Pricing Longevity Bonds using Implied Survival Probabilities*

Daniel Bauer†
DFG Research Training Group 1100, Ulm University
Helmholtzstraße 18, 89069 Ulm, Germany
Phone: +49 (731) 50 31088. Fax: +49 (731) 50 31239
Email: daniel.bauer@uni-ulm.de

Jochen Ruß
Institut für Finanz- und Aktuarwissenschaften
Helmholtzstraße 22, 89081 Ulm, Germany
Email: j.russ@ifa-ulm.de

September 2006

Abstract

For annuity providers, longevity risk, i.e. the risk that future mortality trends differ from those anticipated, constitutes an important risk factor. In order to manage this risk, new financial products will be needed. One of the basic building blocks for such mortality backed securities is the so-called survivor or longevity bond, the future payments of which depend on the survival rates of a certain population.

We propose a methodology for modeling and pricing of longevity bonds. We generalize the ideas of Lin and Cox (2005) and show how to derive *implied survival probabilities* from annuity market quotes. Taking those *implied survival probabilities* as a starting point, we derive the price and the dynamics of longevity bonds by applying the Heath-Jarrow-Morton framework for mortality modeling building on an idea proposed by Miltersen and Persson (2005).

We show how the models within our framework can be calibrated and applied for pricing mortality derivatives.

Keywords: stochastic mortality, forward force of mortality, Heath-Jarrow-Morton framework, implied survival probability.

*The authors thank Richard MacMinn and Hans-Joachim Zwiesler for useful comments as well as Jan Frederik Mai for his support.
†corresponding author

1 Introduction

Recently, it has become clear that mortality improvements occur in an unpredictable manner. The mortality or longevity risk, i.e. the risk that the realized future mortality trend differs from current assumptions, constitutes an important risk factor for insurers offering annuities (for an assessment of future mortality trends, see, e.g., Currie et al. (2004)). This risk is increased by the current problems of state-run pay-as-you-go pension schemes in many countries: The reduction of future benefits from public pension systems and tax incentives for annuitization of private wealth implemented by many governments will lead to an increasing demand for annuities. As a consequence, life insurers will face an increasing longevity risk. One of the most discussed ways of managing this risk is securitization, i.e. isolating the cash flows that are linked to longevity risk and repackaging them into cash flows that are traded in capital markets (see Cowley and Cummins (2005) for an overview of securitization in life insurance).

One of the basic instruments proposed in the academic literature for the securitization of longevity risk is the so-called survivor or longevity bond (see Blake and Burrows (2001), Blake et al. (2006) and references therein). The basic idea is that the coupon payments of such a bond depend on the survival of a certain cohort or population. So far, there have not been any transactions of survivor bonds, however, the European Investment Bank announced the issuance of a longevity bond in 2004. For a detailed description and a discussion, why such bonds were not offered successfully yet, we refer to Cairns et al. (2005a).

The pricing of such a contract is closely related to modeling the stochastic evolution of mortality: For a stochastic mortality model under the physical probability measure, a risk-adjusted price for mortality derivatives, such as longevity bonds, can be obtained by, e.g., the equivalent utility principle (see, e.g., Gerber and Pafumi (1998)). If mortality is modeled under a risk-neutral or risk-adjusted measure, prices of mortality derivatives are given by the expected value of the discounted cash flows (see, e.g., Dahl (2004)).

So far, several stochastic mortality models have been proposed – for a detailed overview and a categorization see Cairns et al. (2005b). Most of those stochastic mortality models are *short rate mortality models*, i.e. they model the spot mortality rate $q(t,x)$ or the spot force of mortality $\mu(t,x)$ (cf. Cairns et al. (2005b)). There are several discrete time models, for example Cairns et al. (2005c) or Renshaw and Haberman (2006), extending the ideas of Yang (2001) and of Lee and Carter (1992), respectively.

Milevsky and Promislow (2001) were the first to propose a stochastic hazard rate or force of mortality. In order to price guaranteed annuitization options in Variable Annuities, they show, first in a discrete framework, how a simple mortality option can be hedged using zero bonds, term life insurance contracts and endowment contracts.[1] Additionally, they price the same option in a continuous-time risk-neutral framework assuming that the short rate and the mortality intensity evolve independently over time according to a Cox-Ingersoll-Ross process and a stochastic Gompertz-type model, respectively. They present numerical results for the option prices.

Dahl (2004) presents a general stochastic model for the mortality intensity. He derives partial differential equations for the market reserve, i.e. the price at which some insurance contract could be sold on the financial market, and for general mortality derivatives in the presence of stochastic mortality. Furthermore, he shows how the systematic mortality risk can be transferred to the financial market. A

[1] However, their arguments depend on the insurer's ability to sell term life insurance contracts and annuity contracts for the same age group, which is rather unrealistic.

specification of the model with an affine term structure is employed in Dahl and Møller (2006). The authors derive risk-minimizing hedging strategies for insurance liabilities, however, in a market without mortality derivatives.

In Biffis (2005), affine jump-diffusion processes are employed to model both, financial and demographic risk factors. Closed form expressions (up to the solution of ordinary differential equations) are derived for numerous life insurance contracts. Furthermore, two specific models for the evolution of the mortality intensity are proposed: a Poisson-Gaussian process and a continuous bidimensional square-root process.

Schrager (2006) presents an affine stochastic mortality model, which simultaneously describes the evolution of mortality for different age groups. He fits a 10-factor model to Dutch mortality data using Kalmann filters and presents valuation approaches for several mortality-dependent contracts.

Most of the above models assume mean-reverting characteristics. However, Luciano and Vigna (2005) propose non mean-reverting affine processes for modeling the mortality intensity. They calibrate simple models to actual mortality data and claim to have a good match, in particular for models with a jump component.

A different approach is taken in Cairns et al. (2005b), which is based on an idea of Pelsser (2003). Here, the "Survivor Credit Offer Rate (SCOR)", i.e. the return of a one-year endowment contract in the case of survival, is modeled similarly to LIBOR market models (see, e.g., Bingham and Kiesel (2003), Chapter 8.5). The authors show how the forward SCOR can be computed and longevity bonds can be priced.

While all of the models so far assume independence of financial and demographic risk factors, Miltersen and Persson (2005) allow for correlations. Following the ideas of Heath-Jarrow-Morton (HJM) (see, e.g., Björk (1999), Chapter 23.1), they model the forward mortality intensity instead of the spot mortality intensity, taking the whole forward mortality curve as an infinite-dimensional state variable. They derive no arbitrage conditions for the drift term. Thus, a specification of the drift is not necessary.

In contrast to these "risk-neutral" approaches, Milevsky et al. (2005) postulate that an issuer of a mortality-contingent claim requires compensation for this risk according to a prespecified instantaneous sharpe ratio. They show that their price has several desirable properties. Furthermore, they obtain some new insights to traditional insurance pricing.

The approach of Lin and Cox (2005) differs substantially from the former models. Instead of using a stochastic mortality model to determine the price of a mortality bond, they use observed annuity prices to determine the *market price of mortality risk* as the shift parameter λ of a Wang transform (see Wang (2002)) of the best estimate mortality probabilities.

Even though we favor their idea of using market prices to derive fair prices for longevity bonds, this application of the Wang transform appears to be poorly motivated. Furthermore, it is not clear how prices or adjusted survival probabilities evolve over time and different age groups. For stochastic mortality models on the other hand, it is not clear how to calibrate them under a pricing measure or how to adjust historical estimates to allow for a risk premium.

The present paper fills this gap. We build on the ideas of Lin and Cox (2005) and propose extensions as well as generalizations to their model. In particular, we present a general no arbitrage framework and define *implied survival probabilities*, which can be used to price longevity bonds. Furthermore, we combine this model with ideas proposed in Miltersen and Persson (2005) in order to obtain a dynamic model for the evolution of mortality or, equivalently, survival probabilities under a risk-neutral measure.

