Pricing Death:
Frameworks for the Valuation and Securitization of Mortality Risk

Andrew J.G. Cairns[ab], David Blake[c], Kevin Dowd[d]

First version: June 2004
This version: September 28, 2004

Abstract

It is now an accepted fact that stochastic mortality – the risk that actual future trends in mortality might differ from those anticipated – is an important risk factor in both life insurance and pensions. As such it affects how fair values, premium rates, and risk reserves are calculated.

This paper makes use of the similarities between the force of mortality and interest rates to show how we can model mortality risks and price mortality-related instruments using adaptations of the arbitrage-free pricing frameworks that have been developed for interest-rate derivatives. In so doing, it develops a range of arbitrage-free (or risk-neutral) frameworks for pricing and hedging mortality risk that allow for both interest and mortality factors to be stochastic. The different frameworks that we describe – short-rate models, forward-mortality models, positive-mortality models and mortality market models – are all based on positive-interest-rate modelling frameworks since the force of mortality can be treated in a similar way to the short-term risk-free rate of interest.

These frameworks can be applied to a great variety of mortality-related instruments, from vanilla survivor bonds to exotic mortality derivatives.

Keywords: stochastic mortality, term structure of mortality, survivor index, spot survival probabilities, spot force of mortality, forward mortality surface, short-rate models, forward mortality models, positive mortality framework, mortality market models, annuity market model, SCOR market model.

[a]Department of Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom.

[b]Corresponding author: E-mail A.Cairns@ma.hw.ac.uk

[c]Pensions Institute, Cass Business School, City University, 106 Bunhill Row, London, EC1Y 8TZ, United Kingdom.

[d]Centre for Risk & Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham, NG8 1BB, United Kingdom.

1   Introduction

A large number of products in life insurance and pensions by their very nature have mortality as a primary source of risk. By this we mean that products are exposed to *unanticipated* changes over time in the mortality rates of the appropriate reference population. For example, annuity providers are exposed to the risk that the mortality rates of pensioners will fall at a faster rate than accounted for in their pricing and reserving calculations, and life insurers are exposed to the risk of unexpected increases in mortality (a recent example being those due to HIV/AIDS). On the asset side of their balance sheets, insurance companies are also exposed to investment risks and, since their investment portfolios are predominantly fixed income, this means that they are heavily exposed to interest-rate risk.

However, there is a huge gap in the tools available to model these two types of risk. On the one hand, the theory and practice of interest-rate modelling is very well developed (see, for example, Vasicek, 1977, Cox, Ingersoll and Ross, 1985, Heath, Jarrow and Morton, 1992, Brace, Gatarek and Musiela, 1997, Jamshidian, 1997, Brigo and Mercurio, 2001, James and Webber, 2002, Rebonato, 2002, and Cairns, 2004b). On the other hand, the state and practice of mortality risk modelling is primitive.

Yet there are important similarities between the force of mortality and interest rates: most obviously, they are both positive and have term structures, but we would argue that they are also similar in being stochastic. These similarities suggest that we should be able to model both types of risks using similar approaches. This paper seeks to develop this insight further. In particular, it makes use of the similarities between mortality and interest rate risks to show how we can model mortality risks and price mortality-related instruments using adaptations of the arbitrage-free pricing frameworks that have been developed for interest-rate derivatives. In so doing, it develops a range of arbitrage-free frameworks for pricing and hedging mortality risk that allow for both interest and mortality factors to be stochastic.

To motivate our discussion, we will first summarise some evidence that confirms that mortality improvements are indeed stochastic. We then briefly discuss the state of the art in mortality modelling, and go on to consider some financial instruments whose values depend on mortality and where it is therefore important to model mortality risk factors in an appropriate way.

The idea that mortality is stochastic is not a new one, and it has been evident for many years that mortality rates have been evolving in an apparently stochastic fashion. The uncertainty of mortality forecasts is illustrated in recent work by Currie, Durban and Eilers (2004) (hereafter CDE) which analysed historical trends in mortality using P-splines. The fitted surface of values for the force of mortality[1] $\hat{\mu}(t, x)$

---

[1]The force of mortality $\mu(t, x)$ is described in more detail at the start of Section 2. Note that t represents the current time, x the age at time t of a specified life. The probability that the

is plotted on a log scale in Figure 1.1 while the development of the force of mortality for specific ages over time relative to values in 1947 is plotted in Figure 1.2. Figure 1.2 reveals some detail that we cannot easily see in Figure 1.1: specifically that the rate of improvement has varied substantially over time, and that the improvements have varied substantially between different age groups. CDE also constructed confidence bounds for the future development of mortality rates. Inevitably these confidence bounds get wider as the forecast horizon lengthens and CDE found that even 15-20 years ahead the bounds are very wide. In general terms, the analysis of CDE, as well as other analyses using the stochastic mortality models discussed below, indicates that future mortality improvements cannot be forecast with any degree of precision. Other studies (e.g., Forfar and Smith, 1987, Macdonald *et al*, 1998, Willetts, 1999, and Macdonald *et al*, 2003) have come to similar conclusions.

A number of recent studies have sought to model mortality as a stochastic process. We shall see presently that all these studies bar one (Lin and Cox (2004)[2]) can be reformulated into one of the more general frameworks that we will describe later in this paper. We will describe these briefly here and in the Appendices. Apart from the one exception, all of these use what we describe as *short-rate models* for mortality: that is, they are modelling the *spot mortality rates*[3], $q(t, x)$, or the *spot force of mortality*, $\mu(t, x)$.

The most coherent group of papers that use the short-rate-modelling framework in discrete time build on the original work of Lee and Carter (1992). This study introduced a simple model for central mortality rates involving both age-dependent and time-dependent terms and applied it to US population data (see Appendix A for further details).[4] The time-dependency is modelled using a univariate ARIMA time-series model implying that changes in the mortality curve at all ages are perfectly correlated. Brouhns, Denuit and Vermunt (2002) applied the same model to Belgian data and also improved some of the statistical aspects of Lee and Carter's work. The possibility of imperfect mortality correlation was investigated by Renshaw and Haberman (2003) who extend the Lee and Carter approach by adding a second time-dependent set of changes.

A second approach that also uses the short-rate-modelling framework in discrete

---

individual will die between times t and $t + dt$ given he has survived until the current time, t, is $\mu(t, x)dt + o(dt)$ as $dt \rightarrow 0$ (that is, approximately $\mu(t, x)dt$ for small dt).

[2]In contrast to the papers mentioned below, Lin and Cox (2004) do not propose a specific model for stochastic mortality. Instead, they apply the Wang (1996, 2000, 2002, 2003) transform to convert deterministic projected mortality rates into risk-neutral probabilities. The use of the Wang transform is gaining in popularity in non-life insurance applications where there is a lack of liquidity in the instruments subject to the underlying risks. However, it is not clear from Lin and Cox (2004) how different transforms for different cohorts and terms to maturity relate to one another to form a coherent whole.

[3]The spot mortality rate $q(t, x)$ is the probability at time t that an individual who is aged x and still alive at time t will die before time $t + 1$.
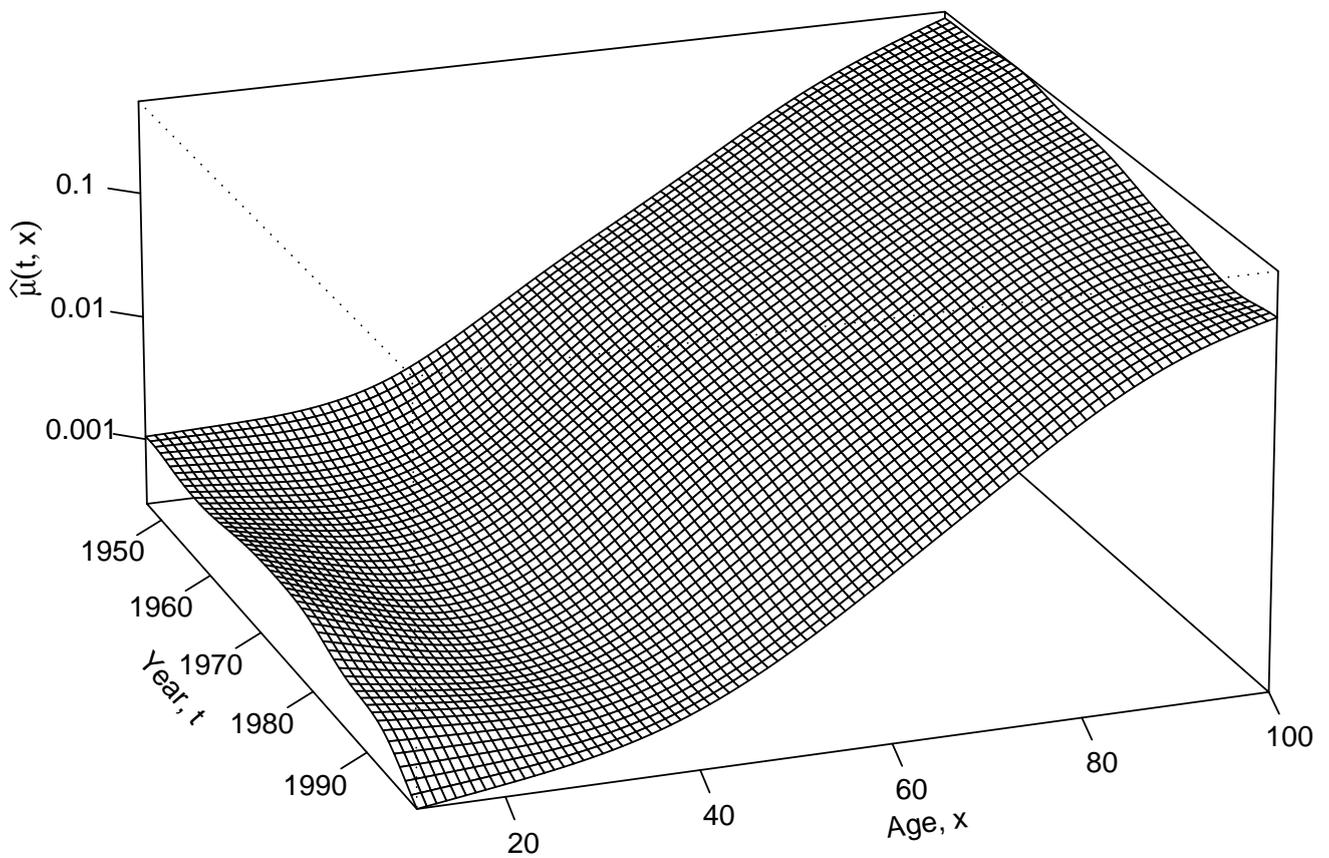
[4]See, also, Lee, 2000b.

