Insurance: Mathematics and Economics 59 (2014) 285–299

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

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

Parametric mortality indexes: From index construction to hedging strategies

Chong It Tan [a,*], Jackie Li [b], Johnny Siu-Hang Li [c], Uditha Balasooriya [a]

[a] Nanyang Business School, Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798, Singapore

[b] Department of Mathematics and Statistics, Curtin University, Perth, Australia

[c] Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1

ARTICLE INFO

Article history:
Received January 2014
Received in revised form
October 2014
Accepted 8 October 2014
Available online 24 October 2014

Keywords:
Cairns–Blake–Dowd model
Mortality indexes
Securitization
Hedging strategies
Longevity risk reduction

ABSTRACT

In this paper, we investigate the construction of mortality indexes using the time-varying parameters in common stochastic mortality models. We first study how existing models can be adapted to satisfy the new-data-invariant property, a property that is required to ensure the resulting mortality indexes are tractable by market participants. Among the collection of adapted models, we find that the adapted Model M7 (the Cairns–Blake–Dowd model with cohort and quadratic age effects) is the most suitable model for constructing mortality indexes. One basis of this conclusion is that the adapted model M7 gives the best fitting and forecasting performance when applied to data over the age range of 40–90 for various populations. Another basis is that the three time-varying parameters in it are highly interpretable and rich in information content. Based on the three indexes created from this model, one can write a standardized mortality derivative called K-forward, which can be used to hedge longevity risk exposures. Another contribution of this paper is a method called key K-duration that permits one to calibrate a longevity hedge formed by K-forward contracts. Our numerical illustrations indicate that a K-forward hedge has a potential to outperform a q-forward hedge in terms of the number of hedging instruments required.

1. Introduction

Pension plan sponsors can mitigate their longevity risk exposures by trading securities that are linked to future realized mortality. To date, the market for such securities is still in its infancy and has yet to overcome a number of challenges. As Blake et al. (2013) pointed out, one of these challenges is the creation of homogeneous and transparent instruments, which allow the market to concentrate liquidity. A significant step in overcoming this challenge is to develop tractable mortality indexes, upon which standardized mortality-linked securities can be written.

Existing mortality indexes produced by investment banks such as J.P. Morgan, Deutsche Börse and Credit Suisse are constructed without using any model. One disadvantage of this construction approach is that the information content of each index is limited, which means the information reflected by an index is either highly aggregate (e.g., the life expectancy at birth) or specific (e.g., the death probability at a certain age).[1] It follows that a large number of such indexes would be needed to effectively hedge longevity risk, which arises from complex and non-parallel shifts in the underlying mortality curve. This problem hinders market development, because liquidity would be diluted across the large spectrum of indexes.

To improve the information content of a mortality index, one may use a model-based construction method, in which mortality indexes are developed from the time-varying parameters in a stochastic mortality model, such as the Lee–Carter model (Lee and Carter, 1992) and the Cairns–Blake–Dowd (CBD) model (Cairns et al., 2006). The model-based approach was first studied by Chan et al. (2014), who argued that the model on which index construction is based must satisfy the new-data-invariant property, which means that when the model is updated with new mortality data, the mortality indexes (time-varying parameters) for the previous years would not be affected. This property is crucially important,

* Corresponding author. Tel.: +65 9793 8939.
E-mail addresses: citan1@e.ntu.edu.sg (C.I. Tan), kakijackie.li@curtin.edu.au (J. Li), shli@uwaterloo.ca (J.S.-H. Li), auditha@ntu.edu.sg (U. Balasooriya).

http://dx.doi.org/10.1016/j.insmatheco.2014.10.005

[1] The mortality index of Credit Suisse is based on the life expectancy at birth of the US population, while that of JP Morgan is based on the death probabilities of four national populations.

because once made public, an index value cannot (and should not) be changed. Chan et al. (2014) found that among the six stochastic models documented by Dowd et al. (2010), the original CBD model (also called Model M5) is the only model that satisfies the new-data-invariant property.

The first objective of this paper is to further investigate the construction of mortality indexes from stochastic mortality models. We begin with a discussion on how models other than the original CBD model can be adapted to satisfy the new-data-invariant property. We then evaluate the fit of the adapted models to the mortality data from 8 national populations on the basis of the Bayesian Information Criterion (BIC). Among the collection of adapted models, we find that the adapted Model M7 (the Cairns–Blake–Dowd model with cohort and quadratic age effects) yields the best BIC values when fitted to data over the age range of 40–90 for various populations. The model also performs the best in a backtest that compares projected values with actual values. On top of that, the time-varying parameters in the model are highly interpretable and are able to reflect the varying age pattern of mortality improvement. For these reasons, we propose to use the three time-varying parameters in the adapted Model M7 jointly as mortality indexes. We call these indexes the 3-factor CBD mortality indexes.

Our second objective is to develop standardized securities written on the 3-factor CBD mortality indexes. In particular, we explain how a security called K-forward, proposed as a concept by Chan et al. (2014), can be used to hedge the longevity risk exposures of a pension plan. The structure of a K-forward is identical to that of a q-forward (see, e.g., Coughlan, 2009), except that the reference rates to which the contracts are linked are different. In more detail, the payoff from a q-forward depends on the realized death probability at a reference age in a reference year, whereas that from a K-forward depends on a realized 3-factor CBD mortality index in a reference year. Compared to a q-forward, a K-forward is an even simpler building block, because its reference rate contains only one parameter (the reference year) instead of two.