The remainder of the paper is organized as follows. In Section 2.2, we introduce the basic definitions, summarize the ideas of Lin and Cox (2005), and point out some shortcomings of their model. In Section 3, we extend their model by proposing our general no arbitrage model. We show how it can be applied to derive *implied survival probabilities* from annuity prices. In Section 4, we combine these theoretical ideas with an idea proposed by Cairns et al. (2005b) and Miltersen and Persson (2005), namely to apply the HJM framework to mortality modeling, and explain how the resulting models can be calibrated. Section 5 discusses applications and extensions of the model. Finally, Section 6 closes with a summary of the main results and an outlook for further research.

2 Basic longevity bonds

This section establishes the definition of a longevity bond used in the remainder of the paper. Afterward, we summarize the ideas of Lin and Cox (2005), adjust their approach to our model, and highlight potential shortcomings of their approach.

2.1 Definitions

Following Cairns et al. (2005b), we use so-called (T, x_0)-bonds as basic building blocks. A (T, x_0)-bond is a financial security paying ${}_Tp_{x_0}$ at time T, where ${}_Tp_{x_0}$ denotes the realized proportion of x_0-year olds at time $t = 0$ who are still alive at time T.[2] However, from a practical point of view, this definition is problematic: In reality, at time T, ${}_Tp_{x_0}$ can only be approximated from a finite amount of data and this approximation will not be available until months or even years after time T (cf. Cairns et al. (2005b)). Furthermore, it is rather unrealistic to assume that there will be bond issues for all combinations of cohorts[3] and maturities. However, at least the latter is postulated in interest modeling as well. There, zero-coupon bonds for any maturity are the basic building blocks, even though some of them are hardly traded. Therefore, we rather consider these problems as implementation problems, which surely have to be analyzed, but should not affect the theory.

Under the assumptions that (T, x_0)-bonds exist for all ages and all maturities, and that the reference population coincides with the population of insured[4], it is straight forward how insurers can hedge their liabilities. For example, when selling annuities to x_0-year olds paying a total amount of $\$K$ annually in arrear, the insurer can hedge his position by buying K (t, x_0)-bonds for $t = 1, 2, 3, \ldots$.

2.2 The model of Lin and Cox (2005)

The price a party is willing to pay for a longevity bond or a similar mortality derivative depends on both, the best estimate of uncertain future mortality trends and on the quality or riskiness of these estimates. This risk induces a *mortality risk premium* that should be priced by the market (cf. Milevsky et al. (2005)). The same argument is made by Cairns et al. (2005c). In their discrete time two-factor model, the authors use a two dimensional shift parameter as the *market price of mortality risk* and calibrate it by employing published pricing data from the EIB/BNP longevity bond.

In order to account for a risk premium, Lin and Cox (2005) employ the Wang transform (cf. Wang (2002)) to adjust the best estimate survival probabilities for

[2]We will denote realized proportions by ${}_Tp_{x_0}$, whereas we use ${}_T\tilde{p}_{x_0}$ for the corresponding probabilities.

[3]For some value of $x_0 \in \{0, 1, 2, ..., \omega\}$, we denote all x_0-year olds in the reference population at time 0 as a cohort.

[4]This assumption enables us to consider only systematic mortality risk, whereas unsystematic and basis mortality risk (cf. Biffis et al. (2005)) are not considered.

the risk inherent in these estimates:[5]

$$_t\tilde{q}^{\lambda}_{x_0} := \Phi\left(\Phi^{-1}\left(_t\tilde{q}^{be}_{x_0}\right) - \lambda\right),$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function, $_t\tilde{q}_{x_0} := 1 -_t \tilde{p}_{x_0}$ is the t-year mortality probability for a x_0-year old, i.e. the probability for an individual aged x_0 to die within the next t years, and λ is the transformation parameter of the Wang transform. In order to distinguish between different probabilities under different measures, we let $_t\tilde{q}^{be}_{x_0}$ denote the best estimate mortality probabilities, and $_t\tilde{q}^{\lambda}_{x_0}$ the adjusted mortality probabilities.

Obviously, for each λ one obtains a different transform. In order to find a suitable transform, λ is derived from market prices of annuities: Using the current yield curve and adjusted survival probabilities, the hypothetical value ä of an annuity contract which pays an amount of K annually is derived as a function of the transform parameter λ, i.e.

$$\ddot{a}_{x_0}(\lambda) = K\sum_{t=1}^{\infty} {_t\tilde{p}^{\lambda}_{x_0}}\, P(0,t) = K\sum_{t=1}^{\infty}\left(1 - \Phi\left(\Phi^{-1}\left(_t\tilde{q}^{be}_{x_0}\right) - \lambda\right)\right) P(0,t).$$

Here, $P(0,t)$ denotes the current (time 0) value of a zero coupon bond with maturity t. Now, λ is determined by equating the hypothetical value and the actual market price of the annuity:

$$\ddot{a}_{x_0}(\lambda) \stackrel{!}{=} \text{Price of an annuity paying K.}$$

Then, $\Pi_0(T, x_0)$, the price of a (T, x_0)-bond at time $t = 0$, is simply given by the expected discounted value under the distorted probability measure, i.e.

$$\Pi_0(T, x_0) = P(0,T)\ {_T\tilde{p}^{\lambda}_{x_0}}. \tag{1}$$

We use synthetic[6] German annuity data, the current (2004) best estimate mortality tables given by the German society of actuaries as well as the German yield curve of December 02, 2005 (Bloomberg Data) for our calculations. We further choose $x_0 = 50$. Our calculations result in a Wang parameter of $\lambda \approx 0.42$, which is substantially higher than the one obtained by Lin and Cox (2005).[7] This indicates that the risk-premium for mortality risk is either higher in Germany or has changed over time, maybe influenced by mortality or interest rate developments.

As pointed out by Pelsser (2004), the Wang transform does not provide *a universal framework for pricing financial and insurance risks*, but basically is just one arbitrary possibility to distort the survival distribution. Figure 1 shows the survival function $_t\tilde{q}_{50}$ based on the best estimate DAV-aggregate mortality table as well as distortion of $_t\tilde{q}_{50}$ with respect to the Wang transform, i.e. with a distortion based on the normal distribution function. Furthermore, we give a distortion based on the Gamma distribution.[8] The structure implied by the Wang transform differs substantially from the distortion with respect to the Gamma distribution. In particular, since it is a positive distribution, the mortality rates are zero for the first

[5]Other authors also used the Wang transform to adjust survival probabilities for risk. See, e.g., Denuit (2005) or Dowd et al. (2006).

[6]The use of synthetic data does not cohere with the idea of taking market data for deriving prices. However, since this paper is intended to introduce the theoretical guideline on how to value mortality derivatives, we consider this choice sufficient.

[7]Lin and Cox (2005) derived a Wang parameter of about 0.18. However, for their calculations they used $x_0 = 65$ as the starting age. For $x_0 = 65$, we obtain a Wang parameter of 0.32, which still is significantly higher.

[8]Since the Gamma distribution is positive, the distortion has the form $\Gamma(a,b)\left\{\|\Gamma(a,b)^{-1}(\cdot) - \lambda\|\right\}$, where $\Gamma(a,b)$ denotes the cumulative distribution function of the Gamma distribution.