Figure 1.1: Fitted values using P-splines for the force of mortality $\hat{\mu}(t, x)$ for the years $t = 1947$ to 1999 and for ages $x = 11$ to 100 from Currie, Durban and Eilers, 2004. (Data: UK males, assured lives.)
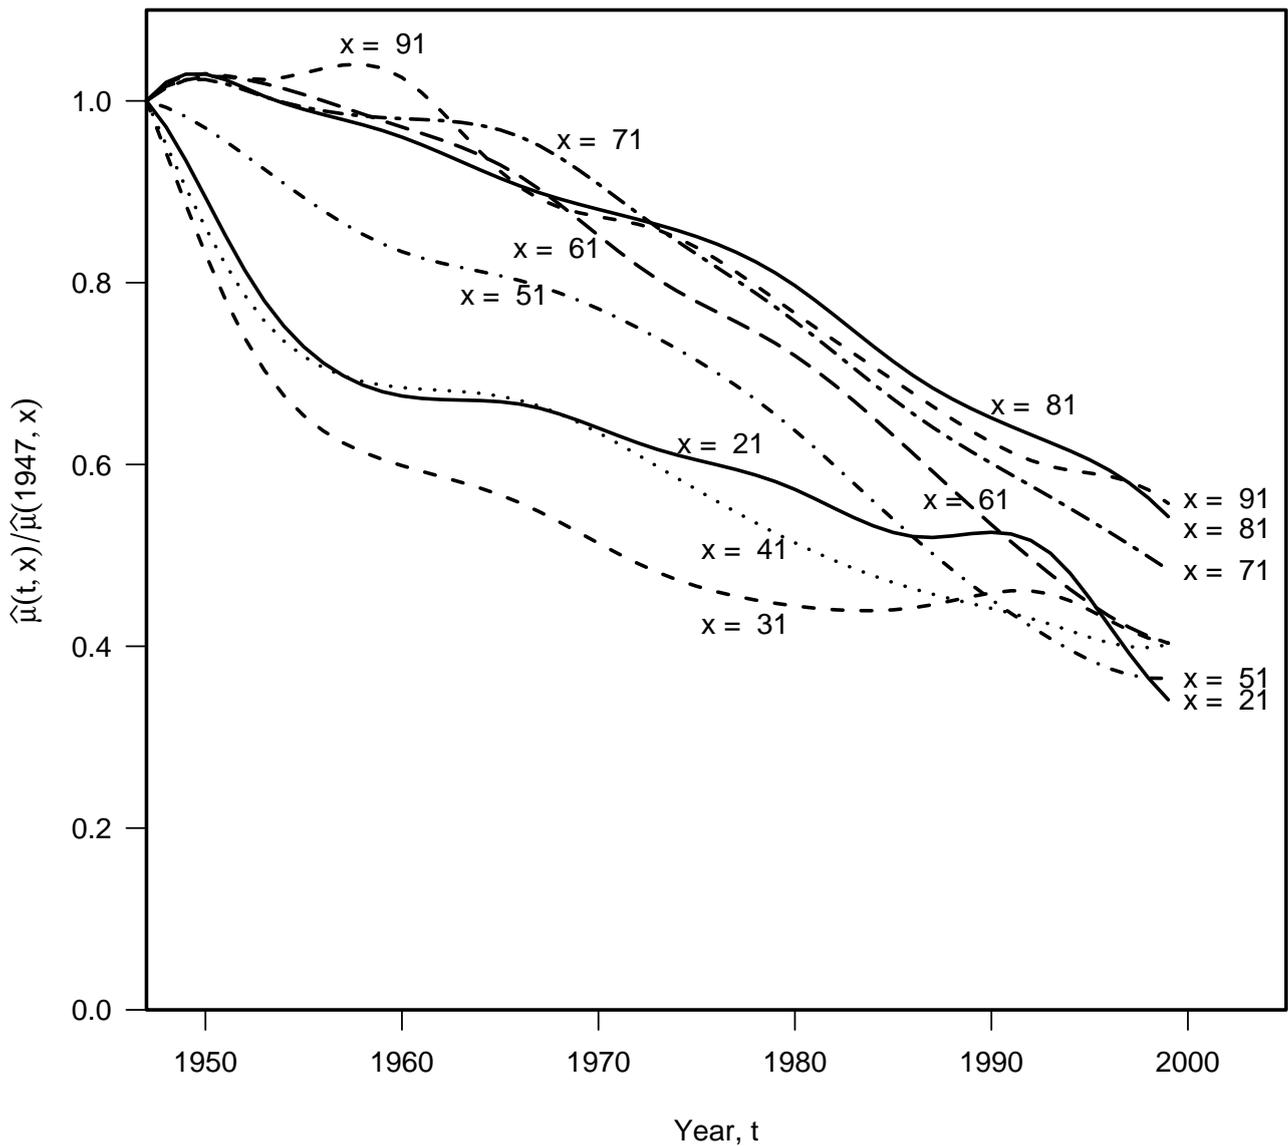
Figure 1.2: $\hat{\mu}(t,x)/\hat{\mu}(1947,x)$: Fitted values using P-splines for the force of mortality $\hat{\mu}(t,x)$, relative to the 1947 value for the years $t = 1947$ to 1999 and for ages $x = $ 21, 31, 41, 51, 61, 71 and 81 from Currie, Durban and Eilers, 2004. Note that the pattern of improvements is different at different ages. (Data: UK males, assured lives.)

time has been proposed by Lee (2000a) and Yang (2001). They take a deterministic projection of the spot mortality rates, $\hat{q}(t, x)$, as given, and then apply an adjustment that evolves over time in a stochastic way. (For further, brief details, see Appendix B.)

Short-rate models for the development in continuous time of the force of mortality have been proposed by Milevsky and Promislow (2001) and Dahl (2004). Milevsky and Promislow (2001) take a more theoretical approach in continuous time which assumes that the force of mortality $\mu(t, x)$ has a Gompertz form $\xi_0(t) \exp(\xi_1 x)$ where the $\xi_0(t)$ term is modelled using a simple mean-reverting diffusion process (see Appendix C). Dahl (2004) on the other hand develops a parsimonious, affine class of processes and we discuss his approach further in Appendix D.

There are many financial applications where it is necessary to take account of the stochastic behaviour of mortality. One example is the calculation of quantile (or value-at-risk) reserves for life-office portfolios, where the uncertain future pattern of liability payments will depend, amongst other things, on the future evolution of the force of mortality $\mu(t, x)$. It is also important to take account of stochastic mortality when reserving for policies that incorporate certain types of guarantee. For example, a guaranteed annuity option is an investment-linked deferred-annuity contract that gives a policyholder the option to convert his accumulated fund at retirement at a guaranteed rate rather than at current market rates. The value of this option most obviously depends upon the level of interest rates at retirement, but also depends upon the mortality table being used by the life office at the time of retirement, which, in turn, depends on mortality forecasts at that time.

Taking account of stochastic mortality is also critical when pricing mortality derivatives. Examples of such contracts include:

- Survivor bonds (where coupon payments are linked to the number of survivors in a given cohort). Long-dated survivor bonds intended to manage longevity risk[5] have recently been revived by Cox, Fairchild and Pedersen (2000)[6], Blake and Burrows (2001) and Lin and Cox (2004). Their origin dates back to Tontine bonds issued by a number of European governments in the 17^{th} and 18^{th} centuries.

---

[5]We use, in this context, the term *longevity risk* to refer specifically to the risk that future survival rates are higher than anticipated. For most of the remainder of the paper, we will use the more general term *mortality risk* to refer to all types of deviation from that anticipated. The most obvious form of this risk is derived from experienced mortality and survival rates between now and some specified future date T. An additional risk applies in cases where cash flows depend upon a mortality table in use on a given future date: for example, a pension contract where a lump sum at T is used to purchase an annuity at prevailing market rates at T.

[6]Cox, Fairchild and Pedersen comment that a number of insurers were proposing to issue survivor bonds as early as 1997. The absence of any issues up to the present time suggest that there are practical problems which still need to be overcome or that the market is not yet ready to invest in such long-term bonds.

- Short-dated, mortality-linked securities (market-traded securities whose payments are linked to a mortality index).[7] The first, widely-marketed, bond of this type was issued by Swiss Re at the start of 2004. This involves a three-year contract (maturing on 1 January 2007) which allows the issuer to reduce its exposure to catastrophic mortality events. The repayment of the principal is linked to a combined mortality index of experienced mortality rates in five countries (France, Italy, Switzerland, the UK and the USA). Under this contract the principal will be at risk "if, during any single calendar year in the risk coverage period, the combined mortality index exceeds 130% of its baseline 2002 level". The credit spread at issue of 135 basis points equates to a risk-neutral probability of about 0.04 that the principal would not be repaid at all. This is equivalent to a catastrophic event that would happen, on average, once every 75 years (treating individual years as being independent). The types of catastrophic mortality events that are large enough to breach the threshold include a severe outbreak of influenza, a major terrorist attack (specifically the use of weapons of mass destruction), or a natural catastrophe. However, the Swiss Re bond addresses a different type of mortality risk (short-term catastrophic mortality risk) from that considered in this paper (unanticipated long-term changes in population mortality). The catastrophe risks being covered by the Swiss Re bond might be correlated with financial markets (past examples include 9/11 or the Kobe earthquake in 1995). In contrast, the systematic mortality risks we consider in this paper are assumed to be uncorrelated with the financial markets.

- Survivor swaps (where counterparties swap a fixed series of payments in return for a series of payments linked to the number of survivors in a given cohort).[8] The case for survivor swaps is made by Dowd *et al* (2004).

- Annuity futures (where prices are linked to a specified future market annuity rate).[9]

- Mortality options (a range of contracts with option characteristics whose payoff

---

[7]The contract specifications of short-dated, mortality-linked bonds are similar to traditional catastrophe bonds (see, for example, Schmock, 1999, Lane, 2000, Wang, 2002, and Muermann, 2004).

[8]A small number of survivor swaps have been arranged on an over-the-counter basis. They are not traded contracts and therefore only provide direct benefit to the counterparties in the transaction.

[9]As an example, suppose that $a(t, x)$ represents the market price at time t of a level annuity of £1 per annum payable monthly in arrears to a male aged x at time t. (This might, for example, be a weighted average of the top 5 prices in the market.) It is proposed that a traded futures market be set up with $a(t, x)$ as the underlying instrument for selected values of x and with a selection of maturity dates stretching out many years into the future. For a given maturity date, the market could be closed out some months or even a year before the maturity date itself, to reduce the impact, for example, of moral hazard, changes in expensing bases, or the movements of individual annuity providers in and out of the market.

depends on an underlying mortality table at the payment date). The guaranteed annuity contract mentioned above is an example of a mortality option, although it is really a complex option involving interest rate risks as well. Contracts of this type are discussed further in Section 6.2.

It is essential that the evolution of the prices of these derivative contracts should accurately reflect the stochastic evolution of $\mu(t, x)$.[10] The evolution of $\mu(t, x)$ can affect prices in two ways. Most obviously, stochastic mortality has an impact on the value of mortality options: the greater the volatility in mortality rates, the greater is the value of a mortality option (as with financial options). However, the second effect is more subtle and relates to the fact that the 'true' values of financial contracts are often non-linear functions of underlying factors. This point often manifests itself through Jensen's inequality, and an example in the present context would be that the price of a contract based on expected cash flow may not be equal to the value of the contract assuming that mortality follows some central projection. Also (in line with the pricing of financial options) it may manifest itself in our calculating expectations using a different probability measure (denoted below by Q) from the real-world or true measure (denoted by P). [11]

Given that mortality is best modelled as a stochastic variable, it is reasonable to suppose that any 'plausible' mortality model would meet the following criteria:

- The model should keep the force of mortality positive.

- The model should be consistent with historical data.

