Available online at www.sciencedirect.com

ScienceDirect

Insurance: Mathematics and Economics 43 (2008) 41–55

INSURANCE MATHEMATICS & ECONOMICS

www.elsevier.com/locate/ime

Heath–Jarrow–Morton modelling of longevity bonds and the risk minimization of life insurance portfolios

Jérôme Barbarin[1]

Université Catholique de Louvain, Institute of Actuarial Science, 6, rue des Wallons, 1348 Louvain-La-Neuve, Belgium

Received May 2007; received in revised form September 2007; accepted 8 September 2007

Abstract

This paper has two parts. In the first, we apply the Heath–Jarrow–Morton (HJM) methodology to the modelling of longevity bond prices. The idea of using the HJM methodology is not new. We can cite Cairns et al. [Cairns A.J., Blake D., Dowd K, 2006. Pricing death: framework for the valuation and the securitization of mortality risk. Astin Bull., 36 (1), 79–120], Miltersen and Persson [Miltersen K.R., Persson S.A., 2005. Is mortality dead? Stochastic force of mortality determined by arbitrage? Working Paper, University of Bergen] and Bauer [Bauer D., 2006. An arbitrage-free family of longevity bonds. Working Paper, Ulm University]. Unfortunately, none of these papers properly defines the prices of the longevity bonds they are supposed to be studying. Accordingly, the main contribution of this section is to describe a coherent theoretical setting in which we can properly define these longevity bond prices. A second objective of this section is to describe a more realistic longevity bonds market model than in previous papers. In particular, we introduce an additional effect of the actual mortality on the longevity bond prices, that does not appear in the literature. We also study multiple term structures of longevity bonds instead of the usual single term structure. In this framework, we derive a no-arbitrage condition for the longevity bond financial market. We also discuss the links between such HJM based models and the intensity models for longevity bonds such as those of Dahl [Dahl M., 2004. Stochastic mortality in life insurance: Market reserves and mortality-linked insurance contracts, Insurance: Math. Econom. 35 (1) 113–136], Biffis [Biffis E., 2005. Affine processes for dynamic mortality and actuarial valuations. Insurance: Math. Econom. 37, 443–468], Luciano and Vigna [Luciano E. and Vigna E., 2005. Non mean reverting affine processes for stochastic mortality. ICER working paper], Schrager [Schrager D.F., 2006. Affine stochastic mortality. Insurance: Math. Econom. 38, 81–97] and Hainaut and Devolder [Hainaut D., Devolder P., 2007. Mortality modelling with Lévy processes. Insurance: Math. Econom. (in press)], and suggest the standard pricing formula of these intensity models could be extended to more general settings.

In the second part of this paper, we study the asset allocation problem of pure endowment and annuity portfolios. In order to solve this problem, we study the "risk-minimizing" strategies of such portfolios, when some but not all longevity bonds are available for trading. In this way, we introduce different basis risks.

JEL classification: G10

Subject classification: IB10

Keywords: Term structure model; No arbitrage; Martingale measure; Hedging strategies; Marked point process; Jump-diffusion

1. Introduction

A zero-coupon longevity bond is a financial security whose single payout occurs at maturity and is equal to the value at this time, of a so-called survivor index. The value at any times t, of a survivor index, is given by the proportion of people still alive at time t in an initially given population. We can distinguish two broad approaches in the literature to the modelling of longevity bond prices. A first approach is known as the "intensity" approach. In a sense, it is the counterpart to the "short-rate" approach in the interest rate term structure literature. Typical examples of these intensity models are, to give a few examples, Biffis (2005), Dahl (2004), Luciano and Vigna (2005), Schrager (2006) and Hainaut and Devolder (in press). The second approach consists of adapting

E-mail address: barbarin@actu.ucl.ac.be.

[1] Axa Belgium Chair in Risk Management.

0167-6687/$ - see front matter
doi:10.1016/j.insmatheco.2007.09.008

the Heath–Jarrow–Morton (HJM) (see Heath et al. (1992)) methodology to the modelling of longevity bond prices. Cairns et al. (2006) is, to our knowledge, the first paper to exploit this idea. These authors studied different specifications for the prices of longevity bonds. In particular, they described the term structure of longevity bond prices as the product of two independent HJM models, one related to the term structure of risk-free zero-coupon bonds and a second one related to the survival probabilities. Miltersen and Persson (2005) extended Blake et al.'s results by removing this independence. However, they assumed the payouts of their longevity bonds at maturity were defined on the survival of a single individual. Bauer (2006) introduced the fact that the actual payouts of the longevity bonds depend on the proportion of individuals alive in a given population, and not on a single individual.