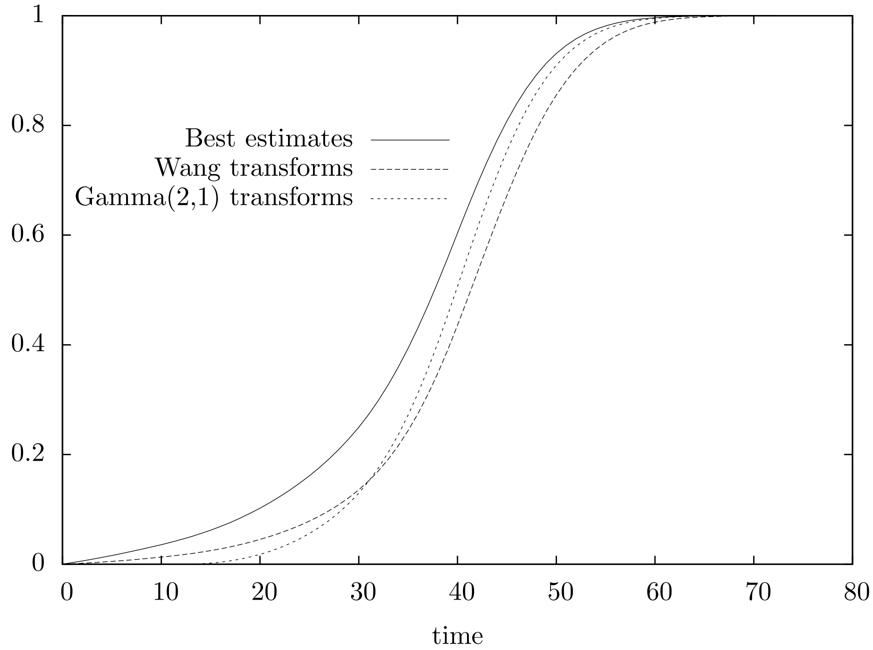

Figure 1: Comparison between different ${}_t\tilde{q}_{50}$

years, i.e. more people tend to survive to age 70 or 80 under this distortion. On the other side, the mortality rates implied by the Wang transform are lower in old ages. This indicates that the Wang transform might be the better choice, since estimates for close dates are probably more accurate and there is more uncertainty about mortalities in high ages. But, of course, this is no adequate motivation for the application of the Wang transform.

Another issue worth exploring is the question, under which conditions annuity prices offer an adequate starting point when pricing longevity bonds, i.e. to assess the relationship between annuity products and longevity bonds. Finally, as pointed out by Cairns et al. (2005b), within the model of Lin and Cox (2005), it is unclear *"how different transforms for different cohorts and terms to maturity relate one to another and form a coherent whole"*.

We will tackle all these issues in the subsequent sections. In particular, we provide a general no arbitrage model providing a theoretical foundation for the relationship between annuity prices and longevity bond prices. Furthermore, we will introduce *implied survival probabilities*, which are different from and provide a more coherent choice than Wang distorted probabilities. Finally, building on an idea introduced by Cairns et al. (2005b) and Miltersen and Persson (2005), we provide a dynamical model, which is able to explain the relationship between prices for different maturities.

3 A no arbitrage model for longevity bonds

This section introduces a model explaining the relationship between annuity prices and longevity bond prices. Therefore, it is necessary to assess the types of risk that insurers and longevity bond providers are exposed to. In the first subsection, we highlight the similarities and differences of the risks they face and, in particular, discuss potential implications for the corresponding risk premiums. Then, we present a no arbitrage model, which provides the theoretical foundation for the definition of *implied survival probabilities*.

3.1 The market price of mortality risk

Biffis et al. (2005) point our three sources of risk affecting insurance securities: basis risk, i.e. the risk that the reference population from which the survival probabilities were estimated differs from the insurer's cohort, systematic mortality risk, and unsystematic mortality risk. Mortality derivatives on the other hand are only exposed to systematic mortality risk. A natural question which arises from this classification is whether all these types of risk are incorporated in the pricing, i.e. whether there exists prices for basis risk, unsystematic, and systematic mortality risk, respectively. The first observation is that under the assumption that the reference population coincides with the underlying population of insureds, basis risk is of a rather unsystematic nature, since it is diversifiable, i.e. it vanishes with an increasing number of insureds in the insurer's portfolio. Therefore, we will consider basis risk as unsystematic mortality risk, and thus distinguish between systematic and unsystematic mortality risk only.

There seems to be no general agreement on whether unsystematic mortality risk is or should be priced. Clearly, for the insurer, this risk is diversifiable[9]. Therefore, in economic theories there is no charge for this risk.[10] There is another argument why this risk should not be priced: If there is a charge for unsystematic mortality risk, this charge will clearly depend on the number of insured in the insurer's portfolio. Thus, large insurers would be able to offer their products considerably cheaper than smaller companies, leading to a market with only large insurers. Since in many markets, there exists a rather large number of insurers, either there is no charge for unsystematic mortality risk, or all insurers are large enough to neglect this type of risk.

There are also several practical arguments why there is no charge for unsystematic mortality risk. Firstly, using the law of large numbers, i.e. diversifying their risk properly, is the core business of insurers, and therefore is probably managed in a sophisticated way. For example, insurers can reduce the variance by adequately marketing certain products or product lines. Also, taking the unsystematic risk can be interpreted as a re-compensation for managing their customers' money and participating in the corresponding returns.

Therefore, in what follows, we assume that unsystematic mortality risk is not priced. Conversely, there seems to be a broad agreement that there exists a market price for systematic mortality risk.[11] However, there seems to be no agreement on the structure and level of this price, and how it should be incorporated when valuating insurance products or mortality derivatives.

In order to solve this market price of mortality risk problem, different methods for pricing mortality risk have been proposed. In their discrete time model, Cairns et al. (2005c) define the market price of risk to be a shift parameter within their model and calibrate it to values of the EIB/BNP Longevity Bond. Milevsky et al. (2005) also argue that there should be an indemnification for taking on mortality or longevity risk. Their idea is to use a prespecified instantaneous Sharpe-ratio in order to compensate investors for holding systematic mortality risk. The model of Schrager (2006) also allows to choose market prices of risk such that the force

[9] Note that here we mean diversifiable in the sense of the law of large numbers and not in the sense of the CAPM.

[10] For example the CCAPM does not price diversifiable risk, see, e.g., Cochrane (2005).

[11] Therefore, in what follows, the term market price of mortality risk will refer to the market price of systematic mortality risk.

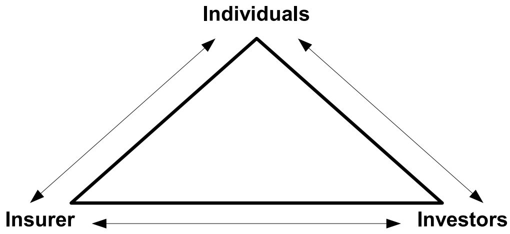

Figure 2: Market model

of mortality under the pricing measure is "prudent" compared to the real world measure. However, he does not show how to calibrate his model. Ballotta and Haberman (2006) assume that the market is completely risk-neutral with respect to mortality risk. However, their choice is rather motivated by the fact, that there is no adequate or standard way to assess the market's view on mortality risk.

All in all, the academic literature seems to agree that there exists a market price of mortality risk. The next subsection presents a method of assessing how this risk is priced by the market.

3.2 The no arbitrage model

We start by assuming that there are three types of market participants, namely an insurer selling endowment contracts[12], x_0-year old individuals who buy endowment contracts, and an investor selling longevity bonds (see Figure 2) in a liquid, frictionless market.[13] Furthermore, we assume a liquid secondary market for endowment policies. We further assume that, at time 0, the individuals buy K one-year endowment contracts with face value 1, where K is a sufficiently large natural number. For each contract, the insurer charges $P(0,1)\,{}_1\tilde{p}_{x_0}^{ins}$, where ${}_1\tilde{p}_{x_0}^{ins}$ denotes the risk-adjusted one year survival probability used by the insurer. Furthermore, in order to hedge the insurer's risk, the management of the insurance company decides to buy K $(1, x_0)$-bonds. In turn, the investor charges $P(0,1)\,{}_1\tilde{p}_{x_0}^{inv}$ per bond, where ${}_1\tilde{p}_{x_0}^{inv}$ is the risk-adjusted one year survival probability used for pricing the $(1, x_0)$-bonds. Thus, the insurer has a net cash flow of $K\,P(0,1)\,\left({}_1\tilde{p}_{x_0}^{ins} -_1 \tilde{p}_{x_0}^{inv}\right)$ at time zero.