- The long-term future dynamics of the model should be *biologically* reasonable.[12]

---

[10] As an aside, it is important to note that the reference population underlying the calculation of the mortality rates is important to both the viability and liquidity of these contracts. Some investors (for example, life offices) will wish to use such contracts to help hedge their mortality risk, but if the reference population is inappropriate, they will be exposed to significant basis risk and the mortality derivative might not provide a good hedge. Other investors, including speculators and hedge funds, will be attracted to mortality-linked securities because their lack of correlation with other assets helps with the diversification of risk on a general portfolio of investments (see, for example, Cox, Fairchild and Pedersen, 2000). These investors may be less interested in using these derivatives for hedging mortality risk but will be interested in liquidity. Adequate liquidity will then require a small number of reference populations, but these will need to be chosen carefully to ensure that the level of basis risk is relatively small for those hoping to use the contracts for hedging purposes. For more detail, see Section 2.1.

[11] We do not discuss in this paper the many practical issues related to the securitization of mortality risks. These issues are discussed elsewhere (see Cummins, 2004, Dowd *et al.*, 2004, and Lin and Cox, 2004).

[12] For example, one might rule out the possibility of an 'inverted' mortality curve: that is, one in which mortality rates fall with age, in contrast to the normal upward slope that we have always observed in the past.

- Long-term deviations in mortality improvements from those anticipated should not be mean-reverting;[13] on the other hand, short-term deviations from the trend due to local environmental fluctuations may be mean-reverting.

- It should be possible to value the most common mortality-linked derivatives using analytical methods or using fast numerical methods.

- The model should be complex enough to deal appropriately with the current pricing, valuation or hedging problem.

Given this list of criteria, it is readily apparent that no one framework dominates the rest: some frameworks fare better by some criteria, and worse by others. There is therefore no strong reason why we should prefer models within one any modelling framework over models in another.

We described above a number of studies that propose specific stochastic models for the future evolution of mortality rates. This paper does not propose a new model. Instead, we seek to provide a general formulation of the problem in continuous time and a range of frameworks for the pricing and valuation of mortality-linked derivatives. Our aim in doing so is to provide extensive foundations for the development in the future of further stochastic models for mortality and to ensure that they are used in an appropriate way in pricing problems. The paper focuses for convenience on the problem of pricing of new securities, but the theory applies equally well to the fair valuation of insurance liabilities that incorporate mortality-linked derivatives.

The layout of the paper is as follows. In Section 2 we introduce the fundamental processes for mortality (the force of mortality process $\mu(t, x)$) and for the risk-free rate of interest $(r(t))$. These processes feed into survivor indices $S(u, y)$ and a risk-free cash account $C(t)$ that play central roles in our analysis. We work with two fundamental types of financial contract:

- pure endowment contracts for a full range of ages *and* terms to maturity; and

- default-free zero-coupon bonds for a full range of terms to maturity.

---

[13] Specifically we take the view that long-run stochastic improvements in mortality, $\mu(t, y)$, should not be mean-reverting to some deterministic projection, $\hat{\mu}(t, y)$. The inclusion of mean reversion would mean that if mortality improvements have been faster than anticipated in the past then the potential for further mortality improvements will be significantly reduced in the future. In extreme cases significant past mortality improvements might be reversed if the degree of mean reversion is too strong. This is clearly a very strong assumption that is difficult to justify on the basis of previous observed mortality changes and with reference to our perception of the timing and impact of, for example, future medical advances. Short-term trends might be detected by analysing carefully recent developments in healthcare and in the pharamaceutical industry, but even then the precise, long-term effects of such advances are difficult to judge. As we peer further into the future it becomes even more difficult to predict what medical advances there might be, when they will happen, and what impacts they will have on survival rates. All of these uncertainties rule out mean reversion in a model for stochastic mortality.

By noting parallels with interest-rate and credit-risk theory, we then describe how pure endowment contracts should be priced if they trade in a perfectly liquid, frictionless and arbitrage-free market.[14]

In Sections 3 to 6 we go on to describe different frameworks that could be employed to build up models for stochastic mortality. Section 3 pulls together all of the short-rate models described above under the one short-rate-modelling framework and discusses how these models can be used to build an arbitrage-free market in mortality derivatives. Sections 4 to 7 then discuss various other modelling frameworks, although to date no specific mortality models have yet been proposed using them.

Each of these frameworks is drawn from the field of interest-rate modelling but with the risk-free rate of interest replaced by the force of mortality. These are all described in theoretical terms: no specific models are proposed or analysed. Rather, the aim is to leave readers with a choice of frameworks within which they can build their own continuous-time stochastic mortality models. Most models built up within one framework can be reformulated within any of the other frameworks (in the same way, for example, that the Vasicek, 1977, short-rate model for $r(t)$ can be re-expressed as a forward-rate model). However, most models rest naturally within one framework, and are awkward to express within the others.

2 The term structure of mortality

In this section we will define the basic components of a model for stochastic mortality. We start by considering the force of mortality, $\mu(t, x)$, at time t for individuals aged x at time t. Traditional static mortality models implicitly assume that $\mu(t, x) \equiv \mu(x)$ for all t and x. Deterministic mortality projections imply that $\mu(t, x)$ is a deterministic function of t and x. By contrast, the models we will consider here will treat $\mu(t, x)$ as a stochastic process.

There are two types of stochastic mortality:

- The first is specific (or unsystematic) mortality risk – the risk that the actual numbers of deaths deviate from anticipated numbers because of the finite number of lives in a given cohort. This type of risk can largely be diversified by investors under the usual assumption that future lifetimes for different individuals are independent random variables.[15] Specific mortality risk therefore

---

[14]We do not claim that real-world markets are perfectly liquid or frictionless. However, we can state that if prices are calculated in the way proposed then even an illiquid market with frictions will be arbitrage-free. Conversely, if we were to propose a pricing framework which violates the conditions in Section 2, then the possibility of arbitrage would emerge over time as the market becomes more liquid or trading costs begin to fall.

[15]Strictly speaking there might be some local dependencies such as thosed between husband

does not lead to a risk premium in the price of mortality derivatives.

- Systematic mortality risk – the risk that the force of mortality evolves in a different way from that anticipated. This type of risk cannot be diversified away and therefore leads to the incorporation of a risk premium.

By drawing parallels with the pricing of financial contracts, we might expect that with mortality derivatives systematic mortality risk should be priced using a risk-neutral probability measure, Q, which is different from the real-world probability measure, P.[16] [17] [18] This intuition turns out to be correct: we will argue below that mortality derivatives need to be priced with reference to such a measure in order for the market to be arbitrage free.

2.1 Basic building blocks: the survivor index

We have previously indicated that our aim is to develop a set of theoretical frameworks to price mortality derivatives. In order to do so, we will make the convenient but over-simplifying assumption that the force of mortality at time t, $\mu(t, y)$, is observable at time t for all y. In reality, we can only estimate $\mu(t, y)$ from a finite amount of data, and this estimate is only calculated and published some months or years after the event. The length of this delay also depends considerably on the reference population: for example, the UK industry-wide Continuous Mortality Investigation tables take longer to compile than tables relating to one specific life office. We recognise that these are important practical issues but we will leave them for future work.

We will use as our basic building block a family of index-linked zero-coupon survivor bonds. The indices we will employ are related to survival probabilities for different ages. Thus we define the survivor index

$$S(u, y) = \exp\left(- \int_0^u \mu(t, y + t)dt \right). \tag{2.1}$$

and wife, or those between people who die in the same event: particularly the more-significant catastrophe risks of the type being covered by the Swiss Re mortality bond (including, for example, deaths cause by natural disasters or terrorist attacks).

[16] P is sometimes alternatively referred to as the true or objective or physical probability measure.

[17] In fact if we require our market in mortality-linked securities to be arbitrage free, even if this market is highly illiquid or has high transaction costs, the use of some risk-neutral measure Q is required of us as a consequence of the *Fundamental Theorem of Asset Pricing*. This states that if the market is arbitrage free then there exists a martingale measure that allows us to recover the market prices.

[18] One example of this is the Swiss Re bond. The 135 basis point spread equates to a risk-neutral probability, approximately, of 0.0135 per annum that the principal will not be paid out. The real-world probability might be rather less than this.

If $\mu(t, x)$ is deterministic then $S(u, y)$ is equal to the probability that an individual aged y at time 0 will survive to age $y + u$. Similarly, if $\mu(t, x)$ is deterministic, the probability that an individual aged x at time t_1 will survive until a later time t_2 is $S(t_2, x - t_1)/S(t_1, x - t_1)$.

In this paper we are mainly concerned with models in which $\mu(t, x)$ is stochastic. Looking forward from time 0, this means that $S(u, y)$ is now a random variable. In this case, $S(u, y)$ can still be regarded as a survival probability, albeit one that can only be observed at time u rather than at time 0. However, it is straightforward to extract a survival probability by taking the expectation of the random variable $S(t, x)$ (equation (2.2) below). We prove this by using a combination of indicator random variables and conditional expectation. Thus, consider an individual aged x at time 0. Let $Y_x(u)$ be a Markov chain which is equal to 1 if the individual is still alive at time u. Also let $\mathcal{M}_t$ be the filtration generated by the evolution of the term-structure of mortality, $\mu(u, x)$, up to time t.[19] The real-world or true survival probability[20] measured at time 0, that an individual aged x at time 0 survives until time u is

$$
\begin{aligned}
p_P(0, u, x) &= E_P[Y_x(u)] \\
&= E_P[E_P(Y_x(u)|\mathcal{M}_u)] \\
&= E_P[S(u, x)].
\end{aligned}
\tag{2.2}
$$

More generally we can define the survival probabilities at time t as follows. Let $p_P(t, u, x)$ be the probability under P that an individual aged x at time 0 and still alive at the current time t survives until time u:

$$
\begin{aligned}
p_P(t, u, x) &= E_P[Y_x(u)|Y_x(t) = 1, \mathcal{M}_t] \\
&= E_P\left[\left. \frac{S(u, x)}{S(t, x)} \right| \mathcal{M}_t\right].
\end{aligned}
$$

For the alternative risk-neutral probability measure Q, we can define the corresponding survival probabilities:

$$
\begin{aligned}
p_Q(t, u, x) &= E_Q[Y_x(u)|Y_x(t) = 1, \mathcal{M}_t] \\
&= E_Q\left[\left. \frac{S(u, x)}{S(t, x)} \right| \mathcal{M}_t\right].
\end{aligned}
$$

We are now in a position to consider the pricing of index-linked zero-coupon survivor bonds. There is (potentially) a different bond for each maturity date T and for each age x at time 0. We refer to a specific bond as the (T, x)-bond for compactness.

---

[19]In other words $\mathcal{M}_t$ gives us full information about changes in mortality up to and including time t, but no information about how mortality rates will develop after time t.

[20]Note that $E_P[S(u, x)] > \exp\left[-\int_0^u E_P(\mu(t, y + t))dt\right]$ by Jensen's inequality. Also, if $\bar{\mu}(t, y+t)$ is a deterministic, best estimate at time 0 of future mortality (for example, the median) then normally we will still have $E_P[S(u, x)] \neq \exp\left[-\int_0^u \bar{\mu}(t, y + t)dt\right]$.

