the analysis. As possible incoherence may arise from this two-step procedure, Czado et al. (2005) integrated both steps into a Bayesian version of the model developed by Lee and Carter (1992) in order to avoid this deficiency. After Czado et al. (2005) and Pedroza (2006) formulated the Lee–Carter method as a state-space model, using Gaussian error terms and a random walk with drift for the mortality index. See also Girosi and King (2008), Kogure et al. (2009), Kogure and Kurachi (2010) and Li (2014) for related works. However, the practical implementation of Bayesian methods often requires computer-intensive Markov Chain Monte Carlo (MCMC) simulations. This is why we propose in this paper a simple credibility model ensuring robustness over time while keeping the computational issues relatively easy and allowing for the flexibility of time series modeling. It is worth stressing that the time factor is here treated as such, and not as a parameter to be estimated from past mortality statistics using regression techniques before entering time series models. In this way, we recognize the hidden nature of the time factor and its intrinsic randomness.

Whereas most mortality studies consider both genders separately, the model proposed in this paper easily combines male and female mortality statistics. This is particularly useful in practice when both genders are usually involved. In insurance applications, for instance, separate analyses could lead to miss this strong dependence pattern, which considerably reduces possible diversification effects between male and female policyholders inside the portfolio. In demographic projections, combining male and female data is necessary to ensure consistency in gender-specific mortality forecast. This problem has been considered by several authors

* Corresponding author.
E-mail address: edo.schinzinger@uni-ulm.de (E. Schinzinger).

http://dx.doi.org/10.1016/j.insmatheco.2016.04.004
0167-6687/

in the literature. Let us mention Carter and Lee (1992) who fitted the Lee and Carter (1992) model to male and female populations separately and then measured the dependence between the two gender-specific time factors. These authors considered three models for the pair of time factors: a bivariate random walk with drift, a single time factor common to both genders and a co-integrated process where the male index follows a random walk with drift and there exists a stationary linear combination of both time factors.

More generally, the credibility model proposed in this paper is able to pool several populations to produce mortality forecasts for a group of countries. In such a context, Yang and Wang (2013) assumed that the time factors followed a vector error correction model. See also Zhou et al. (2013). Other models incorporate a common factor for the combined population as a whole, as well as additional factors for each sub-population. The common factor describes the main long-term trend in mortality change while the additional factors depict the short-term discrepancy from the main trend inside each sub-population. See Li and Lee (2005) who proposed the augmented common factor model generalized by Li (2013) to several factors. The model structure proposed in Delwarde et al. (2006), by Debón et al. (2011), and by Russolillo et al. (2011) only includes a single, common time factor. As argued in Carter and Lee (1992), this simple arrangement may enforce greater consistency and is a parsimonious way to model both populations jointly. However, it also implies that the death rates of the two populations are perfectly associated, an assumption with far-reaching consequences in risk management.

Our paper innovates in that the new multi-population mortality projection model we propose is based on mortality improvement rates instead of levels. Recently, several authors suggested to target improvement rates to forecast future mortality, instead of the death rates. While the time dependence structure of death rate models is dominated by the continuing downward trend, the improvement rates are already trend adjusted. See, e.g., Mitchell et al. (2013) or Börger and Aleksic (2014). The model developed in this paper appears to be useful for studying securitization mechanisms, as shown by the Kortis bond issued by Swiss Re in 2010. The payoff of this first longevity trend bond is linked to the divergence in mortality improvement rates between two countries (US and UK) and thus nicely fits our proposed model.

Furthermore, the model is fitted properly, recognizing the hidden nature of time factors which are not treated as unknown parameters to be estimated from the mortality data. Mortality projections are derived by means of the predictive distribution of the time index, i.e. its a posteriori distribution given past observations. This is the credibility feature of the proposed approach. New data feed this predictive distribution as they become available and so help to update mortality projections. This recognizes the dynamic aspect of mortality forecasting and avoids re-fitting the entire model based on new data. To the best of our knowledge, this dynamic updating approach has not been used so far and our numerical illustrations demonstrate its advantages compared to classical frequentist approaches.

The remainder of this paper is organized as follows. Section 2 gives a short introduction to evolutionary credibility models. Section 3 carefully presents the credibility model proposed to project future mortality. In Section 4, we discuss the covariance structure of the model and address the identifiability problem the model may encounter. Section 5 describes a two-step model fitting concept, which studies period and age effects separately. Section 6 is devoted to empirical illustrations. First, we fit the mortality experience of the G5 countries using our proposed methodology. Then, we study the index governing the payoff of the Swiss Re Kortis bond. Finally, we perform successive forecasts for the Belgian population to illustrate how newly available data can be incorporated in revised forecasts. We compare the results to the official forecasts published yearly by the Federal Planning Bureau, the Belgian agency in charge of mortality projections.