At time one, the insurer has to pay out $K\,{}_1p_{x_0}$ to the individuals and receives $K\,{}_1p_{x_0}$, and therefore has a neutral position. Therefore, assuming an arbitrage-free market, we get ${}_1\tilde{p}_{x_0}^{ins} =_1 \tilde{p}_{x_0}^{inv}$:

- Assume ${}_1\tilde{p}_{x_0}^{ins} >_1 \tilde{p}_{x_0}^{inv}$. Then the insurer realizes $K\,P(0,1)\,\left({}_1\tilde{p}_{x_0}^{ins} -_1 \tilde{p}_{x_0}^{inv}\right)$ as a riskless gain at time zero.
- Assume ${}_1\tilde{p}_{x_0}^{ins} <_1 \tilde{p}_{x_0}^{inv}$. Then the investor buys the individuals' endowment contracts for $KP(0,T)\frac{{}_1\tilde{p}_{x_0}^{ins}+{}_1\tilde{p}_{x_0}^{inv}}{2}$ and sells the K longevity bonds. Thus, he has a neutral position at time 1 and both, him and the individuals have a free lunch at time 0.

Therefore, the survival probabilities used in the pricing of insurance products should also be used when pricing longevity bonds. This is in the spirit of Lin and Cox (2005), who used annuity data to calibrate their pricing model. Even though

[12]Endowment contracts are the basic building blocks for any annuity. Therefore, without loss of generality, we use endowment contracts in this subsection.

[13]Note that here as well as in the following, we ignore any charges for both, bond and insurance products.

some of the assumptions, in particular the assumption of a liquid secondary market for insurance policies may seem far fetched, there is evidence that large investors try to employ arbitrage possibilities in insurance markets. Insurance arbitrage markets are a profitable trend that Wall Street is eying.[14] For example Coventry First, a company specializing in the secondary market for insurance policies, announced in 2005 that they plan to buy Variable Annuities in the future if their intrinsic value exceeds the surrender value.

The next subsection shows, how *implied survival probabilities* can be defined and derived from market data in order to price longevity bonds consistent with our no arbitrage model.

3.3 Implied survival probabilities

In the previous subsection we derived that the adjusted survival probabilities to be used for pricing longevity bonds should equal the adjusted survival probabilities used for pricing annuity or endowment contracts. We will denote these survival probabilities by ${}_t\tilde{p}_{x_0}^{imp}$ and call them risk-neutral survival probabilities, since they lead to arbitrage-free prices in the sense of Section 3.2, or *implied survival probabilities*, since they are implied by the annuity market.

Of course, annuity prices from different providers may differ, for example due to credit risk reasons. However, in a competitive market, quotes offered by big, stable players should be close. Thus, when taking into account only prices which are not or only hardly affected by credit and similar types of risk, the implied survival probabilities are sufficiently well defined for all practical purposes.

Hence, similar to equation (1), we have for the price of a (T, x_0)-bond:

$$\Pi_0\left(T, x_0\right) = P\left(0, T\right)\, {}_T\tilde{p}_{x_0}^{imp}. \tag{2}$$

Using a synthetic insurance product as explained in Section 2.2, we can uniquely calculate all required annuity quotes. In practical applications, one would have to use market quotes. Practical issues of using interpolation methods where certain annuity quotes – especially for older ages – do not exist or, e.g., least square methods when quotes are not unique are beyond the scope of this paper. Therefore, we will leave them to future work.

We calculate the implied survival probabilities as follows: We are given the prices ${}_{|n}a_{x_0}$ for deferred annuities for x_0-year olds paying \$1 a year starting in $n = 0$ (immediate annuity), 1, 2, 3 etc. years. From Section 3.2, we know that the following relationship holds:

$${}_{|n}a_{x_0} = \sum_{t=n}^{\infty} P(0,t)\, {}_t\tilde{p}_{x_0}^{imp}.$$

Thus, we get ${}_n\tilde{p}_{x_0}^{imp} = \frac{{}_{|n}a_{x_0} - {}_{|n+1}a_{x_0}}{P(0,n)}$. Similarly, endowment policies or temporary annuities could be used.

Figure 3 shows best estimate survival probabilities, Wang-transformed survival probabilities from Section 2.2, and implied survival probabilities. We observe that, in comparison to implied survival probabilities, the Wang-transformed survival probabilities overestimate the survival probabilities in lower ages (up to the age of about 90 years),[15] i.e. less people die under the Wang-transformed probabilities over the first about 40 years. Conversely, the Wang-transformed survival probabilities are lower than the implied survival probabilities for high ages, i.e. the implied mortality rates ${}_t\tilde{q}_{x_0}^{imp} = 1 -_t \tilde{p}_{x_0}^{imp}$ are lower in old ages. With the same argument

[14] http://en.wikipedia.org/wiki/Arbitrage, March 2006.

[15] Note that in Figure 3, ${}_t\tilde{q}_{50} = 1 -_t \tilde{p}_{50}$ is plotted.

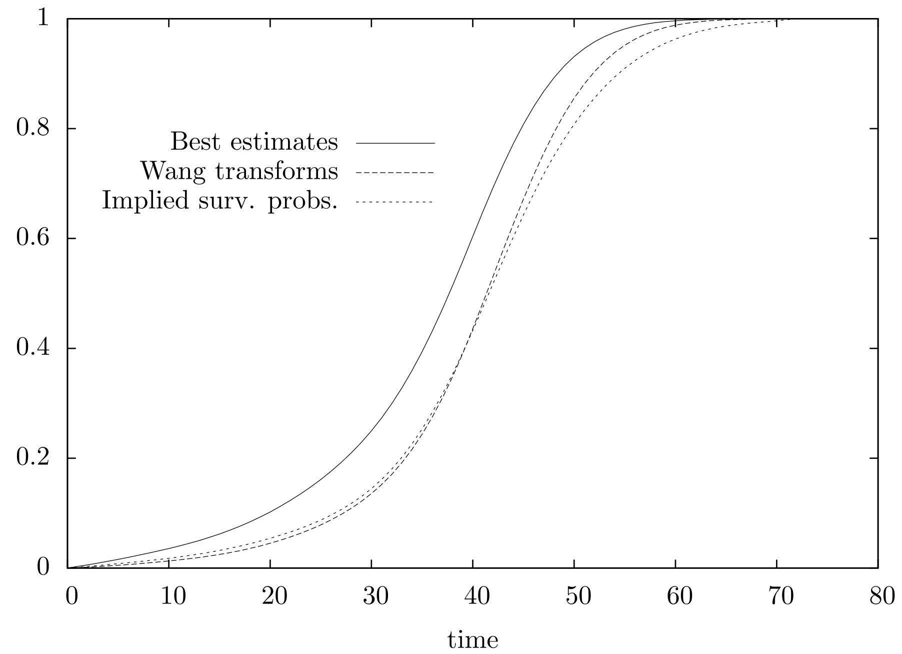

Figure 3: Comparison between implied and Wang ${}_t\tilde{q}_{50}$

that led to the conclusion that the Wang transform is more appropriate than the Gamma transform in Section 2.2, i.e. that estimates for close dates are probably more accurate and there is more uncertainty about mortalities in high ages, we can now conclude that the use of implied survival probabilities is more adequate than using the Wang transform. This is also consistent with the results of Cairns et al. (2005c), who find that longevity bonds with longer maturity and higher ages should ask for a higher risk premium.

Yet, this approach still lacks a model for the evolution of the longevity bond prices or, equivalently, a model for the evolution of the survival probabilities over time. However, when comparing our formula (2) to Lemma 3.1 in Bauer (2006)

$$\Pi_0\left(T, x_0\right) = P\left(0, T\right)\ E_{Q_T}\left[{}_T p_{x_0}\right]$$

we get
$$ {}_T\tilde{p}_{x_0}^{imp} = E_{Q_T}\left[{}_T p_{x_0}\right], \tag{3}$$

where $E_{Q_T}[\cdot]$ is the expectation under the T-forward neutral measure (see, e.g., Björk (1999), Section 24.1). This coheres with Milevsky and Promislow (2001), who pointed out that the technical rates used by companies' actuaries, which are closely related to our definitions of implied mortality rates, are actually forward rates. We will further explore this relationship and use it in order to obtain a model for the dynamics of survival probabilities over time in the next section.

