Insurance: Mathematics and Economics 53 (2013) 122–133

Contents lists available at SciVerse ScienceDirect

Insurance: Mathematics and Economics

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

Mortality surface by means of continuous time cohort models

Petar Jevtić [a], Elisa Luciano [b], Elena Vigna [b,*]

[a] University of Torino, Italy
[b] University of Torino, Collegio Carlo Alberto, Italy

HIGHLIGHTS

- We model the mortality surface with continuous-time, cohort-based stochastic intensities.
- We specify to two-factor, Ornstein–Uhlenbeck intensities.
- We model and calibrate correlation across different cohorts.
- Fit of historical data is good, both deterministic and stochastic forecast reliable.
- Differential-Evolution algorithm provides cohort correlations high but not perfect.

ARTICLE INFO

Article history:
Received December 2012
Received in revised form
April 2013
Accepted 14 April 2013

JEL classification:
C1
C13
C38
C53
J11

Keywords:
Stochastic mortality
Age effect
Cohort effect
Differential Evolution algorithm
Mortality forecasting

ABSTRACT

We study and calibrate a cohort-based model which captures the characteristics of a mortality surface with a parsimonious, continuous-time factor approach. The model allows for imperfect correlation of the mortality intensity across generations. It is implemented on UK data for the period 1900–2008. Calibration by means of stochastic search and the Differential Evolution optimization algorithm proves to yield robust and stable parameters. We provide in-sample and out-of-sample, deterministic as well as stochastic forecasts. Calibration confirms that correlation across generations is smaller than one.

1. Introduction

Insurance companies and pension funds are exposed to mortality risk and hope for the development of a liquid and transparent longevity-linked capital market. Active trading of mortality derivatives would help them assessing and hedging the risks they are exposed to, in the same manner as financial models and markets help them mutualize financial risks. Mortality-risk appraisal consisting in an accurate but easy-to-handle description of human survivorship is fundamental in this respect.

In spite of this need, no consensus has been reached yet on the best model for mortality risk modeling. A number of successful proposals have been put forward. Most of these models, starting from the celebrated (Lee and Carter, 1992) model and its several extensions – that include for instance Brouhns et al. (2002) and Renshaw and Haberman (2003), up to the more recent Cairns et al. (2006b) two-factor model – are discrete-time descriptions of survivorship evolution. In some cases though the adoption of a continuous-time approach proves useful. This is the case when we couple the appraisal of mortality and financial risk, and we adopt some financial model such as Black–Scholes or Duffie et al. (2000). Another motivation for adopting a continuous-time description is the search for closed-form evaluation formulas for insurance products and their derivatives. Continuous-time stochastic mortality models for single generation were considered by a number of researchers, including Milevsky and Promislow (2001), Dahl (2004), Biffis (2005), Cairns et al. (2006a), Schrager (2006) and Luciano and Vigna (2008).

* Corresponding author. Tel.: +39 0116705014.
E-mail addresses: petar.j@vektor.net (P. Jevtić), elisa.luciano@unito.it (E. Luciano), elena.vigna@unito.it (E. Vigna).

0167-6687/$ – see front matter
http://dx.doi.org/10.1016/j.insmatheco.2013.04.005

Be it discrete or continuous-time based, a proper description of mortality risk should capture several dimensions. Consider survival probabilities over a given horizon. A satisfactory model should capture their evolution when changing the horizon, for a fixed initial age and cohort (or generation), and its evolution over cohorts, for fixed initial age and horizon. Introducing the cohort dimension though adds a level of complexity to the problem, since it calls for a notion of correlation across generations, which is by no means easy to capture. In principle one has a whole "mortality surface" to model. How to do this while keeping a satisfactory trade-off between the accuracy – or the fit – and the tractability of the model is an open issue. A theoretical extension of the continuous-time single-generation model to the mortality surface appears in Biffis and Millossovich (2006). This is followed by Blackburn and Sherris (2012) who focus also on the calibration aspect. More specifically, in their attempts to calibrate the whole surface in continuous time they make the assumption of perfect correlation across generations. However, common intuition suggests that correlation among close generations is high but not perfect. This suggestion is often implemented in actuarial practice.

In order to reconcile the calibration of the whole mortality surface with common actuarial practice, this paper fits the mortality surface by means of a continuous-time cohort model, that is able to capture correlations across generations. As a relevant consequence, this model provides the actuary with a calibrated correlation among generations rather than a "best estimate" one. Given the same initial age, the intensities of several generations are written in terms of factors, identified via Principal Component Analysis (PCA). Differential Evolution algorithm is a robust stochastic search and optimization algorithm which already proved its use across a wide range of engineering applications. We use it to fit the mortality surface with an extreme precision. Provided that we fully exploit the power of this stochastic search algorithm, we discover that the fitted parameters are extremely robust, stable and lead to correlations across generations that is high but less than one.

The paper unfolds as follows. In Section 2, we review mortality modeling via affine mortality intensities for a single generation. Then, we develop ex-novo a model for the mortality intensities of several generations, i.e. we model the mortality surface. Section 3 specifies a simple two-factor model for modeling the mortality surface, which will then be calibrated. In Section 4, we discuss the criteria that a good mortality model for the mortality surface should satisfy. In Section 5, we proceed to the empirical part. We use PCA to identify the number of relevant factors and apply it to UK males data from the Human Mortality Database. We review the Evolutionary Approach to the global minimum/maximum search and use it in Section 6 to calibrate a two-factor model to a number of UK generations born between 1900 and 1950. We discuss all the key criteria introduced in Section 4. In Section 7 we use polynomial interpolation to further improve parsimoniousness of the model. In Section 8, we conclude and outline further research.