The (T, x)-bond pays the amount $S(T, x)$ at time T. This payment is well defined in the sense that $S(T, x)$ is an observable quantity at time T. The (T, x)-bond is an example of what financial mathematicians call a *tradeable asset*[21]: that is, an asset that pays no coupons or dividends and whose price at any time $t < T$ represents the total return on an investment in that asset.[22]

To price such bonds we also need to make reference to the term-structure of interest rates. Let $P(t, T)$ represent the price at time t of a zero-coupon bond that pays 1 at time T. The instantaneous forward rate curve at time t is given by

$$f(t, T) = -\frac{\partial}{\partial T} \log P(t, T)$$

and the instantaneous risk-free rate of interest is

$$r(t) = \lim_{T \to t} f(t, T)$$

(see, for example, Cairns, 2004b). The cash (or money-market) account invests at the risk-free rate of interest. Its value at time t is denoted by $C(t)$ with

$$\begin{aligned} dC(t) &= r(t)C(t)dt \\ \Rightarrow C(t) &= C(0) \exp\left(\int_0^t r(u)du \right). \end{aligned} \quad (2.3)$$

Let $\mathcal{F}_t$ be the filtration generated by the term-structure of interest rates up to time t, and $\mathcal{H}_t$ be the combined filtration for both the term-structure of interest rates and mortality. If there exists a measure Q (the risk-neutral measure) equivalent to the real-world measure P with

$$P(t, T) = E_Q\left[\frac{C(t)}{C(T)} \,\middle|\, \mathcal{F}_t \right]$$

(which implies that $P(t, T)/C(t)$ is a Q-martingale) then the dynamics of the zero-coupon bond prices are arbitrage free.

Now let $\tilde{B}(t, T, x)$ represent the price at time t of the (T, x)-bond that pays $S(T, x)$ at time T. If there exists a measure Q equivalent to the real-world measure P with

$$\tilde{B}(t, T, x) = E_Q\left[\frac{C(t)}{C(T)} S(T, x) \,\middle|\, \mathcal{H}_t \right]$$

for all T and x then the dynamics of the index-linked zero-coupon bond prices are also arbitrage free. This formula matches those of Milevsky and Promislow (2001) and Dahl (2004) but encompasses a much wider range of models.

---

[21]To financial economists this would be more commonly known as a pure discount asset.

[22]For an asset that does pay dividends or coupons a tradeable asset can be created by reinvesting the dividends in the underlying asset itself.

Assumption 1

We now make the assumption that the dynamics of the term structure of mortality are independent of the dynamics of the term-structure of interest rates.

This assumption will allow us to separate the pricing of mortality risk from the pricing of interest-rate risk. It follows that

$$
\begin{aligned}
\tilde{B}(t,T,x) &= E_Q \left[\frac{C(t)}{C(T)} \;\middle|\; \mathcal{F}_t \right] E_Q \left[S(T,x) \mid \mathcal{M}_t \right] \\
&= P(t,T)B(t,T,x)
\end{aligned}
$$
$$
\text{where } \; B(t,T,x) = E_Q \left[S(T,x) \mid \mathcal{M}_t \right].
$$

Thus $B(t,T,x)$ is a martingale under Q. We can also assume that the $B(t,T,x)$ processes are strictly positive (barring the possibility of catastrophic events that wipe out the entire population).

This allows us to make three further observations.

- $B(t,T,x)/B(t,t,x) = p_Q(t,T,x)$. Since we can regard the $B(t,T,x)$ as spot prices we will refer to the $p_Q(t,T,x)$ as *spot survival probabilities*.

- We can use the $B(t,T,x)$ to define the *forward force of mortality surface* (we will sometimes shorten this to *forward mortality surface*):

$$
\bar{\mu}(t,T,x+T) = -\frac{\partial}{\partial T} \log B(t,T,x).
$$

  Conversely, knowledge of the forward mortality surface allows us to price the bonds:
$$
\frac{B(t,T,x)}{B(t,t,x)} = \exp \left[-\int_t^T \bar{\mu}(t,u,x+u)du \right].
$$
  If we take $T = t$, we get the *spot force of mortality*:
$$
\mu(t,x+t) = \bar{\mu}(t,t,x+t).
$$

- Let us assume that the dynamics of the term structure of mortality are governed by an n-dimensional Brownian motion $\tilde{W}(t)$ under Q. The martingale property of $B(t,T,x)$ together with its positivity allows us to write down the stochastic differential equation for $B(t,T,x)$ in the following form
$$
dB(t,T,x) = B(t,T,x)V(t,T,x)'d\tilde{W}(t)
$$
  where $V(t,T,x)$ is family of previsible[23] vector processes that specify the volatility term structure of bond prices.

---

[23]To describe a process, $X(t)$, as previsible means that the value of $X(t)$ is known or observable by time t.

We will now consider the possible frameworks[24] which we can use to model the dynamics of the $B(t, T, x)$ processes. These correspond to a variety of frameworks used in modelling interest rates (see, for example, Cairns, 2004b):

- short-rate modelling framework for the dynamics of $\mu(t, y)$ (which correspond to short-rate models for the risk-free rate of interest, $r(t)$, including those of Vasicek, 1977, Cox, Ingersoll and Ross, 1985, and Black and Karasinski, 1991);

- forward-mortality modelling framework for the dynamics of the forward mortality surface, $\bar{\mu}(t, T, x+T)$ (corresponding to the framework of Heath, Jarrow and Morton, 1992);

- positive-mortality modelling framework for the spot survival probabilities, $p_Q(t, T, x)$ (corresponding to the positive-interest framework developed by Flesaker and Hughston, 1996, Rogers, 1997, and Rutkowski, 1997);

- market modelling framework for forward survival probabilities or forward annuity prices (corresponding to the LIBOR and swap market models of Brace, Gatarek and Musiela, 1997, and Jamshidian, 1997).

Towards the end of the paper we also discuss the parallels between pricing mortality derivatives and credit risk. We note that there are many similarities which allow the transfer to our context of some intensity-based models that have been developed for pricing credit risk.

We stress that the purpose of the following sections is to expand the theoretical foundations of this relatively new field. We leave for further work the development of new models within the different frameworks considered here. We also leave for others the process of fitting these models to historical and market data.

3 Short-rate modelling framework

Models built up within this framework[25] specify directly the dynamics of $\mu(t, y)$. Existing models for the term-structure of mortality within this framework include Lee and Carter (1992), Lee (2000a), Yang (2001), Brouhns, Denuit and Vermunt (2002), and Renshaw and Haberman (2003) in discrete time, and Milevsky and

---

[24] In this paper we use the terms *framework* and *model* in a precise way. A framework is a general approach to modelling. We use the word model to mean a specific model developed within a given framework. For example, we have the Vasicek (1977), Cox, Ingersoll and Ross (1985), and Black and Karasinski (1991) *models* within the short-rate modelling *framework*. The framework normally defines what we can use as the basic building blocks in the modelling process.

[25] See footnote 24.

Promislow (2001) and Dahl (2004) in continuous time. In continuous time we model the force of interest which will have the general stochastic differential equation

$$d\mu(t,y) = a(t,y)dt + b(t,y)'d\tilde{W}(t) \tag{3.1}$$

where $a(t,y)$ and $b(t,y)$ (an $n \times 1$ vector) are previsible processes and $\tilde{W}(t)$ is a standard n-dimensional Brownian motion under the risk-neutral measure Q. We then have (see, for example, Milevsky and Promislow, 2001, or Dahl, 2004)[26]

$$\begin{aligned} \frac{B(t,T,x)}{B(t,t,x)} &= p_Q(t,T,x) \\ &= E_Q\left[\exp\left(-\int_t^T \mu(u, x+u)du\right) \,\Big|\, \mathcal{M}_t\right]. \end{aligned} \tag{3.2}$$

We can make the following observations about this framework:

- We have specified that $\tilde{W}(t)$ and $b(t,y)$ are $n \times 1$ vectors. This means that we can allow for the possibility that short-term changes in the term-structure of mortality can be different at different ages. Different rates of change at different ages can also be achieved through the $a(t,y)$ drift function.

- $a(t,y)$ and $b(t,y)$ might depend on other diffusion processes which are themselves adapted to $\mathcal{M}_t$. Note that this dependence allows $b(t,y) \equiv 0$, in which case the force of mortality curve evolves in a smooth fashion over time. However, the evolution of the force of mortality curve is still stochastic because of its dependence on the stochastic drift rate $a(t,y)$. Other models might assume that $b(t,y) \neq 0$, in which case the force of mortality curve exhibits a degree of local volatility.

- The assumption that $b(t,y) \equiv 0$ is equivalent to the assumption that the volatility function $V(t,T,x)$ for the $B(t,T,x)$ processes tends to zero as $T \to t$. Thus, the shortest-dated bonds will have a very low volatility.

- This framework includes models that assume that $\mu(t,y)$ takes some parametric form (for example, the Gompertz-Makeham model $\mu(t,x) = \xi_0(t) + \xi_1(t)e^{\xi_2(t)x}$). We can model the parameters in this curve as diffusion processes. This class is a specific example of the type noted above where $a(t,y)$ and $b(t,y)$ themselves depend on other diffusion processes.

The framework includes the affine class of models for $\mu(t,x)$ considered by Dahl (2004), under which the spot survival probabilities have the closed form

$$p_Q(t,T,x) = \exp\left[A_0(t,T,x) - A_1(t,T,x)\mu(t,x+t)\right]$$

---

[26]If we are modelling the spot survival probabilities, $p_Q(t, t+1, x)$, or the spot mortality rates, $q_Q(t, t+1, x) = 1 - p_Q(t, t+1, x)$ in discrete time, then the equivalent of equation (3.2) in discrete time is $p_Q(t,T,x) = E_Q\left[p_Q(t, t+1, x) \times \ldots \times p_Q(T-1, T, x+(T-1-t))|\mathcal{M}_t\right]$.

with $n = 1$ dimension. Dahl provides sufficient conditions on $a(t, y)$ and $b(t, y)$ (equation 3.1) that result in this affine representation for $p_Q(t, T, x)$. These conditions match those of Duffie and Kan (1996) for interest-rate models (see, also, Vasicek, 1977, and Cox, Ingersoll, and Ross, 1985). An important criterion of mortality models is that the spot survival probability function $p_Q(t, T, x)$ is decreasing in T – otherwise this would imply the potential for future negative mortality rates.

One potential drawback of this affine class is that the only models that ensure that $p_Q(t, T, x)$ is decreasing in T require the use of a mean-reverting process for $\mu(t, y)$. This mean reversion might be towards a time-dependent, but deterministic, local mean-reversion level, in which case mortality improvements can be systematically built into the model. However, mean-reverting mortality probabilities are questionable, as discussed previously in Footnote 13.

4 Forward mortality modelling framework

The next set of models are forward mortality models.

Suppose that we have the two stochastic differential equations:

$$dB(t, T, x) = B(t, T, x)V(t, T, x)'d\tilde{W}(t) \tag{4.1}$$

$$\text{and } d\bar{\mu}(t, T, x + T) = \alpha(t, T, x + T)dt + \beta(t, T, x + T)'d\tilde{W}(t) \tag{4.2}$$

where $V(t, T, x)$, $\alpha(t, T, x + T)$ and $\beta(t, T, x + T)$ are previsible processes. Now we might ask if we can specify $V(t, T, x)$, $\alpha(t, T, x+T)$ and $\beta(t, T, x+T)$ freely. However, by drawing parallels with the forward-interest-modelling framework of Heath, Jarrow and Morton (1992) (HJM), we can see that, in fact, there will need to be some form of relationship between $V(t, T, x)$, $\alpha(t, T, x+T)$ and $\beta(t, T, x+T)$ to be sure that we have an arbitrage-free framework for pricing mortality-linked derivatives. Before we develop the mathematical form of this relationship we can remark that the presence of age as an additional dimension means that our framework provides a richer and more complex modelling environment than does the classical HJM framework.