4 The Heath-Jarrow-Morton approach for mortality modeling

The HJM framework for interest rate modeling, originally introduced by Heath et al. (1992), is based on modeling the dynamics of the forward rate curve. This framework can be interpreted as a generic description of the arbitrage-free movements of the forward curve, driven by Brownian motion.[16] As such, the HJM-framework basically unifies all continuous interest rate models (see Filipović (2001)). This is the key motivation for applying the HJM-framework to mortality modeling: Since no specific model is proposed, the question whether the application is adequate is obsolete.

The idea of applying the HJM-methodology to mortality modeling is not new (see Cairns et al. (2005b) and Miltersen and Persson (2005)). However, Cairns et al. (2005b) assume that the dynamics of the term structure of mortality and the dynamics of the term structure of interest rates are independent. We believe that this assumption is not adequate. As pointed out by Miltersen and Persson (2005), even though the death of an individual policy holder can for all practical matters be considered independent of the development of the financial market, there is evidence that there are significant interdependencies between the force of mortality in the future and the development of the economy in general, and thus also the development of the financial markets.[17]

Furthermore, it is widely accepted that the current increasing attention to stochastic mortality modeling is largely due to the low interest rate environment. When rates are low, the values of the liabilities of insurers increase and, thus, insurers become more concerned about their future liabilities (see, e.g., Dowd et al. (2006)). Furthermore, various options embedded in insurance products are mortality contingent. These are more valuable and therefore exercised more often when rates are low. Examples for such options are guaranteed annuity options (GAOs) or guaranteed minimum income benefits (GMIBs) within variable annuity contracts. Thus, when rates are low, insurers are stronger affected by the mortality-dependent obligations promised within those guarantees. Therefore, even if the development of future mortality and of the financial markets were uncorrelated under the physical probability measure, there still might be some dependence under the pricing measure.

Miltersen and Persson (2005) provide a model which allows for correlations of the dynamics of the term structure of mortality and the dynamics of the term structure of interest rates. They propose a generic model for the forward force of mortality and, in particular, derive an equivalent to the well-known HJM drift condition for mortality modeling (see, e.g., Björk (1999), Proposition 23.2). However, their calculations are taken out from an individual level, i.e. for insurance contracts, rather than for longevity bonds. Thus, in what follows, we follow Bauer (2006), who provides a model for an arbitrage-free family of longevity bonds, even though the results, and in particular the HJM drift condition, remain the same as in Miltersen and Persson (2005).

We will start out by introducing the forward force of mortality and provide a generic model for the corresponding dynamics. Subsequently, we will further explore the relationship (3) and apply it to obtain a generic model for pricing mortality derivatives. Finally, we discuss the implementation and parameter estimation of those models.

4.1 Modeling the forward force of mortality

In what follows, we assume an arbitrage-free market of both, regular bonds and longevity bonds as in Bauer (2006). In particular, we assume that all longevity bonds have the same reference cohort of x_0-year olds at time 0 and fixed inception time zero, i.e. we assume that all longevity bond payoffs are of the form ${}_tp_{x_0}$ at time t. Furthermore, we assume that all quantities are smooth and regular enough,

[16] However, there exist extensions to a more general classes of driving processes, e.g. Lévy-processes (see, e.g., Schoutens (2002)).

[17] See Miltersen and Persson (2005) for details.

such that all derivatives and integrals used in this section are well defined. For all technical details and proofs, we refer to Bauer (2006).

The forward force of mortality with maturity T as from time t is defined as

$$\tilde{\mu}_t\left(T, x_0\right) := -\frac{\partial}{\partial T} \log \left\{ \frac{\Pi_t\left(T, x_0\right)}{P(t,T)} \right\}$$

where $\Pi_t\left(T, x_0\right)$ is the value of a (T, x_0)-bond at time t. This implies (see Bauer (2006), Proposition 3.1)

$$\frac{\Pi_t\left(T, x_0\right)}{P(t,T)} = \frac{\Pi_t\left(S, x_0\right)}{P(t,S)} \exp \left\{ -\int_S^T \tilde{\mu}_t\left(u, x_0\right) du \right\},\ T \geq S \geq t. \tag{4}$$

Now, we further assume that, for every fixed $T > 0$, the forward forces[18] $f\left(\cdot, T\right)$ and $\tilde{\mu}_\cdot\left(T, x_0\right)$ have stochastic differentials which, under the risk-neutral measure Q are given by

$$\begin{aligned} d\, f\left(t,T\right) &= \alpha^F\left(t,T\right) dt + \sigma^F\left(t,T\right) dW(t), \\ f\left(0,T\right) &= f^*\left(0,T\right), & (5) \\ d\, \tilde{\mu}_t\left(T,x_0\right) &= \alpha^L\left(t,T\right) dt + \sigma^L\left(t,T\right) dW(t), \\ \tilde{\mu}_0\left(T,x_0\right) &= \tilde{\mu}_0^*\left(T,x_0\right), & (6) \end{aligned}$$

where W is a finite dimensional $Q-$Wiener process, and the integrands are adequately regular adapted processed. Since Q is a risk neutral measure, the time zero price of a longevity bond is given by

$$\Pi_0\left(T, x_0\right) = E_Q\left[\exp\left\{-\int_0^T r(s)\, ds\right\}\ {}_T p_{x_0}\right], \tag{7}$$

where r denotes the short rate of interest. Conversely, (4) implies

$$\Pi_0\left(T, x_0\right) = \exp\left\{-\int_0^T \tilde{\mu}_0\left(u, x_0\right) du\ -\ \int_0^T f\left(0, u\right) du\right\}. \tag{8}$$

In order for these formulas to hold simultaneously, we arrive at the HJM-drift condition for mortality modeling[19]

$$\begin{aligned} \alpha^L\left(t,T\right) &= \sigma^L\left(t,T\right)\left(\int_t^T \sigma^F\left(t,s\right)' + \sigma^L\left(t,s\right)' ds\right) \\ &\quad + \sigma^F\left(t,T\right)\left(\int_t^T \sigma^L\left(t,s\right)' ds\right). \end{aligned} \tag{9}$$

When pricing mortality derivatives, we have to additionally model the financial market. We assume that the forward rates evolve according to (5). Note, that for the interest rates, there is a similar drift condition, where the drift is specified by the volatility process $\sigma^F\left(t,T\right)$ (see, e.g., Björk (1999)). Thus, the only step remaining is a specification of $\sigma^F\left(t,T\right)$ and $\sigma^L\left(t,T\right)$. However, the model still needs to be calibrated. Besides calibrating the volatilities, which we will come back to in Section 4.3, we have to determine initial conditions for (5) and (6), respectively. The prior can be derived from the prevailing forward structure of interest rates as implied by the prevailing yield curve. The next subsections deal with the question of finding an initial condition for the latter.

[18]Here $f\left(t,T\right)$ denotes the instantaneous forward rate or forward force of interest with maturity T contracted at t, see, e.g., Björk (1999).

[19]Here, x' denotes the transposed of x.

4.2 Implied survival probabilities and the HJM approach

When changing from the risk-neutral measure Q to the T-forward neutral measure by a change of numéraire from the money account $B_t := \exp\left\{-\int_0^T r(s)\,ds\right\}$ to the T-Bond $p(t,T)$, from (7) we obtain

$$\Pi_0(T, x_0) = P(0,T)\ E_{Q_T}\left[{}_T p_{x_0}\right],$$

which, as already pointed out in (3), yields

$${}_T\tilde{p}_{x_0}^{imp} = E_{Q_T}\left[{}_T p_{x_0}\right].$$

Note that we did not make an independence assumption. Conversely, we still allow for dependence between the evolution of mortalities and interest rates. In particular, this dependence results from using the same Wiener process for the dynamics in (5) and (6).

Furthermore, we know that

$$P(0,T) = \exp\left\{-\int_0^T f(0,s)\ ds\right\}$$

as well as

$$\Pi_0(T, x_0) = \exp\left\{-\int_0^T \tilde{\mu}_0(u, x_0)\ du\ -\ \int_0^T f(0,u)\ du\right\}.$$

from (8), which, altogether, leads to

$${}_T\tilde{p}_{x_0}^{imp} = E_{Q_T}\left[{}_T p_{x_0}\right] = \exp\left\{-\int_0^T \tilde{\mu}_0(s, x_0)\ ds\right\}. \tag{10}$$

Thus, we have

$${}_T\tilde{p}_{x_0}^{imp} = \exp\left\{-\int_0^T \tilde{\mu}_0^*(s, x_0)\ ds\right\}.$$

Since these implied survival probabilities are given exogenously, we can use (10) to define an initial condition for (6).

