Available online at www.sciencedirect.com
SCIENCE DIRECT®
Insurance: Mathematics and Economics 38 (2006) 427–440
INSURANCE MATHEMATICS & ECONOMICS
www.elsevier.com/locate/ime

Mortality-dependent financial risk measures

Kevin Dowd [a,*], Andrew J.G. Cairns [b], David Blake [c]

[a] *Centre for Risk and Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, UK*
[b] *Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK*
[c] *Pensions Institute, Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, UK*

Received June 2005; received in revised form October 2005; accepted 17 November 2005

Abstract

This paper uses a recently developed two-factor stochastic mortality model to estimate financial risk measures for four illustrative types of mortality-dependent financial position: investments in zero-coupon longevity bonds; investments in longevity bonds that pay annual survivor-dependent coupons; and two examples of an insurer's annuity book that are each hedged by a longevity bond, one based on the annuity book and hedge having the same reference cohort, and the other not. The risk measures estimated are the value-at-risk, the expected shortfall and a spectral risk measure based on an exponential risk-aversion function. Results are reported on a model calibrated on data provided by the UK Government Actuary's Department, both with and without underlying parameter uncertainty.

Keywords: Mortality risk; Longevity bonds; Value-at-risk; Coherent risk measures; Spectral risk measures

1. Introduction

There has been considerable interest in recent years in models of mortality. Much of this interest has arisen from the belated recognition – evident, for example, in the Equitable Life fiasco[1] – that actuaries in the past have tended to pay insufficient attention to aggregate mortality risk. It is also becoming clear that many insurance companies have considerable exposure to this risk, and that they currently lack the tools to price and manage this risk as effectively as they should. It is therefore not surprising that a number of mortality risk models have been proposed in the past few years,[2] and that there have also been proposals for mortality derivatives, most particularly for longevity bonds (LBs).[3] The first publicly offered mortality derivative – the Swiss Re LB – was then issued in December 2003,[4] and this was followed by the European

* Corresponding author. Tel.: +44 115 846 6682; fax: +44 115 846 6667.
E-mail address: Kevin.Dowd@nottingham.ac.uk (K. Dowd).

[1] In this case, the firm offered guaranteed annuity options based on 1950s mortality tables, and its failure to take proper account of mortality risk was a key factor contributing to its being forced to close to new business in 2000.

[2] These models include those of Milevsky and Promislow (2001), Yang (2001), Cairns et al. (2006a), Dahl (2004) and Lin and Cox (2005b).

[3] Longevity bonds were first proposed by Blake and Burrows (2001) (under the name of survivor bonds), and further developed by Lin and Cox (2005a) among others. Mortality swaps were suggested by Lin and Cox (2004) and Dowd et al. (2006).

[4] This was an otherwise conventional bond whose principal payment was linked to adverse mortality risk scenarios. The face value was $400 million and the maturity of the issue was 4 years. Investors received a floating coupon rate of US LIBOR plus 135 basis points, but the principal payment was at risk if the weighted average of general population mortality across five reference countries (US, UK, France, Italy and Switzerland) should exceed 130% of the 2002 level. Since mortality is improving, the chances of such high mortality were judged to be very

0167-6687/$ – see front matter
doi:10.1016/j.insmatheco.2005.11.003

Investment Bank/Bank National de Paris longevity bond announced in November 2004.[5] Financial institutions have also started to trade mortality derivatives over-the-counter, and the indications are that a wholesale market in aggregate mortality risk is beginning to take shape.

These developments are encouraging, but insurance companies still have a problem: how do they assess the magnitude of the financial risks implied by their mortality exposures? This paper seeks to provide an answer to this problem. The answer suggested has three key ingredients: a stochastic model of aggregate mortality; some financial risk measures; and a simulation framework that enables us to estimate these financial risk measures for a mortality-dependent portfolio based on the underlying mortality model. In addition, since the parameters of the mortality model are (inevitably) uncertain, our estimation framework should also take account of the uncertainty attached to estimates of the parameters (i.e., parameter risk).