2. Evolutionary credibility models

Following the book of Bühlmann and Gisler (2005) and focusing on the aspects that will be needed later on, this section gives a short introduction to evolutionary credibility modeling.

Consider a time series $(r_t, \Theta_t)_{t\in\mathbb{N}}$ with a w-variate stochastic observation process $(r_t)_{t\in\mathbb{N}}$ and a v-variate stochastic state (or risk profile) process $(\Theta_t)_{t\in\mathbb{N}}$ on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The state process (Θ_t) is unobservable but shall follow a known dynamics. We are now at time $T \in \mathbb{N}$ and the aim is to predict future states Θ_{T+k}, $k \in \mathbb{N}$, and conditional future expected observations $\mathbb{E}[r_{T+k} \mid \Theta_{T+k}]$, $k \in \mathbb{N}$. The past observations $r_1, \dots, r_T$ are the available information at time T.

Let all (r_t, Θ_t), $t \in \mathbb{N}$ be square integrable. Then the credibility estimator for Θ_{T+k}, given the observations till time T, is defined as the orthogonal projection

$$\mu_{T+k|T} := \mathrm{Pro}\,(\Theta_{T+k} \mid L(1, r_1, \dots, r_T)) \tag{2.1}$$

with respect to the set

$$L(1, r_1, \dots, r_T) := \left\{ a + \sum_{t=1}^{T} A_t r_t : a \in \mathbb{R}^v, A_t \in \mathbb{R}^{v,w} \right\}$$

in the Hilbert space of square integrable random variables. In other words, $\mu_{T+k|T}$ is the unique element in the linear space $L(1, r_1, \dots, r_T)$ that satisfies

$$\mathbb{E}[(\mu_{T+k|T} - \Theta_{T+k})(X - \Theta_{T+k})] = 0 \quad \text{for all } X \in L(1, r_1, \dots, r_T).$$

So $\mu_{T+k|T}$ is the best linear predictor of Θ_{T+k} in terms of $1, r_1, \dots, r_T$.

Furthermore, (r_t, Θ_t) is assumed to have a state-space representation of the form

$$r_t = G\Theta_t + W_t, \tag{2.2}$$

$$\Theta_{t+1} = F\Theta_t + V_t \tag{2.3}$$

with $G \in \mathbb{R}^{w,v}$, $F \in \mathbb{R}^{v,v}$ and white noise processes (W_t) and (V_t). The two white noise processes shall be serially uncorrelated and also uncorrelated with each other. Their joint covariance matrix thus has the structure

$$\mathbb{C}\mathrm{ov}\left(\begin{pmatrix} V_t \\ W_t \end{pmatrix}, \begin{pmatrix} V_s \\ W_s \end{pmatrix} \right) = \begin{pmatrix} Q & 0 \\ 0 & R \end{pmatrix}, \tag{2.4}$$

$Q \in \mathbb{R}^{v,v}$ and $R \in \mathbb{R}^{w,w}$, if $t = s$ and zero else.

Under all these assumptions, the credibility estimator $\mu_{T+k|T}$ for Θ_{T+k} can be calculated in a recursive way, see Theorem 10.3 in Bühlmann and Gisler (2005). Starting from an initial value $\mu_{1|0} = \mathbb{E}[\Theta_1]$, the estimate is sequentially updated by the newest observation through the recursive formula

$$\mu_{t|t} = \mu_{t|t-1} + A_t(r_t - G\mu_{t|t-1}) \tag{2.5}$$

for an appropriate matrix A_t. The step from t to $t+1$ then follows the evolution rule (2.3), i.e.

$$\mu_{t+1|t} = F\mu_{t|t}. \tag{2.6}$$

The credibility estimator $\mu_{T+1|T}$ is obtained by iterating this procedure for $t = 1, \dots, T$. Finally, $\mu_{T+k|T}$ and the credibility estimator for $\mathbb{E}[r_{T+k} \mid \Theta_{T+k}] = G\Theta_{T+k}$ are given by

$$\mu_{T+k|T} = F^{k-1}\mu_{T+1|T},$$

$$\mathrm{Pro}\,(G\Theta_{T+k} \mid L(1, r_1, \dots, r_T)) = G\mu_{T+k|T},$$

respectively. This formula is also known as the Kalman recursion or the Kalman filter algorithm, cf. Brockwell and Davis (2006), and is implemented in the statistical software R.

A particular example of a stochastic process that can be expressed as an evolutionary credibility model is an autoregressive