It is worth noting that under this condition, the measure Q is risk-neutral in two senses: On the one hand it produces prices which are implied by an arbitrage free longevity bond market (drift condition (9)), and on the other hand it gives time-zero prices which are consistent with our no arbitrage model from Section 3.2.

So far, besides the generic dynamics from (5) and (6), we have not specified any particular model or parametrization. Thus, the foregoing could be understood as a framework rather than a particular model. Also, as pointed out earlier, relaxing the dynamics conditions to more general driving processes, as for example Lévy processes (see Schoutens (2002)), would not generally alter the framework. However, in order to apply the model, we need further specifications. The next subsection discusses this in more detail. In particular, we propose a method to obtain an initial forward mortality structure from (10) and explain, how a given volatility structure can be calibrated.

4.3 Practical implementation and parameter estimation

In order to use the above model, we need an initial condition for the dynamics in (6) and (5), i.e. we need continuous representations for $\tilde{\mu}_0(\cdot, x_0) : t \mapsto \tilde{\mu}_0(t, x_0)$

and $f(0,\cdot) : t \mapsto f(0,t)$, respectively.[20] Of course, we cannot uniquely estimate a continuous function from a given discrete amount of data. For determining the term structure of interest rates, a popular method to solve this problem has been proposed by Nelson and Siegel (1987). They derive the yield curve based on a given parameterical representation, i.e. they let

$$y(0,T) = y(0,T,\beta),$$

where $y(0,T)$ denotes the time 0 interest with respect to maturity T, β is a parameter vector, and $y(0,\cdot,\beta)$ is a given function with parameter set β. We will proceed analogously: We will employ the relationship (10) by considering a given parametric function for the mortality intensity and determining the parameters to fit the exogenously given data, i.e. the finite number of implied survival probabilities.

In Ballotta and Haberman (2006), a parametric representation for the force of mortality as employed by British actuaries is presented:

$$\begin{aligned} \tilde{\mu}_0(t,x_0) &= a_1 + a_2 R(x_0+t) + \\ &\quad \exp\left\{b_1 + b_2 R(x_0+t) + b_3\left(2\,R(x_0+t)^2 - 1\right)\right\} \end{aligned} \tag{11}$$

where

$$R(x) = \frac{x-70}{50},\ x \geq 50.$$

However, this functional representation does not match our data from Section 3.3 well. Therefore, we use the following modification to better match our data:

$$\begin{aligned} \tilde{\mu}_0(t,x_0) &= a_1 + a_2 R(x_0+t) + a_3 \ln\left\{\frac{x_0+t}{50}\right\} \\ &\quad \exp\left\{b_1 + b_2 R(x_0+t) + b_3\left(2\,R(x_0+t)^2 - 1\right)\right. \\ &\quad\quad \left.+ b_4\left(3R(x_0+t)^3 - 1\right)\right\}. \end{aligned} \tag{12}$$

In order to fit the data, we approximate the $\tilde{\mu}_0(k,x_0)$ by the respective one-year mortality probabilities: From

$$\begin{aligned} \tilde{p}^{imp}_{x_0+t} &= \exp\left\{\int_t^{t+1} -\tilde{\mu}_0(s,x_0)\,ds\right\} \\ \Rightarrow -\log\left\{\tilde{p}^{imp}_{x_0+t}\right\} &= \int_t^{t+1} \tilde{\mu}_0(s,x_0)\,ds, \end{aligned}$$

we get

$$\tilde{\mu}_0(t,x_0) \approx -\ln\left\{1 - q^{forward(0)}_{x_0+t}\right\}$$

as an approximation. An iterative fit of the data resulted in the parameter values displayed in Table 1.

In Figure 4, the fitted curves for both, the function presented in Ballotta and Haberman (2006) and our modification, as well as the data are plotted. Noticeably, the modification provides a better match.

The functional representation of the forward mortality intensity (12) enables us to compute the value of (T,x_0)-bonds for any maturity T for our given cohort of x_0-year olds:

$$\Pi_0(T,x_0) = P(0,T)\,\exp\left\{-\int_0^T \tilde{\mu}_0(s,x_0)\,ds\right\}.$$

[20] To simplify notation, we drop the $\cdot^*$ in what follows.

	Parameter value	Asymptotic standard error
a_1	-0.497846	0.06174
a_2	-0.258478	0.07212
a_3	1.13593	0.1518
b_1	-8.71617	0.8579
b_2	-4.65033	0.1517
b_3	-3.74638	0.3805
b_4	-2.85213	0.397

Table 1: fitted parameter values

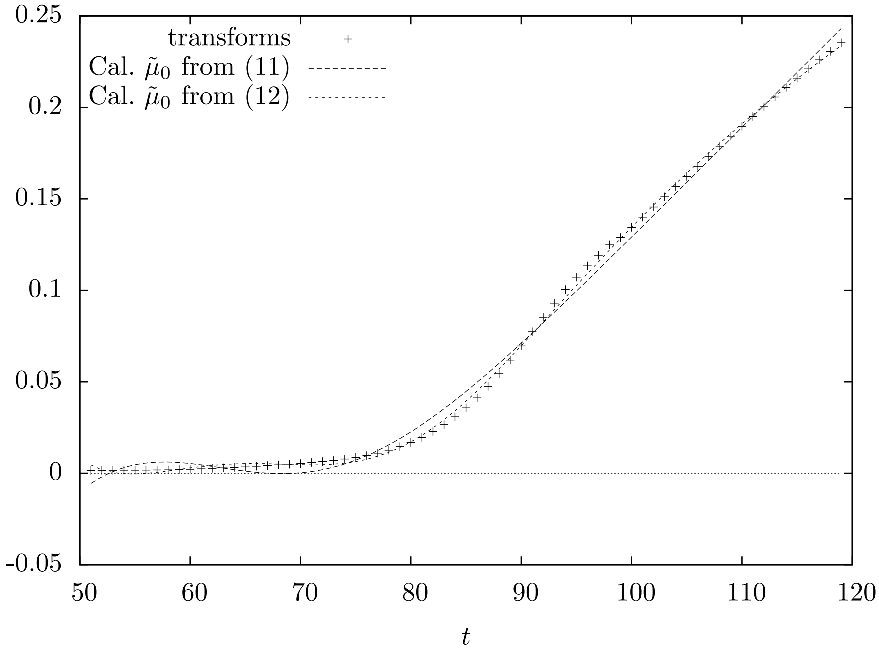

Figure 4: Parameterizations for $\tilde{\mu}_0\,(t, 50)$

Note that the parametrization in (12) may not be be a consistent parametrization in the sense of Filipović (2001). Furthermore, we do not claim to have a better match in general.[21] Also, the asymptotic errors are rather large, so the model might be over-parametrized. We leave the further exploration of this issue for future work.

Having specified the initial conditions for the dynamics, the only remaining problem is to determine and calibrate an adequate volatility structure. When assessing the volatility of mortality, various questions need to be answered:

- How many factors should be considered, i.e. what different types of possible shocks are there?
- What influence do these different types of shocks have for the close and for the far future, respectively?

[21] We do claim to have a better in-sample match than Ballotta and Haberman (2006), since our parameterization contains theirs as a special case. However, we did not examine out-of-sample effects.

- What form does the dependence of interest rates and mortality rates have under the pricing measure?
- Etc.

Some of these questions will require joint efforts of medical as well as of economical and mathematical experts. However, the problem of calibrating a given volatility structure clearly is a mathematical problem. As usual in financial modeling, this problem can be tackled in two ways: by estimating the volatility from historic time series or by using market prices to derive implied volatilities. Since a liquid market for mortality derivatives such as options on longevity bonds does not yet exist, the latter seems to be impossible. However, in Section 3.3 we showed how prices for longevity bonds can be derived from annuity data. Thus, the idea to calibrate the volatility structure to market prices of options offered within annuity contracts comes naturally.