This paper is organized as follows. Section 2 summarizes the mortality model used in the study: it covers the main features of the model and addresses the associated issues of calibration and simulation. Section 3 looks at applications of this model to estimate the risk measures of four types of mortality-dependent financial position: zero-coupon LBs that make a single mortality-dependent payment; coupon-paying LBs that make annual mortality-dependent payments; an annuity book that is hedged by a t-year coupon-paying LB predicated on the same reference population, where t is taken from 1 to 50, and an annuity book that is hedged by a t-year coupon-paying LB predicated on a different reference population. Section 4 explains the risk measures to be estimated—the VaR, ES and spectral risk measures. Section 5 presents two alternative sets of estimates of these risk measures, one ignoring parameter risk, and the other taking account of it. Section 6 concludes.

low, so investors obtained a high coupon rate in return for assuming some degree of exposure to extreme mortality risk.

[5] This instrument was issued by the EIB and managed by BNP Paribas. The face value was £540 million, and involved time t coupon payments that were tied to an initial annuity payment of £50 million indexed to the time t survivor rates of English and Welsh males aged 65 years in 2003. The EIB/BNP bond was closer in spirit to a 'classic' longevity bond, because it tied coupon payments to a survivor index and dealt with likely mortality risks rather than extreme ones.

2. A stochastic mortality model

Let $S(t,x)$ be the survivor rate at time t of a cohort aged x in year 0. For any given x, $S(0, x) = 1$ and we expect $S(t, x)$ to diminish as t gets bigger and eventually go to 0 as t gets very large. We also know that if $q(t, x)$ is the realized mortality rate in year $t+1$ (that is, from time t to time $t+1$) of our cohort, then

$$S(t+1, x) = (1 - q(t, x))S(t, x) \tag{1}$$

We assume that $q(t, x)$ is governed by the following two-factor Perks stochastic process:

$$q(t, x) = \frac{e^{A_1(t+1)+A_2(t+1)(t+x)}}{1 + e^{A_1(t+1)+A_2(t+1)(t+x)}} \tag{2}$$

where $A_1(t+1)$ and $A_2(t+1)$ are themselves stochastic processes that are measurable at time $t+1$ (see Perks, 1932; Benjamin and Pollard, 1993). Cairns et al. (2006b) generate empirical results showing that this mortality model provides a good fit to realized male mortality data in England and Wales. Their results also indicate that a two-factor model of UK mortality fits the data better than a one-factor one.

Now let $A(t) = (A_1(t), A_2(t))'$ and assume that $A(t)$ is a random walk with drift:

$$A(t+1) = A(t) + \mu + CZ(t+1) \tag{3}$$

where μ is a constant 2×1 vector of drift parameters, C the constant 2×2 lower triangular matrix reflecting volatilities and correlations, and $Z(t)$ is a 2×1 vector of independent standard normal variables. Cairns et al. (2006b) also show that if we use the UK Government Actuary's Department (GAD) data for English and Welsh males over 1961–2002, then the least squares estimates of our parameters are:

$$\hat{\mu} = \begin{bmatrix} -0.04340 \\ 0.000367 \end{bmatrix} \tag{4a}$$

$$\hat{V} = \hat{C}\hat{C}' = \begin{bmatrix} 0.01067000, -0.00016170 \\ -0.00016170, 0.00000259 \end{bmatrix} \tag{4b}$$

We can recover $\hat{C}$ from $\hat{V}$ using a Choleski decomposition, and all that remains is to specify a suitable starting value $A(0)$. The results of Cairns et al. suggest that we might take $A(0) \approx (-11.0, 0.107)'$ if we take 2003 as our starting point (i.e., if we set $t = 0$ for the end of 2003).[6]

[6] To elaborate: Fig. 2 in Cairns et al. (2006b) plots estimated $A(0)$ values against time for the years 1961–2002, and the values $A(0) \approx (-11.0, 0.107)'$ for the year 2003 are obtained by simple extrapolation.