From equation (4.2) we have

$$\bar{\mu}(t,T,x+T) \;=\; \bar{\mu}(0,T,x+T)+\int_0^t \alpha(s,T,x+T)ds+\int_0^t \beta(s,T,x+T)'d\tilde{W}(s)$$

$$\Rightarrow \mu(t,x+t) \;=\; \bar{\mu}(0,t,x+t)+\int_0^t \alpha(s,t,x+t)ds+\int_0^t \beta(s,t,x+t)'d\tilde{W}(s)$$

$$\text{and }\; S(t,x) \;=\; \exp\left[-\int_0^t \mu(s,x+s)ds\right]$$

$$=\; \exp\left[-\int_0^t \bar{\mu}(0,u,x+u)du-\int_0^t\int_u^t \alpha(u,s,x+s)ds\,du\right.$$

$$\left.-\int_0^t\int_u^t \beta(u,s,x+s)'ds\,d\tilde{W}(u)\right].$$

Next note that

$$B(t,T,x) \;=\; S(t,x)\exp\left[-\int_t^T \bar{\mu}(t,s,x+s)ds\right]$$

$$=\; \exp\left[-\int_0^T \bar{\mu}(0,u,x+u)du-\int_0^t\int_u^T \alpha(u,s,x+s)ds\,du\right.$$

$$\left.-\int_0^t\int_u^T \beta(u,s,x+s)'ds\,d\tilde{W}(u)\right]. \qquad (4.3)$$

Now define $V(u,T,x)=-\int_u^T \beta(u,s,x+s)'ds$. We can then apply Ito's formula to $B(t,T,x)$ in equation (4.3) to get the SDE

$$dB(t,T,x) \;=\; B(t,T,x)\left[\left(\frac{1}{2}|V(t,T,x)|^2-\int_t^T \alpha(t,s,x+s)ds\right)dt\right.$$

$$\left.+V(t,T,x)'d\tilde{W}(t)\right].$$

(This confirms our earlier claim that $V(t,T,x)=-\int_u^T \beta(u,s,x+s)'ds$.)

Now we require the drift under Q to be zero. Therefore

$$\frac{1}{2}|V(t,T,x)|^2 \;=\; \int_t^T \alpha(t,s,x+s)ds$$

and by taking the partial derivative with respect to T we get

$$\alpha(t,T,x+T) \;=\; -V(t,T,x)'\beta(t,T,x+T).$$

As with the other frameworks, the challenge is to specify an appropriate form for $\beta(t,T,x+T)$ or $V(t,T,x)$. The chosen formulation needs to ensure that the forward mortality surface remains strictly positive. This requirement is most easily achieved by making $\beta(t,T,x+T)$ explicitly dependent on the current forward mortality surface. In addition, the chosen form needs to ensure that the spot force of mortality curve, $\mu(t,y)$, retains an appropriate shape (for example, that it is increasing with age).

5   The positive-mortality framework

We now turn to our third class of models, the positive mortality framework.

Let $\tilde{P}$ be some measure equivalent to Q, and let $A(t,x)$ be some family of $\mathcal{M}_t$ adapted, strictly-positive supermartingales.

Define

$$p_Q(t,T,x) = \frac{B(t,T,x)}{B(t,t,x)} = \frac{E_{\tilde{P}}[A(T,x)|\mathcal{M}_t]}{A(t,x)}. \tag{5.1}$$

The strict positivity of $A(t,x)$ means that $p_Q(t,T,x)$ is positive. The supermartingale property of $A(t,x)$ ensures that the $p_Q(t,T,x)$ are less than or equal to 1 and decreasing in $T > t$. It is straightforward to demonstrate (for example, through the application of the Radon-Nikodym derivative $dQ/d\tilde{P}$) that the resulting dynamics of $B(t,T,x)$ are appropriate for an arbitrage-free pricing model (see, also, Rogers, 1997, and Rutkowski, 1997). Within this pricing framework, the drift of $A(t,x)$ under $\tilde{P}$ is equal to $-\mu(t,x+t) \times A(t,x)$. (In the corresponding positive-interest model the drift of $A(t)$ is equal to $-r(t) \times A(t)$ – see, for example, Cairns, 2004b.)

Equation (5.1) appears deceptively simple as a pricing formula. However, the effort comes in specifying a model for the processes $A(t,x)$ and in calculating the expectations. (For examples in interest-rate modelling see Flesaker and Hughston, 1996, Rogers, 1997, and Cairns, 2004a.)

A special case of this framework is an adaptation of Flesaker and Hughston (1996) (FH). Let $N(t,s,x)$ for $0 < t < s$ be a family of strictly-positive martingales under $\tilde{P}$. Define

$$A(t,x) = \int_t^\infty N(t,s,x)ds.$$

The martingale property of $N(t,s,x)$ means that

$$E_{\tilde{P}}[A(T,x)|\mathcal{M}_t] = \int_T^\infty N(t,s,x)ds \tag{5.2}$$

$$< A(t,x). \tag{5.3}$$

It follows from (5.3) that $A(t,x)$ satisfies the Rogers/Rutkowski requirements for a strictly-positive supermartingale.

Combining equations (5.2) and (5.1) we now get

$$p_Q(t,T,x) = \frac{\int_T^\infty N(t,s,x)ds}{\int_t^\infty N(t,s,x)ds}.$$

From a computational point of view this involves, at worst, the numerical evaluation of a one-dimensional integral, no matter what the model for $N(t,s,x)$. Our problem is now one of devising an appropriate model for the family of martingales $N(t,s,x)$.

It is common in interest-rate-derivatives markets to calibrate the initial term structure of the model to the observed interest-rate term structure, and we can apply this approach to the mortality term structure too. Suppose then that we take as given at time 0 the market prices of the zero-coupon bonds, $P(0, T)$, and the (T, x)-bonds, $\tilde{B}(0, T, x)$, for all x and $T > 0$. From this we can derive the implied spot survival probabilities $p_Q(0, T, x) = \tilde{B}(0, T, x)/P(0, T)$. The initial values for the family $N(t, T, s)$ can then be calibrated as follows:

$$N(0, T, x) = -\frac{\partial}{\partial T} p_Q(0, T, x) = \bar{\mu}(0, T, x + T) p_Q(0, T, x).$$

This initial calibration is unique up to a strictly-positive, constant scaling factor.

By analogy with interest-rate modelling, this framework might contain natural model formulations that are difficult to identify in other frameworks. For example, the Cairns (2004a) interest-rate model can be reformulated as a short-rate model. However, the short-rate formulation is rather clumsy compared with the positive-interest formulation.

6   Mortality market modelling framework

We come now to the mortality market models, and begin with some preliminaries about the types of model covered by this framework.

As with the previous frameworks, market models are formulated in continuous time. However, in contrast to the previous frameworks, market models give us the dynamics for a restricted set of assets or forward rates (for example, the prices $\tilde{B}(t, T, x)$ for $T \in \{1, 2, 3, \dots\}$).[27]

Within an interest-rate context, one of the key steps in the development of a market model (see, for example, Brace, Gatarek and Musiela, 1997, and Jamshidian, 1997) is making a change from the risk-neutral probability measure Q to a suitable pricing measure. This is done by changing the numeraire from cash to a different tradeable asset. For example, with the LIBOR market model we use a zero-coupon bond, $P(t, T)$ as the numeraire. In this section we will discuss first a possible change of numeraire in the mortality-modelling context and then show how this change of numeraire can be used to develop mortality market models.

---

[27]We may *infer* the prices of other assets not explicitly modelled by interpolation. However, these inferences are not exact since the market is incomplete outside of the market variables that are explicitly modelled.

6.1 Introduction: change of numeraire

Recall that the processes $B(t, T, x)$ in a zero-interest-rate environment are martingales under Q with SDE's

$$dB(t, T, x) = B(t, T, x)V(t, T, x)'d\tilde{W}(t)$$

for appropriate previsible volatility functions $V(t, T, x)$.[28]

Now consider some, strictly-positive, tradeable assets as numeraires. As a specific first example consider $B(t, \tau, y)$ as the numeraire. We then consider processes of the type

$$Z(t, T, x) = \frac{B(t, T, x)}{B(t, \tau, y)}.$$

For most problems it is likely that the most productive choice of y will be x itself (since then $Z(\tau, \tau, x) = p_Q(\tau, T, x)$). If we then apply Ito's formula and the product rule we find that

$$dZ(t, T, x) = Z(t, T, x)\Big(V(t, T, x) - V(t, \tau, x)\Big)'\big(d\tilde{W}(t) - V(t, \tau, x)dt\big).$$

Now define a new process $W^{\tau,x}(t) = \tilde{W}(t) - \int_0^t V(s, \tau, x)ds$. Provided that $V(t, \tau, x)$ satisfies the Novikov condition we can use the Girsanov theorem (see, for example, Karatzas and Shreve, 1998) to infer that there exists a measure $P_{\tau,x}$ equivalent to Q under which $W^{\tau,x}(t)$ is a standard Brownian motion. In this case

$$dZ(t, T, x) = Z(t, T, x)\Big(V(t, T, x) - V(t, \tau, x)\Big)'dW^{\tau,x}(t),$$

so that $Z(t, T, x)$ is a martingale under $P_{\tau,x}$.

In what follows we will make more sophisticated choices for the numeraire, but the basic techniques described above will remain the same.

6.2 Annuity market models

The annuity market models that we will describe below (with and without stochastic interest) follow in the footsteps of Jamshidian (1997) in that they are formulated to provide a simple solution for a very specific type of contract. Jamshidian formulated his model specifically to provide a simple formula for the value of a swaption. Here we will derive a formula for the value of a simple form of guaranteed annuity option. The approach that we take here is also in the spirit of Pelsser (2003) who tackles the issue of pricing guaranteed annuity options. However, Pelsser concentrates on the inherent interest rate and equity risks, whereas we concentrate here on mortality and interest-rate risks.

---

[28]Readers who are familiar with interest-rate market models can consider the cash (or money-market) account (equation 2.3) as being the numeraire when pricing under Q. In this zero-interest-rate environment the cash account is equal to 1 for all time.

6.2.1 Zero interest rates

For simplicity we will restrict ourselves initially to a market where interest rates are set to zero: this will help clarify the basic argument. Non-zero stochastic interest rates will be added later.

Let

$$F(t,x) = \frac{B(t,T,x)}{\sum_{s=T+1}^{\infty} B(t,s,x)}$$

be a forward annuity rate under which survivors at T pay £1 at T and receive back $F(t,x)$ at times $T+1, T+2, \ldots$ so long as they are still alive at each of those dates.[29] In the assumed zero-interest-rate environment this contract has zero value at time t. Note specifically that $F(T,x) = 1/\sum_{s=T+1}^{\infty} p_Q(T,s,x)$ is the spot (market) annuity rate at T.