To ensure that a hedge formed by standardized instruments is effective, there is a need to calibrate it. Following the lines of Cairns et al. (2008), Cairns (2011), Coughlan et al. (2007), Plat (2009, 2010), Li and Luo (2012), Lin and Tsai (2013), Tsai et al. (2010) and Wang et al. (2010), we contribute a measure called key K-duration, which measures a liability's price sensitivity to a specific segment in the time trend of a 3-factor CBD mortality index. The required notional amounts of K-forwards can be determined readily by equating the key K-durations of the portfolio of K-forwards and the liability being hedged. The key K-duration measure is parallel to Li and Luo's (2012) key q-duration, which measures the change in a liability's value due to a small change in a death probability. It also has a close resemblance to Cairns' (2011) approximate deltas (with respect to the time-varying parameters in the original CBD model) and to Plat's (2009) minimum variance hedge ratios that are derived by considering the shifts of the two time-varying parameters in the original CBD model.

The calculation of key K-durations does not require any simulation, so the execution of the proposed hedging strategy is quick and requires minimal computational effort. Our numerical illustrations indicate that the proposed hedging strategy is effective in reducing a portfolio's longevity risk exposure, even if parameter uncertainty and sampling risk are taken into account. Furthermore, although the mortality indexes are derived from the adapted Model M7, our proposed hedging strategy also works well under scenarios simulated from other stochastic mortality models, suggesting that the success in hedging is largely independent of the simulation model and is likely to be achievable in practice.

The advantage of our proposed hedging strategy is particularly apparent when the portfolio being hedged involves individuals who were born in different years. First, a K-forward hedge is easier to execute in comparison to a q-forward hedge, which requires the hedger to determine the key cohorts in the portfolio. Second, our numerical illustrations suggest that compared to a q-forward hedge, a K-forward hedge giving a comparable hedge effectiveness involves a smaller number of securities. This helps the market to concentrate liquidity, thereby facilitating market development.

Table 1
Specifications of the six stochastic mortality models under consideration.

Model M1: The Lee–Carter model (Lee and Carter, 1992)	
$\ln(m_{x,t}) = \beta_x^{(1)} + \beta_x^{(2)}\kappa_t^{(2)}$	(2 constraints)
Model M2: The Renshaw–Haberman model (Renshaw and Haberman, 2006)	
$\ln(m_{x,t}) = \beta_x^{(1)} + \beta_x^{(2)}\kappa_t^{(2)} + \beta_x^{(3)}\gamma_{t-x}^{(3)}$	(4 constraints)
Model M3: The Age–Period–Cohort model (Osmond, 1985)	
$\ln(m_{x,t}) = \beta_x^{(1)} + n_a^{-1}\kappa_t^{(2)} + n_a^{-1}\gamma_{t-x}^{(3)}$	(3 constraints)
Model M5: The original Cairns–Blake–Dowd (CBD) model (Cairns et al., 2006)	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x})$	(No constraint)
Model M6: The CBD model with a Cohort effect term (Cairns et al., 2009)	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \gamma_{t-x}^{(3)}$	(2 constraints)
Model M7: The CBD model with Cohort effect and quadratic age effect terms (Cairns et al., 2009)	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \kappa_t^{(3)}((x-\bar{x})^2 - \hat{\sigma}_x^2) + \gamma_{t-x}^{(4)}$	(3 constraints)

The remainder of this paper is organized as follows. Section 2 discusses the construction of mortality indexes using common stochastic mortality models. Section 3 specifies a K-forward contract, defines the key K-duration measure and details the proposed hedging strategy. Section 4 illustrates the proposed methods with a hypothetical pension plan involving one single cohort and investigates important issues such as sampling risk. Section 5 presents the generalization to multiple birth cohorts. Section 6 concludes the paper.

2. Constructing mortality indexes

In this section, we revisit the problem of model-based mortality index construction, which was previously studied by Chan et al. (2014). The conventions below are used throughout the discussion:

- $m_{x,t} = \frac{D_{x,t}}{E_{x,t}}$ is the central death rate at age x in year t;
- $D_{x,t}$ is observed number of deaths at age x in year t;
- $E_{x,t}$ is the matching exposures at age x in year t;
- $q_{x,t}$ is the probability that a person aged x at time t will die between time t and $t+1$;
- $\beta_x^{(i)}$, $i = 1, 2, 3$, are age-specific parameters;
- $\kappa_t^{(i)}$, $i = 1, 2, 3$, are time-varying parameters;
- $\gamma_c^{(i)}$, $i = 3, 4$, where $c = t - x$ denotes year of birth, are cohort-related parameters;
- n_a is the number of ages covered in the sample age range;
- $\bar{x}$ is the mean age over the sample age range;
- $\hat{\sigma}_x^2$ is the mean of $(x-\bar{x})^2$ over the sample age range.

The candidate models under consideration are the six stochastic mortality models discussed by Dowd et al. (2010). The specifications of the six models are displayed in Table 1. In fitting these models (except Model M5), we need to impose identifiability constraints to stipulate parameter uniqueness. The number of identifiability constraints needed for each model is displayed in parentheses in Table 1. We refer interested readers to Cairns et al. (2009) for a deeper discussion on the identifiability constraints.