Unfortunately, none of these papers properly defines the prices of the longevity bonds they are supposed to study. See Barbarin (2007) for a detailed discussion of these papers. The trouble comes from the fact that they fail to take into account the effect of the actual mortality in the population on the prices of their longevity bonds. The first and, probably, the main goal of this paper is, accordingly, to describe a coherent theoretical setting in which we can rigorously apply the HJM methodology to the modelling of longevity bond prices. For this, it is necessary to introduce an explicit model of the (actual) number of deaths in a given population, through a more or less general (marked) point process. Modelling this mortality explicitly is fundamental to offering a coherent and rigorous definition of longevity bond prices.

A second important objective of this paper is to describe a more realistic financial market. We extend the existing literature in various directions. Firstly, we take into account an additional effect of the actual mortality on the longevity bond prices. We can distinguish two main effects of this actual mortality. The first is a purely mechanical effect. If 5% of the population of the survivor index on which a longevity bond is defined dies, then, everything else being equal, the price of the longevity bond will also decline by 5%. This simple mechanical and proportional effect clearly appears, for example, in the pricing formula of longevity bonds in the intensity models. This first effect can obviously only depend on the actual mortality in the population of the survivor index on which the longevity bond is defined. However, there is also a possibly much more subtle second effect which has (to my knowledge) never been taken into account either in the intensity models or in the HJM models. We know from financial economic theory that the price at time t of a longevity bond with maturity T should depend on the survival probabilities up to time T, possibly adjusted for the risk, in the population on which the survivor index is defined. *A priori*, there is no reason to believe these (risk-adjusted) survival probabilities might not also depend on the actual mortality in the population. Investors, if they are rational, should update their estimates of the survival probabilities (even under the real measure) if they observe a mortality which is not in line with their previous expectations. They could even modify their correction for risk if the actual mortality makes them think they are bearing more (or less) systematic risk than they previously thought. Unlike the first effect, this second effect does not necessarily depend only on the actual mortality in the survivor index on which the longevity bond is defined; it may also depend on the mortality in other parts of the population. To be more concrete, let us assume we have a longevity bond defined on 40-year-old individuals and assume the time to maturity is 20 years. The price of this bond should depend not only on the actual current mortality of 40-year-old individuals, but also on the actual current mortality of 45, 50, 55, etc.-year-old individuals, because investors can use this information to update their expected future survival probabilities of the current 40-year-old population. This second effect does not appear in the existing intensity models because these models assume, at least implicitly, the random times of death are conditionally independent with respect to a certain filtration (which does not include the actual mortality) and to the risk-adjusted measure (the risk-neutral martingale measure). To me, these assumptions are not obvious.

Secondly, we simultaneously consider different survivor indices, each defined on a different age group of the overall population. We can thus define a longevity bond term structure on each of these indices. In other words, unlike the existing literature, we do not consider a single longevity bond term structure but multiple longevity bond term structures. This is a natural extension, since longevity bonds defined on several indices would allow the insurance companies and pension funds to better tailor their assets to their liabilities. As the longevity bonds market matures, more and more longevity bonds, defined on different survivor indices, will naturally be issued. Since each longevity bond depends on the actual mortality in the whole population, complex dependencies could appear between the prices of the different longevity bonds. In particular, it is interesting to study the conditions yielding the absence of arbitrage between the different longevity bond term structures.

Finally, we extend the existing literature by allowing the different survivor indices to be heterogeneous. In the current "intensity" models, the survivor index is always defined on individuals of the same age and who are homogeneous (identical forces of mortality). We do not make these assumptions here. In particular, our survivor indices can be defined on individuals of different ages. As we said, it is very likely that longevity bonds defined on different survivor indices will be issued in the future, but the number of different survivor indices will probably remain relatively low. In order to cover a large proportion of the insured population with a limited number of indices, it might be interesting to simply define larger survivor indices. For example, one longevity bond could be defined on the population currently between 20 and 35 years old, another on the population currently between 35 and 45 years old, etc. Survivor indices could be defined on larger or smaller age groups according to the needs of the insurance companies or the pension funds. Obviously, the population of such survivor indices would be heterogeneous from the point of view of mortality. Our model allows such "heterogeneous" longevity bonds to be dealt with.

The third aim is to study the asset allocation problem for portfolios of endowments and annuities. In particular,