In the following, we show how volatilities can be calibrated using products with guaranteed minimum income benefits (GMIBs) as offered within variable annuity contracts. These guarantees give the policyholder the option to annuitize a prespecified amount (often the single up-front premium compounded at a certain annual "roll-up" rate, e.g., 6%) at the future date T at a prespecified annuitization rate. Note that the guarantee only applies if the amount is annuitized at the specified conversion rate. Thus, at maturity T, the policyholder has the three following options: to take the current account value as a lump sum payment, to annuitize the current account value at prevailing annuitization rates at time T, or to annuitize the guaranteed amount at the conversion rates as offered within the GMIB resulting in a guaranteed annuity GA. Clearly, the first two options are of the same value, so the option to annuitize the guaranteed amount will be excercised if the value of the resulting annuity exceeds the account value at T. For a mored detailed introduction to options within variable annuity contracts, we refer to Bauer et al. (2006).

Proceeding similar to Ballotta and Haberman (2006), for the value of an annuity of \$1 starting to pay at the future date T given that the insured is still alive at time T we have:

$$\begin{aligned} a_{x_0+T} &= B_T\, E_Q\left[\sum_{k=0}^{\infty} B_k^{-1}\, {}_k p_{x_0+T} \mid \mathcal{F}_T\right] \\ &= \sum_{k=0}^{\infty} P(T,k)\, E_{Q_k}\left[{}_k p_{x_0+T} \mid \mathcal{F}_T\right] \\ &= \sum_{k=0}^{\infty} \exp\left\{-\int_T^k f(T,s) \;+\; \tilde{\mu}_T(s,x_0)\, ds\right\} \end{aligned}$$

Thus, we obtain for the time zero value of a Variable Annuity contract including a Guaranteed Minimum Income Benefit (GMIB) with a guaranteed annuity payment of GA and no Guaranteed Minimum Death Benefit (GMDB)[22]:

$$\begin{aligned} V_0^{GMIB} &= E_Q\left[\exp\left\{-\int_0^T r_s\, ds\right\}\, {}_T p_{x_0}\, \max\left\{\mathrm{GA}\, a_{x_0+T}, A_T\right\}\right] \\ &\qquad + E_Q\left[\exp\left\{\int_0^{\tau_{x_0}} r_s\, ds\right\}\, {}_T q_{x_0}\, A_{\tau_{x_0}}\right] \\ &= E_Q\left[\exp\left\{-\int_0^T r_s\, ds\right\}\, {}_T p_{x_0}\, \max\left\{\mathrm{GA}\, a_{x_0+T}, A_T\right\}\right] \end{aligned} \tag{13}$$

[22] That is, in case of death, only the current account value is paid out.

$$+E_Q\left[\int_0^T \exp\left\{-\int_0^t r_s \ + \ \tilde{\mu}_s\left(s,x_0\right)\,ds\right\}\tilde{\mu}_t\left(t,x_0\right)\,A_t\,dt\right]$$

where τ_{x_0} denotes the remaining lifetime of an x_0 year old at time 0 and A_t is the insured's account value at time t. Assuming an arbitrage free market and a fair valuation of the contract, we postulate that

$$V_0^{GMIB} \stackrel{!}{=} A_0, \tag{14}$$

where the GMIB-option is financed by a continuously deducted option charge proportional to the account value. This fee can be observed in the market; e.g., in January 2004, a 55-year old investing \$500,000 in the Equitable Accumulator product (see Equitable Accumulator Product Prospectus (2004)) with a GMIB for a fee of 0.65% was entitled to a guaranteed annuity payment of \$48,353 annually when annuitizing at the age of 65 years. Thus, equation (14) presents a equilibrium condition, which enables us to derive *implied volatilities* for the underlying processes. In particular, given specifications for the interest and asset processes, we are able to compute the implied volatility for the mortality process.

Unfortunately, aside from some special cases, there is no closed form representation for the value of a contract including a GMIB option (see equation (13)). However, numerical methods can be applied to derive the value and, hence, implied volatilities. Methods similar to those proposed in Bauer et al. (2006), where "deterministic mortalities" were used, can be employed.

Recapitulating, this section provided the methodology needed to completely determine and calibrate a stochastic mortality model, which offers a pricing mechanism in an arbitrage free market of mortality derivatives, which is consistent with our no arbitrage model from Section 3.2. However, there are still practical issues that need to be further explored. In the next section, we will show how these models can be applied to price more complex mortality derivatives and provide an outlook on possible extensions.

5 Application and possible extensions

When discussing the reasons why the BNP/EIB longevity bond could not be successfully offered, i.e. why the demand in the market was rather limited, one often presented argument is that the payoff structure was suboptimal: On the one hand the bond includes a fixed maturity date, i.e. it is not an adequate hedging instrument for a life-long annuity, and, on the other hand, it is paying the full ${}_tp_{x_0}$ rather than just a cap[23] and, thus, pays out too much. Therefore, the (T,x_0)-bonds as introduced here may not constitute adequate instruments either. Other instruments as for example survivor swaps (see Dowd et al. (2006)) or an option-type bond contract (see Lin and Cox (2005)) have been proposed, which are likely to present better, i.e. more marketable alternatives.

However, this does not mean that the framework introduced in the last section is inadequate. Conversely, when taking (T,x_0)-bonds as a starting point, various derivatives such as, e.g., survivor swaps can be priced. For example, when assuming deterministic volatilities, i.e. when assuming that σ^F and σ^L in (5) and (6) are deterministic, continuous functions in both parameters, we are able to obtain closed form solutions for multiple types of options. As an example, we consider payoffs of the type

$$P_T = \left({}_Tp_{x_0} - K\right)^+ = \max\left\{{}_Tp_{x_0} - K, 0\right\},$$

[23] Here, a cap is a contractual agreement where the granter has an obligation to pay cash to the holder if a specified quantity exceeds a mutually agreed level at some future date or dates (see Musiela and Rutkowski (1997), page 390 for interest rate caps).

where K is a fixed constant, the so-called strike. This call option type payoff is similar to the product introduced in Lin and Cox (2005), and K could for example be fixed at a certain expected level, e.g. $K = {}_T\tilde{p}_{x_0}^{be}$. Then, we obtain for the time zero price of this option (see Bauer (2006)):

$$P_0 = E_Q\left[B_T^{-1}\ P_T\right] = \Pi_0\left(T, x_0\right)\ \Phi\left(d_1\right)\ +\ K\, P\left(0,T\right)\ \Phi\left(d_2\right),$$

where

$$\begin{aligned} d_1 &= \frac{\log\left\{\frac{\Pi_0(T,x_0)}{K\,P(0,T)}\right\}\ +\ \frac{1}{2}\Sigma^2\left(T\right)}{\sqrt{\Sigma^2\left(T\right)}}, \\ d_2 &= d_1 - \sqrt{\Sigma^2\left(T\right)}, \end{aligned}$$

and

$$\Sigma^2\left(T\right) = \int_0^T \left\| \int_t^T \sigma^L\left(t,s\right)\ ds \right\|\ dt$$

where $\|\cdot\|$ denotes the Euclidean norm. Other options can be priced analogously, or, when no closed form solutions exist, by applying numerical methods.

So far, we considered only one single cohort. However, our approach can also be used to model all ages simultaneously. We can interpret $\tilde{\mu}_t\left(T, x\right)$ for a given t not only as a function in T, but rather as a two-dimensional function in T and x. Thus, after some extensions, the argumentation may still be valid. In particular, when including the dependence on x also in the volatility terms similar to Schrager (2006), we can model the interdependence between different bonds for different cohorts. However, we leave the further exploration of this issue for future work.

6 Conclusions

This paper presents a methodology for modeling and pricing mortality derivatives. Besides the valuation of so-called (T, x_0)-bonds as implied by no arbitrage arguments, our framework allows for a simultaneous modeling of the evolution of interest and mortality rates using the well-known HJM approach. Thus, arbitrary mortality derivatives such as options or survivor swaps can be priced. Furthermore, the framework allows for correlations between the dynamics of the term structure of interest rates and the dynamics of the term structure of mortality. For various reasons such correlations may be necessary when pricing mortality derivatives.