2. The mortality model

In Section 2.1, we illustrate the stochastic mortality intensity setup for one generation only—as is standard in this kind of literature. In Section 2.2, we specify how to move on from the description of the mortality intensity of one generation to the mortality intensities of several generations. This procedure enables us to describe the whole mortality surface. In Section 2.3, we restrict ourselves to constant-parameter dynamics of the Ornstein–Uhlenbeck type. The general mortality model is described in Section 2.4, and a simplified version of it is presented in Section 2.5.

2.1. The affine mortality framework for the single generation

As in the standard unidimensional framework of stochastic mortality (see e.g. Biffis, 2005; Dahl, 2004) we describe the mortality of a given generation by means of a Cox or doubly stochastic counting process. Intuitively, the time of death is supposed to be the first jump time of a Poisson process with stochastic intensity.

Let us introduce a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$, where $\mathbb{P}$ is the real-world probability measure. The filtration $\{\mathcal{F}_t : 0 \leq t \leq T\}$ satisfies the usual properties of right-continuity and completeness. On this space, let us consider a predictable process $\mu(t, x)$, which represents the mortality intensity of an individual belonging to a given generation, initial age x at (calendar) time t. Her death is the first stopping time of a doubly stochastic process with intensity $\mu(t, x)$.

We model the intensity $\mu(t, x)$ of the given generation and initial age x as a function $R(\mathbf{X})$ of a vector of state processes

$$\mathbf{X}(t) = [X_1(t), \ldots, X_n(t)]^\top .$$

For notational simplicity, in the rest of this section we will omit the argument x. Therefore, we have that

$$\mu(t) = R(\mathbf{X}(t)). \tag{1}$$

Moreover, in order to keep the model mathematically tractable, we put ourselves in the affine framework of Duffie et al. (2000) (sometimes referred to as DPS). In this setting $\mathbf{X}$ is a Markov process in some state space $D \subset \mathbb{R}^n$ and it is the solution to the stochastic differential equation

$$d\mathbf{X}(t) = \boldsymbol{\lambda}(\mathbf{X}(t))dt + \boldsymbol{\sigma}(\mathbf{X}(t))d\mathbf{Z}(t),$$

where $\mathbf{Z}$ is an $(\mathcal{F}_t)$-standard Brownian motion in $\mathbb{R}^n$, $\boldsymbol{\lambda} : D \to \mathbb{R}^n$, $\boldsymbol{\sigma} : D \to \mathbb{R}^{n\times n}$, $\boldsymbol{\lambda}$, $\boldsymbol{\sigma}$, and $R : D \to \mathbb{R}$ are affine:

- $\boldsymbol{\lambda}(x) = \mathbf{K}_0 + \mathbf{K}_1 x$, for $\mathbf{K} = (\mathbf{K_0}, \mathbf{K_1}) \in \mathbb{R}^n \times \mathbb{R}^{n\times n}$,
- $(\boldsymbol{\sigma}(x)\boldsymbol{\sigma}(x)^\top)_{ij} = (\mathbf{H}_0)_{ij} + (\mathbf{H}_1)_{ij} \cdot x$, for $\mathbf{H} = (\mathbf{H}_0, \mathbf{H}_1) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times n\times n}$,
- $R(x) = r_0 + \boldsymbol{r}_1 x$, where $(r_0, \boldsymbol{r}_1) \in \mathbb{R} \times \mathbb{R}^n$.

The advantage of this affine choice is that it is possible to calculate in closed form the expectation of functionals of the state variables. In fact, we have

$$\mathbb{E}[e^{-\int_t^T R(\mathbf{X}(s))ds} \mid \mathcal{F}_t] = e^{\alpha(t;T)+\beta(t;T)\cdot\mathbf{X}(t)}, \tag{2}$$

where the coefficients $\alpha(\cdot; T), \beta(\cdot; T) : \mathbb{R}^+ \to \mathbb{R}^n$ satisfy the complex-valued ODEs

$$\beta'(t; T) = \mathbf{r}_1 - \mathbf{K}_1^\top \beta(t; T) - \frac{1}{2}\beta(t; T)^\top \mathbf{H}_1 \beta(t; T),$$

$$\alpha'(t; T) = r_0 - \mathbf{K_0}\beta(t; T) - \frac{1}{2}\beta(t; T)^\top \mathbf{H_0}\beta(t; T),$$

with boundary conditions $\alpha(T, T) = \beta(T, T) = 0$.

In the actuarial context, if the intensity is given by (1), the expectation (2) is the survival probability from t to T, conditional on being alive at t:

$$S(t, T) = \mathbb{E}_t\left[e^{-\int_t^T R(X(s))ds}\right] = \mathbb{E}_t\left[e^{-\int_t^T \mu(s)ds}\right]. \tag{3}$$

2.2. Transition from single generations to the whole mortality surface

In the previous section we have described the mortality intensity of one given generation. However, our main aim is to describe the whole mortality surface, that is composed by a number of different generations. We need then to label each generation with a proper index $i \in \mathbf{I} \subset \mathbb{N}$. Each generation has its own mortality intensity and the intensities of different generations are correlated.