This suggests (by analogy with Brace, Gatarek and Musiela, 1997, and Jamshidian, 1997) the use of a different numeraire $X(t) = \sum_{s=T+1}^{\infty} B(t,s,x)$. Since $X(t)$ is a strictly-positive martingale we can write its SDE as

$$dX(t) = X(t)V_X(t)'d\tilde{W}(t)$$

for an appropriate previsible volatility function $V_X(t)$. Then

$$
\begin{aligned}
dF(t,x) &= F(t,x)\Big(V(t,T,x) - V_X(t)\Big)'\big(d\tilde{W}(t) - V_X(t)dt\big) \\
&= F(t,x)\gamma(t,x)'dW^X(t)
\end{aligned}
$$

where $\gamma(t,x) = \Big(V(t,T,x) - V_X(t)\Big)$ and $W^X(t) = \tilde{W}(t) - \int_0^t V_X(s)ds$ is a standard Brownian motion under an appropriate measure P_X equivalent to Q.

The standard modelling assumption for market models is to specify that $\gamma(t,x)$ is a deterministic function. It follows in this case that $F(s,x)$ for $t < s \leq T$ is log-normally distributed under P_X with

$$
\begin{aligned}
E_{P_X}[F(s,x)|\mathcal{M}_t] &= F(t,s) \\
\text{and } Var_{P_X}[\log F(s,x)|\mathcal{M}_t] &= \int_t^s |\gamma(u,x)|^2 du.
\end{aligned}
$$

Now consider an annuity contract that includes a guaranteed annuity rate. In the open market £1 at time T will secure a pension of $F(T,x)$ per annum from time T payable annually in arrears (assuming no expenses and a fair price). The contract

---

[29]Readers may be more familiar with a deferred annuity contract. Under this contract survivors at t pay £1 at t in return for a defined series of payments at times $T+1, T+2, \ldots$ payable only to those who are still alive at each of those dates. In contrast, with the forward contract, the purchase price is not paid until time T, and then only by those who are still alive at that time.

also includes a guarantee that the amount of the pension will be K per annum (the guaranteed annuity rate) if this rate is higher than the open market rate.

When we wish to value the guarantee we need to consider carefully the nominal amount being converted into an annuity at T. We claim that the appropriate amount is $S(T, x)$. To see why, suppose that we have a group of $N(t, x)$ lives at time t aged $x + t$. At time T, $N(T, x)$ of these individuals will still be alive. Suppose that each of these survivors will have available a nominal amount of £1 for conversion into an annuity at T. Then, given $\mathcal{M}_T$, $N(T, x)$ will have a binomial distribution with parameters $N(t, x)$ and $S(T, x)/S(t, x)$ and expected value $\kappa S(T, x)$ where $\kappa = N(t, x)/S(t, x)$. The argument is concluded with the developments leading up to equation (6.1) below where we see that $N(T, x)$ is conditionally independent of the mortality table in use at time T. This allows us to replace $N(T, x)$ by $\kappa S(T, s)$, thereby justifying our original claim that the appropriate amount is $S(T, x)$.

The total value of the contract at T is

$$N(T, x) \max\{F(T, x), K\} \sum_{s=T+1}^{\infty} p_Q(T, s, x).$$

The value of the option itself at T is therefore

$$G(T) = \frac{N(T, x)(K - F(T, x))_+}{F(T, x)}.$$

Now the option itself is a tradeable asset with price $G(t)$ at time t, so $G(t)/X(t)$ is a P_X martingale. Hence

$$
\begin{aligned}
\frac{G(t)}{X(t)} &= E_{P_X}\left[\frac{G(T)}{X(T)} \,\Big|\, \mathcal{M}_t\right] \\
&= E_{P_X}\left[\frac{N(T, x)(K - F(T, x))_+}{F(T, x)X(T)} \,\Big|\, \mathcal{M}_t\right] \\
&= E_{P_X}\left[E_{P_X}\left(\frac{N(T, x)(K - F(T, x))_+}{F(T, x)X(T)} \,\Big|\, \mathcal{M}_T\right) \,\Big|\, \mathcal{M}_t\right] \\
&= E_{P_X}\left[E_{P_X}\left(N(T, x) \mid \mathcal{M}_T\right) \frac{(K - F(T, x))_+}{F(T, x)X(T)} \,\Big|\, \mathcal{M}_t\right] \\
&= E_{P_X}\left[\kappa S(T, x) \frac{(K - F(T, x))_+}{F(T, x)X(T)} \,\Big|\, \mathcal{M}_t\right] \qquad (6.1) \\
&= E_{P_X}\left[\kappa(K - F(T, x))_+ \mid \mathcal{M}_t\right] \\
\Rightarrow G(t) &= \kappa X(t)\left(K\Phi(-d_2) - F(t, x)\Phi(-d_1)\right) \\
\text{where } d_1 &= \frac{\log F(t, x)/K + \frac{1}{2}\sigma_F^2}{\sigma_F} \\
d_2 &= d_1 - \sigma_F \\
\sigma_F^2 &= \int_t^T |\gamma(u, x)|^2 du
\end{aligned}
$$

and $\Phi(z)$ is a cumulative distribution function for the standard Normal distribution.

We have explained here how the swap market model can be adapted to mortality modelling. However, it remains for an empirical study to determine whether the assumption of a deterministic $\gamma(u, x)$ is reasonable or not.

6.2.2   Stochastic interest rates

Now consider the case with stochastic interest rates. In this case we have

$$
\begin{aligned}
dP(t, s) &= P(t, s)(r(t)dt + V_P(t, s)'d\tilde{Z}(t)) \\
dB(t, s, x) &= B(t, s, x)V_B(t, s, x)'d\tilde{W}(t)
\end{aligned}
$$

where $\tilde{Z}(t)$ and $\tilde{W}(t)$ are independent Brownian motions. Application of the product rule gives us

$$
d(P(t, s)B(t, s, x)) = P(t, s)B(t, s, x)(r(t)dt + V_P(t, s)d\tilde{Z}(t) + V_B(t, s, x)d\tilde{W}(t)).
$$

Now consider the annuity contract described above with a guaranteed annuity rate of K. The actual annuity rate at time T per £1 at T is $F(T, x)$ where

$$
F(t, x) = \frac{P(t, T)B(t, T, x)}{\sum_{s=T+1}^{\infty} P(T, s)B(T, s, x)}.
$$

(With $t = T$ this equates to $F(T, x) = 1/\sum_s P(T, s)p_Q(T, s, x)$.) This suggests the use of the numeraire

$$
X(t) = \sum_{s=T+1}^{\infty} P(T, s)B(T, s, x)
$$

with

$$
\begin{aligned}
dX(t) &= X(t)\left[r(t)dt + V_{PX}(t)d\tilde{Z}(t) + V_{BX}(t)d\tilde{W}(t)\right] \\
\text{where } V_{PX}(t) &= X(t)^{-1}\sum_{s=T+1}^{\infty} V_P(t, s)P(t, s)B(t, s, x) \\
\text{and } V_{BX}(t) &= X(t)^{-1}\sum_{s=T+1}^{\infty} V_B(t, s, x)P(t, s)B(t, s, x).
\end{aligned}
$$

Under the measure P_X, the prices of all tradeable assets discounted by $X(t)$ are martingales. Specifically this implies that $F(t, x)$ is a P_X-martingale with SDE under P_X

$$
dF(t, x) = F(t, x)[\gamma_P(t, x)dZ^X(t) + \gamma_B(t, x)dW^X(t)]
$$

for suitable previsible processes $\gamma_P(t, x)$ and $\gamma_B(t, x)$. In the annuity market model we assume that $\gamma_P(t, x)$ and $\gamma_B(t, x)$ are deterministic functions. As before we

assume that the nominal amount to be converted into an annuity at T is $S(T,x)$. (The argument in the previous section (6.2.1) converting actual numbers of lives surviving to T into $S(T,x)$ applies equally well here.) The value of the option component is denoted by $G(t)$ with

$$G(T) = \frac{S(T,x)(K - F(T,x))_+}{F(T,x)}.$$

The martingale property implies that

$$
\begin{aligned}
\frac{G(t)}{X(t)} &= E_{P_X}\left[\left.\frac{G(T)}{X(T)}\ \right|\ \mathcal{H}_t\right] \\
&= E_{P_X}\left[\left.\frac{S(T,x)(K - F(T,x))_+}{P(T,T)B(T,T,x)}\ \right|\ \mathcal{H}_t\right] \\
&= E_{P_X}\left[(K - F(T,x))_+ | \mathcal{H}_t\right].
\end{aligned}
$$

With the assumption that the volatility functions $\gamma_P(t,x)$ and $\gamma_B(t,x)$ are determinstic this gives us the pricing formula

$$
\begin{aligned}
G(t) &= X(t)\left(K\Phi(-d_2) - F(t,x)\Phi(-d_1)\right) \\
\text{where}\quad d_1 &= \frac{\log F(t,x)/K + \frac{1}{2}\sigma_F^2}{\sigma_F} \\
d_2 &= d_1 - \sigma_F \\
\text{and}\quad \sigma_F^2 &= \int_t^T (|\gamma_P(u,x)|^2 + |\gamma_B(u,x)|^2)du.
\end{aligned}
$$

It can be seen, therefore, that the annuity-market model offers a simple but powerful tool that can enable us to tackle some important questions involving annuity guarantees.

6.3   The SCOR market model

We will now consider a market model that looks directly at annualised forward mortality rates. This type of model is less tractable than the annuity market model (Section 6.2) if we use it to value annuity guarantees. On the other hand, this type of model can be applied much more easily to a wider class of product.

As before we will start by considering the situation where interest rates are equal to zero. We now define the concept of survival credits (previously utilised by Blake, Cairns and Dowd, 2003). These are, in effect, bonuses payable to survivors within a pool of life office policyholders in a way which ensures that no systematic profits or losses accrue to the life office. The survival credit payable to survivors at $t+1$ is calculated at time t by the life office based on the latest mortality tables available

at time t. In the event that actual survivorship from t to $t+1$ differs from that anticipated, the variation or risk over that year is borne by the life office.

The risk-neutral survival probability from t to $t+1$ measured at time t is $p_Q(t, t+1, x)$ and this implies that the actuarially and financial-economically fair survival credit payable at $t+1$ is

$$\frac{1 - p_Q(t, t+1, x)}{p_Q(t, t+1, x)}.$$

This represents a fair subdivision (as far as it can be anticipated at time t) of the amount invested at t by those who die before $t+1$ amongst those who survive to $t+1$.[30] It equals the odds at t of failing to survive to $t+1$ given survival to t.

This survivor credit is reminiscent of the τ-LIBOR (the London Interbank Offer Rate with duration or tenor τ) in the money markets which is equal (in a world with non-zero interest rates) to $L = (1 - P(t, t+\tau))/\tau P(t, t+\tau)$. The τ-LIBOR contract states that for each £1 deposited at t, £$1 + \tau L$ will be returned at $t + \tau$. For this reason we will refer to

$$L(T_{k-1}, T_{k-1}, T_k, x) = \frac{1 - p_Q(T_{k-1}, T_k, x)}{(T_k - T_{k-1})p_Q(T_{k-1}, T_k, x)} \tag{6.2}$$

as the *Survivor Credit Offer Rate* (or *SCOR*). In general we will assume that $T_k - T_{k-1} = 1$ for all k.[31]

Note also that we can rewrite (6.2) as

$$L(T_{k-1}, T_{k-1}, T_k, x) = \frac{B(T_{k-1}, T_{k-1}, x) - B(T_{k-1}, T_k, x)}{(T_k - T_{k-1})B(T_{k-1}, T_k, x)}.$$

This allows us to define the forward SCOR as follows

$$L(t, T_{k-1}, T_k, x) = \frac{B(t, T_{k-1}, x) - B(t, T_k, x)}{(T_k - T_{k-1})B(t, T_k, x)}. \tag{6.3}$$

Under a forward SCOR contract arranged at t we are fixing in advance the survivor credit that will be payable at T_k to survivors at T_k. That is for each £1 payable by survivors at T_{k-1}, those still alive at T_k will receive £$1 + (T_k - T_{k-1})L(t, T_{k-1}, T_k, x)$ at T_k. This will have zero value at t using the risk-neutral pricing approach discussed in Section 2.

---

[30]This is rather like a pool of annuitants as considered by Blake, Cairns and Dowd (2003), but here we are not making any assumptions about how much income is paid out of the fund to the survivors. However, Blake, Cairns and Dowd do not consider in detail the possibility of stochastic mortality.

[31]Note that, while $p_Q(T_{k-1}, T_k, x)$ must lie between 0 and 1, $L(T_{k-1}, T_{k-1}, T_k, x)$ can lie between 0 and ∞. This means that we can model $L(T_{k-1}, T_{k-1}, T_k, x)$, if we so choose, as a log-normal random variable.

For simplicity of notation in what follows, let us assume that $T_k - T_{k-1} = 1$ for all k and denote

$$L_k(t) \equiv L(t, T_{k-1}, T_k, x).$$

From equation (6.3) we see that $L_k(t)$ is equal to the value of a tradeable asset or portfolio $(B(t, T_{k-1}, x) - B(t, T_k, x))$ with the tradeable asset $B(t, T_k, x)$ as the numeraire.[32] As noted at the start of this section on market models, this implies that there exists a measure P_{T_k} equivalent to Q under which the prices of all tradeable assets divided by $B(t, T_k, x)$ are martingales and under which $W^{T_k}(t) = \tilde{W}(t) - \int_0^t V_B(u, T_k, x)du$ is a standard Brownian motion.

Application of the product rule to $L_k(t)$ (following a similar argument in Cairns, 2004b, Section 9.1) gives us

$$
\begin{aligned}
dL_k(t) &= \frac{B(t, T_{k-1}, x)}{B(t, T_k, x)} (V_B(t, T_{k-1}, x) - V_B(t, T_k, x))' \left\{ d\tilde{W}(t) - V_B(t, T_k, x)dt \right\} \\
&= L_k(t)V_L(t)'dW^{T_k}(t) \tag{6.4}
\end{aligned}
$$

where

$$
\begin{aligned}
W^{T_k}(t) &= \tilde{W}(t) - \int_0^t V_B(u, T_k, x)du \\
\text{and } V_{Lk}(t) &\equiv V_L(t, T_{k-1}, T_k, x) \\
&= (V_B(t, T_{k-1}, x) - V_B(t, T_k, x)) \frac{(1 + L_k(t))}{L_k(t)}. \tag{6.5}
\end{aligned}
$$

With reference to equation (6.4), first we note that the martingale property implies that $E_{P_{T_k}}[L_k(u)|\mathcal{M}_t] = L_k(t)$ for $t < u < T_{k-1}$. Second, if we make the usual market model assumption that $V_{Lk}(t)$ is a deterministic function then $L_k(u)$, given $\mathcal{M}_t$ for $t < u < T_{k-1}$, is log-normal under P_{T_k} with $Var_{P_{T_k}}[\log L_k(u)|\mathcal{M}_t] = \int_t^u |V_{Lk}(s)|^2 ds$.

Equation (6.5) can be rearranged to give

$$V_B(t, T_{k-1}, x) - V_B(t, T_k, x) = \frac{L_k(t)}{1 + L_k(t)} V_{Lk}(t). \tag{6.6}$$

Bearing in mind the relationship between the $W^{T_k}(t)$ and $\tilde{W}(t)$, for $l > k$ we can use (6.6) to show that

$$dW^{T_l}(t) = dW^{T_k}(t) + \sum_{j=k+1}^{l} \frac{L_j(t)}{1 + L_j(t)} V_{Lj}(t)dt.$$

---

[32]The fact that $L_k(t)$ can be expressed as the ratio of the prices of two tradeable assets is central to the no-arbitrage arguments. If either quantity was not a tradeable asset then the arguments would break down.

Thus, expressing the dynamics under P_{T_1}, we have

$$dL_1(t) \;\; = \;\; L_1(t)V_{L1}(t)dW^{T_1}(t)$$

and for $k > 1$

$$dL_k(t) \;\; = \;\; L_k(t)V_{Lk}(t)\left(dW^{T_1}(t) + \sum_{j=2}^{k}\frac{L_j(t)}{1+L_j(t)}V_{Lj}(t)dt\right).$$

These equations can be used as the basis for simulation of the $L_k(t)$ in discrete time. In addition we can simulate under the real world measure P by replacing $dW^{T_1}(t)$ by $dW(t) + \lambda(t)dt$ for a suitable process $\lambda(t)$.

Once again, this class of model offers us a powerful toolkit. However, we need to test potential models against historical data. At the same time, there are challenges in finding a suitable market-price-of-risk process $\lambda(t)$ that is statistically justifiable and that leaves the model reasonably tractable.

7   Credit risk modelling framework

Finally, we consider credit-risk models.

To start, we note that the zero-coupon survivor bond with price $\tilde{B}(t, T, x)$ at time t is similar to a zero-coupon corporate bond that pays 1 at T if there has been no default and 0 if the bond has defaulted. There are many models that address the problem of how to price such bonds (see, for example, the textbooks by Schönbucher, 2003, or Lando, 2004). In the present context, the most useful models for default risk that could be translated into a stochastic mortality model are *intensity-based* models (see, for example, Schönbucher, 2003, Chapter 7). In these models the default intensity, $\lambda(t)$ corresponds to the force of mortality $\mu(t, x+t)$. Thus from the theoretical point of view pricing can be approached in the same way.

However, there are some differences between mortality risk and credit risk that will be reflected in the type of model used:

- In a credit risk context different companies are equivalent to different cohorts in the mortality model. In mortality modelling there is a natural ordering of the different cohorts (that is, by current age) with strong correlations between adjacent cohorts. In contrast, there is no natural ordering of individual companies and, although defaults may be correlated, the additional structure in a mortality model will be absent.

- The default intensity in a credit model is likely to be modelled as a mean-reverting process that is also possibly time-homogeneous. In contrast, mor-

tality models are certainly time inhomogeneous and need to incorporate non-mean-reverting elements. This has the important implication that Cox, Ingersoll and Ross (1985)-type models can be used for credit-risk models, but not for mortality-based models.

- The default intensity is likely to be correlated with the interest-rate term structure, whereas the force of mortality is unlikely to be.

We can conclude that credit-risk modelling does have something to offer us in the mortality context. However, we need to use models that reflect the differences described above, or else we must adapt suitable credit-risk models to handle mortality risk.

8   Conclusions

We have presented here a number of theoretical frameworks that could be used for pricing many different types of mortality derivatives. More specifically, these frameworks demonstrate how an arbitrage-free (or risk-neutral) valuation methodology can be used to price a great variety of mortality-related instruments, from vanilla survivor bonds to exotic mortality derivatives. These frameworks build on those already established for conventional interest-rate derivatives, and therefore help to show how the existing literature on the valuation of interest-rate derivatives can be adapted to the valuation of derivatives involving mortality risk factors. They also provide a basis for the future development of specific stochastic mortality models, which can be developed within the frameworks offered here.

Our valuation methodology is of course the same, in principle, as that used to price derivatives based on other underlyings. As such, it is also open to exactly the same arguments over its merits. In particular, the use of arbitrage-free methods is always problematic if markets are incomplete, as is certainly the case with mortality derivatives markets. However, risk-free approaches provide natural benchmark valuations, and one can also argue heuristically that they will become easier to justify in any given context as markets become less incomplete over time. We also recognize that many of the assumptions underpinning these frameworks (such as liquid, frictionless markets) do not hold in practice. Nevertheless, it is still the case that, in such imperfect markets, *if* prices evolve in the way suggested by these pricing frameworks, *then* the model will be arbitrage free. We are not assuming that the market must be complete, or that transactions costs must be zero or that assets are infinitely divisible, and so on.

Needless to say, many challenges remain for future work. These challenges centre around five main issues:

- Models: We need to investigate which models give adequate statistical descriptions of historical mortality data, and which models generate 'reasonable' mortality dynamics.

- The number of risk factors: We need to determine how many risk factors are needed to provide a satisfactory model of the mortality term structure. Is one risk factor adequate, or do we need to have two or more factors to accommodate imperfectly correlated mortality improvements at different ages? One can easily argue that multiple factors might be desirable because the factors that affect mortality (medical advances, outbreaks of disease, etc.) are likely to have different impacts on people of different ages.

- Liquidity vs. basis risk: As with other derivatives, there is the usual tradeoff between basis risk and liquidity. If we want to encourage the development of market liquidity, then we would encourage standardized mortality instruments trading on organized exchanges; however, such instruments will embody considerable basis risk, which will reduce their usefulness as hedge instruments. Conversely, the basis risk for the primary client associated with tailor-made (or OTC) instruments will be low, but at the cost of poor liquidity. However, there are also other factors to consider. For example, the trading of mortality derivatives can be encouraged by choosing 'good' reference populations for the mortality indices (e.g., a good reference population might one be typical of the population as a whole, or typical of the annuitant population). Furthermore, the market for mortality derivatives might be helped by the fact that mortality risks have low market beta, and the willingness of investors to buy up the Swiss Re mortality-related bond is certainly encouraging.

- Index specification and moral hazard: The choice of index also needs to take account of possible moral hazard. For example, can the index be manipulated by the issuer of a security or by investors? Issuers of mortality-linked securities can learn from past experience of cat bond issuance, where moral hazard is always present to some degree.[33] A carefully-chosen index will not only increase liquidity, as discussed above, but it will also reduce moral hazard. A consequence, though, is that this reduction is typically accompanied by increased basis risk for those wishing to use the security for hedging. With cat bonds, one of the principal ways to reduce moral hazard is to link the payments to a non-company-specific index rather than to the ceding company's experience. The same arguments apply to mortality-related securities, although, perhaps, to a lesser degree.[34]

---

[33]For a discussion of moral hazard as it relates to the issue of cat bonds, see, for example, the papers by Doherty (1997), and Doherty and Richter (2002).

[34]This issue has already been address in part, by choosing in this paper to link payments to the survivor index (equation 2.1) rather than to the actual number of survivors in a particular cohort.

- The measurement-publication lag: How do we allow for the time lag between the measurement date and the date when mortality rates for that date have been graduated and made public? For example, is there something that can be learned here from catastrophe derivatives, where information gradually emerges after a catastrophic event? Connected to this issue is the need to specify the contract in a way that minimises insider-trading moral hazard associated with time lags in the release of information. This is a critical issue, not least because some early attempts to introduce non-life-insurance-linked catastrophe derivatives failed on precisely this point.[35]

In summary, mortality risks are an important new frontier in quantitative financial risk. They offer intellectual challenges to those who study them and serious financial benefits to those who trade them – and one would hope that they might also bring significant benefits to customers down the line.

Acknowledgements

References

Black, F., and Karasinski, P. (1991) "Bond and option pricing when short rates are log-normal," Financial Analysts Journal July-Aug, 52-59.

Blake, D., Cairns, A.J.G., and Dowd, K. (2003) "Pensionmetrics II: Stochastic pension plan design during the distribution phase," Insurance: Mathematics and Economics 33, 29-47.

Blake, D., and Burrows, W. (2001) "Survivor bonds: Helping to hedge mortality risk," Journal of Risk and Insurance 68, 339-348.

Brace, A., Gatarek, D., and Musiela, M. (1997) "The market model of interest-rate dynamics," Mathematical Finance 7, 127-155.

Brigo, D., and Mercurio, F. (2001) *Interest Rate Models: Theory and Practice.* Springer: Berlin.

Brouhns, N., Denuit, M., and Vermunt J.K. (2002) "A Poisson log-bilinear regression approach to the construction of projected life tables," Insurance: Mathematics

---

[35]We note though, that the Swiss Re mortality-related bond was issued successfully, implying that investors must have been reassured on the issue of moral hazard.

and Economics 31, 373-393.

Cairns, A.J.G. (2004a) "A family of term-structure models for long-term risk management and derivative pricing," Mathematical Finance To appear.

Cairns, A.J.G. (2004b) *Interest-Rate Models: An Introduction.* Princeton University Press: Princeton.

CMI (1978) "Proposed standard tables for life office pensioners and annuitants," Continuous Mortality Investigation Reports 3, 1-30.

CMI (1999) "Projection factors for mortality improvement," Continuous Mortality Investigation Reports 17, 89-108.

Cox, S.H., Fairchild, J.R., and Pedersen, H.W. (2000) "Economic aspects of securitization of risk," ASTIN Bulletin 30, 157-193.

Cox, J., Ingersoll, J., and Ross, S. (1985) "A theory of the term-structure of interest rates," Econometrica 53, 385-408.

Cummins, J.D. (2004) "Securitization of life insurance assets and liabilities," Working Paper, The Wharton School.

Currie I.D., Durban, M. and Eilers, P.H.C. (2004) "Smoothing and forecasting mortality rates," Working Paper, Heriot-Watt University.

Dahl, M. (2004) "Stochastic mortality in life insurance: Market reserves and mortality-linked insurance contracts," Insurance: Mathematics and Economics 35, 113-136.

Doherty, N.A. (1997) "Innovations in managing catastrophe risk," The Journal of Risk and Insurance 64, 713-718.

Doherty, N.A., and Richter, A. (2002) "Moral hazard, basis risk, and gap insurance," The Journal of Risk and Insurance 69, 9-24.

Dowd, K., Blake, D., Cairns, A.J.G., and Dawson, P. (2004) "Survivor swaps," In preparation.

Duffie, D., and Kan, R. (1996) "A yield-factor model of interest rates," Mathematical Finance 6, 379-406.

Flesaker, B., and Hughston, L.P. (1996) "Positive interest," Risk 9(1), 46-49.

Forfar, D.O., and Smith, D.M. (1987) "The changing shape of English life Tables," Transactions of the Faculty of Actuaries 40, 98-134.

Heath, D., Jarrow, R. and Morton, A. (1992) "Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation," Econometrica 60, 77-105.

James, J., and Webber, N. (2000) *Interest Rate Modelling.* Wiley: Chichester.

Jamshidian, F. (1997) "LIBOR and swap market models and measures," Finance and Stochastics 1, 293-330.

JMIC (1974) "Considerations affecting the preparation of standard tables of mortality," Journal of the Institute of Actuaries 101, 135-201.

Karatzas, I., and Shreve, S.E. (1998) *Methods of Mathematical Finance.* Springer-Verlag: New York.

Lando, D. (2004) *Credit Risk Modeling: Theory and Applications.* Princeton University Press: Princeton.

Lane, M.N. (2000) "Pricing risk transfer transactions," ASTIN Bulletin 30, 259-293.

Lee, P.J. (2000a) "A general framework for stochastic investigations of mortality and investment risks," Presented at the Wilkiefest, Heriot-Watt University, March 2000.

Lee, R. (2000b) "The Lee-Carter method for forecasting mortality, with various extensions and applications," North American Actuarial Journal 4, 80-93.

Lee, R.D., and Carter, L.R. (1992) "Modeling and forecasting U.S. mortality," Journal of the American Statistical Association 87, 659-675.

Lin, Y., and Cox, S.H. (2004) "Securitization of mortality risks in life annuities," Working paper, Georgia State University.

Macdonald, A.S., Cairns, A.J.G., Gwilt, P.L., and Miller, K.A. (1998) "An international comparison of recent trends in population mortality," British Actuarial Journal 4, 3-141.

Macdonald, A.S., Bartlett, D., Berman, C., Daykin, C., Grimshaw, D., Savill, P. (2003) "Mortality improvements and the cohort effect," CMI Working Papers 1 and 2; Presented to the Staple Inn Actuarial Society on 11 March 2003; 46 pages. (Available online at www.sias.org.uk.)

Milevsky, M.A., and Promislow, S.D. (2001) "Mortality derivatives and the option to annuitise," Insurance: Mathematics and Economics 29, 299-318.

Muermann, A. (2004) "Catastrophe derivatives," To appear in *The Encyclopaedia of Actuarial Science*, Wiley, Chichester.

Pelsser, A. (2003) "Pricing and hedging guaranteed annuity options via static option replication," Insurance: Mathematics and Economics 33, 283-296.

Rebonato, R. (2002) *Modern Pricing of Interest-Rate Derivatives.* Princeton University Press: Princeton.

Renshaw, A.E., and Haberman, S. (2003) "Lee-Carter mortality forecasting with age-specific enhancement," Insurance: Mathematics and Economics 33, 255-272.

Rogers, L.C.G. (1997) "The potential approach to the term-structure of interest rates and foreign exchange rates," Mathematical Finance 7, 157-164.

Rutkowski, M. (1997) "A note on the Flesaker & Hughston model of the term structure of interest rates," Applied Mathematical Finance 4, 151-163.

Schmock, U. (1999) "Estimating the value of the WINCAT coupons of the Wintherthur Insurance Convertible Bond: A study of the model risk," ASTIN Bulletin 29, 101-163.

Schönbucher, P. (2003) *Credit Derivatives Pricing Models: Models Pricing and Implementation.* Wiley: Chichester.

Vasicek, O. (1977) "An equilibrium characterisation of the term structure," Journal of Financial Economics 5, 177-188.

Wang, S.S. (1996) "Premium calculation by transforming the layer premium density," ASTIN Bulletin 26, 71-92.

Wang, S.S. (2000) "A class of distortion operations for pricing financial and insurance risks," Journal of Risk and Insurance 67, 15-36.

Wang, S.S. (2002) "A universal framework for pricing financial and insurance risks," ASTIN Bulletin 32, 213-234.

Wang, S.S. (2003) "Equilibrium pricing transforms: new results using Buhlmann's 1980 economic model," ASTIN Bulletin 33, 57-73.

Willetts, R.C. (1999) "Mortality in the next millenium," Presented to the Staple Inn Actuarial Society on 7 December 1999. (Available online at www.sias.org.uk.)

Yang, S. (2001) "Reserving, pricing and hedging for guaranteed annuity options," PhD Thesis, Heriot-Watt University, Edinburgh.

Appendices

A The Lee and Carter model for stochastic mortality

Lee and Carter (1992) investigate the dynamics of the observed central mortality rates $m(t, x)$ for integer t and x. Their model breaks $m(t, x)$ down into a log-bilinear model

$$\log m(t, x) = a(x) + b(x)k(t)$$

with the translation and scaling constraints that $\sum_x b(x) = 1$ and $\sum_{t=T_0}^{T_1} k(t) = 0$. $a(x)$ and $b(x)$ are non-parametric functions without any smoothing applied or functional form. Stage 1 of the modelling process estimates the functions $a(x)$, $b(x)$ and $k(t)$ without reference to a dynamic model for $k(t)$. For a given functional form for $a(x)$ and $b(x)$, $k(t)$ is estimated directly using the data for date t in isolation and without any assumption about its dynamic form. This is repeated until we optimise the fit over $a(x)$ and $b(x)$ subject to the scaling constraints above. Stage 2 of the modelling process then fits an ARIMA process to the $k(t)$.

B The Lee and Yang model for stochastic mortality

Lee (2000a) and Yang (2001) proposed the following model for stochastic mortality.

Suppose that a deterministic forecast of annual mortality rates is made at time 0. Thus $\hat{q}(x, t)$ represents the probability (as estimated at time 0) that an individual aged x at time t will die before time $t + 1$ for each integer x and t. The actual mortality experience is modelled as

$$q(x, t) = \hat{q}(x, t) \exp \left[X(t) - \frac{1}{2}\sigma_Y^2 + \sigma_Y Z_Y(t) \right]$$

$$\text{where } X(t) = X(t - 1) - \frac{1}{2}\sigma_X^2 + \sigma_X Z_X(t)$$

and $Z_X(t)$ and $Z_Y(t)$ are mutually independent sequences of i.i.d. standard normal random variables.

It follows that the $X(t)$ models the stochastic trend in the development of the mortality curve while the $-\frac{1}{2}\sigma_Y^2 + \sigma_Y Z_Y(t)$ models one-off environmental variations in mortality (such as a major flu epidemic). From the limited data available Yang found that σ_Y was not significantly different from 0.

C   The Milevsky and Promislow model for stochastic mortality

Milevsky and Promislow (2001) model the force of mortality in the form $\mu(t, x) = \xi_0 \exp(\xi_1 x + Y_t)$ where Y_t is an Ornstein-Uhlenbeck process with SDE $dY_t = -\alpha Y_t dt + \sigma dW_t$. Essentially this is equivalent to a Gompertz model with a time-varying scaling factor.

D   The Dahl model for stochastic mortality

Dahl (2004) models the process for $\mu(t, x + t)$ as follows

$$d\mu(t, x + t) = \alpha^\mu(t, x, \mu(t, x + t))dt + \sigma^\mu(t, x, \mu(t, x + t))d\tilde{W}(t).$$

He finds that if the drift and volatility are of the form

$$
\begin{aligned}
\alpha^\mu(t, x, \mu(t, x + t)) &= \delta^\alpha(t, x)\mu(t, x + t) + \zeta^\alpha(t, x) \\
\text{and } \sigma^\mu(t, x, \mu(t, x + t)) &= \sqrt{\delta^\sigma(t, x)\mu(t, x + t) + \zeta^\sigma(t, x)}
\end{aligned}
$$

for some deterministic functions $\delta^\alpha(t, x)$, $\delta^\sigma(t, x)$, $\zeta^\alpha(t, x)$ and $\zeta^\sigma(t, x)$ then

$$p_Q(t, T, x) = e^{A(t,T,x) - B(t,T,x)\mu(t,x+t)}$$

where the deterministic functions $A(t, T, x)$ and $B(t, T, x)$ are derived from differential equations involving $\delta^\alpha(t, x)$, $\delta^\sigma(t, x)$, $\zeta^\alpha(t, x)$ and $\zeta^\sigma(t, x)$.