The resulting models produce prices which are arbitrage-free in a double sense: On the one hand, the initial prices of mortality derivatives are consistent with our no arbitrage model presented in Section 3, which describes the interrelations of the annuity market and the market of mortality derivatives via so-called *implied survival probabilities*; on the other hand, the prices are consistent within the hypothetical market of mortality derivatives. The latter is implied by imposing a HJM drift condition for mortality derivatives as presented in Section 4.

We introduce a method to derive the forward force of mortality curve analogously to methods from interest rate modeling via a parametrization of this curve, propose a specific parametrization, which can be used as a starting point of the dynamic model, and calibrate it to synthetic data. Furthermore, we show the volatilities of a specific model can be calibrated to existing products in the annuity market. In particular, we derive a valuation formula for a variable annuity contract including a guaranteed minimum benefit under stochastic mortalities, which can be evaluated numerically. After equating it with existing market prices, this formula presents a condition to which the volatility of mortality can be calibrated. We discuss how these models can be applied to value more complex mortality derivatives. As an example, we present the closed form solution for the value of an call option on a (T, x_0)-bond assuming a deterministic volatility structure. An instrument with this type of payoff may come closer to the demand of annuity providers, and thus may be more successful in the market than the EIB/BNP longevity bond as proposed in 2004.

So far, the framework only allows for modeling and pricing derivatives on the number of survivors of one single cohort. However, we explain how the framework could be extended to allow for simultaneously modeling mortality derivatives for multiple cohorts, or even the whole population. This may be achieved by making volatilities not only dependent on the time and the maturity, but also of age which is added as another variable. However, this issue needs to be further explored.

Besides the use as a pricing tool, the framework also offers the possibility to derive hedging strategies for particular models. Even though this may be of less importance at this time since the underlying instruments are not yet traded, this may be of interest in the future: If there is a liquid market for longevity bonds, similar to the fixed income market, hedging will become an important risk management tool.

The next issues that need to be addressed are clear straight away: In order to obtain a workable model, we need to adequately specify the volatility structure. As already mentioned earlier in the text, this is a rather difficult problem which clearly exceeds the limits of financial engineering and statistics. Volatilities need to reflect possible movements of the future term structure of mortalities. Therefore, this task has to be tackled by combining statistical aspects with insights from epidemiologists and health professionals. Also, instead of using synthetic data, market data should be used to derive implied survival probabilities. This is a considerable task and requires both, a thorough knowledge of the annuity market and sophisticated methods how to deal with missing or ambiguous data. Furthermore, numerous practical problems have to be solved, for example how payoffs of (T, x_0)-bonds can be determined at or close to maturity T.

In spite of these open issues, we believe that our approach offers important insights and that our framework presents an adequate method for pricing and modeling mortality derivatives.

References

Ballotta, L., Haberman, S., 2006. The fair valuation problem of guaranteed annuity options: The stochastic mortality environment case. Insurance: Mathematics and Economics, 38: 195–214.

Bauer, D., Kling, A., Russ, J., 2006. A Universal Pricing Framework for Guaranteed Minimum Benefits in Variable Annuities. Working Paper, Ulm University.

Bauer, D., 2006. An Arbitrage-Free Family of Longevity Bonds. Working Paper, Ulm University. Available at: www.mortalityrisk.org.

Biffis, E., 2005. Affine processes for dynamic mortality and actuarial valuation. Insurance: Mathematics and Economics, 37:443–468.

Biffis, E., Denuit, M., Devolder, P., 2005. Stochastic Mortality Under Measure Change. Working Paper, CASS Business School and Université Catholique de Louvain.

Bingham, N., Kiesel, R., 2003. Risk-Neutral Valuation. Springer Finance. Springer, London.

Björk, T., 1999. Arbitrage Theory in Continuous Time. Oxford University Press, Oxford, UK.

Blake, D., Burrows, W., 2001. Survivor Bonds: Helping to Hedge Mortality Risk. The Journal of Risk and Insurance, 68: 339–348.

Blake, D., Cairns, A.J., Dowd, K., 2006. Living with Mortality: Longevity Bonds and other Mortality-Linked Securities. Presented to the Faculty of Actuaries, 16 January 2006.

Cairns, A.J., Blake, D., Dowd, K., 2005a. Pricing the Risk on Longevity Bonds. Life and Pensions, October 2005: 41–44.

Cairns, A.J., Blake, D., Dowd, K., 2005b. Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk. ASTIN Bulletin, to appear.

Cairns, A.J., Blake, D., Dowd, K., 2005c. A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty. To appear in The Journal of Risk and Insurance.

Cochrane, J.H., 2005. Asset Pricing. Princeton University Press, Princeton.

Cowley, A., Cummins, J.D., 2005. Securitization of Life Insurance Assets and Liabilities. The Journal of Risk and Insurance, 72: 193–226.

Currie, I.D., Durban, M., Eilers, P.H.C, 2004. Smoothing and forecasting mortality rates. Statistical Modelling, 4: 279–298.

Dahl, M., 2004. Stochastic mortality in life insurance: market reserves and mortality-linked insurance contracts. Insurance: Mathematics and Economics, 35: 113–136.

Dahl, M., Møller, T., 2006. Valuation and Hedging of Life Insurance Liabilities with Systematic Mortality Risk. Insurance: Mathematics and Economics, 39: 193–217.

Denuit, M., 2005. Life Insurance Pricing with Dynamic Mortality. Presentation at the Scientific Conference on Insurance and Finance of the DGVFM, available at *www-m4.ma.tum.de/dgvm05/denuittalk.pdf*.

Dowd, K., Blake, D., Cairns, A.J.G., Dawson, P., 2006. Survivor Swaps. The Journal of Risk and Insurance, 73:1–17.

Equitable, 2004. Equitable Accumulator Product Prospectus. The Equitable Life Assurance Society of the United States, New York.

Filipović, D., 2001. Consistency Problems for Heath Jarrow Morton Interest Rate Models. Lecture Notes in Mathematics, 1760. Springer, Berlin.

Gerber, H.U., Pafumi, G., 1998. Utility Functions: From Risk Theory to Finance. North American Actuarial Journal, 2:74–100.

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

Lee, R.D, Carter, L.R., 1992. Modeling and forecasting US mortality. Journal of the American Statistical Association, 87: 659–675.

Lin, Y., Cox, S., 2005. Securitization of Mortality Risks in Life Annuities. The Journal of Risk and Insurance, 72: 227–252.

Luciano, E., Vigna, E., 2005. Non mean reverting affine processes for stochastic mortality. ICER Applied Mathematics Working Paper Series, Working Paper no. 4/2005.

Milevsky, M.A., Promislow, S.D. , 2001. Mortality Derivatives and the Option to Annuitize. Insurance: Mathematics and Economics, 29:299 – 318.

Milevsky, M.A., Promislow, S.D., Young, V.R., 2005. Financial Valuation of Mortality Risk via the Instantaneous Sharpe Ratio. Working Paper, York University and University of Michigan.

Miltersen, K.R., Persson, S.A. , 2005. Is mortality dead? Stochastic force of mortality determined by no arbitrage. Working Paper, University of Bergen.

Musiela, M., Rutkowski, M., 1997. Martingale Methods in Financial Modelling. Application of Mathematics, 36. Springer, Berlin.

Nelson, C.R., Siegel, A.F. , 1987. Parsimonious modelling of yield curves. Journal of Business, 60: 473–489.

Pelsser, A. , 2003. Pricing and Hedging Guaranteed Annuity Options via Static Option Replication. Insurance: Mathematics and Economics, 33: 283–296.

Pelsser, A. , 2004. On the Applicability of the Wang Transform for Pricing Financial Risks. Working Paper, Erasmus University Rotterdam.

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

Schoutens, W., 2003. Lévy Processes in Finance: Pricing Financial Derivatives. Wiley Series in Probability and Statistics. Wiley, Chichester.

Schrager, D., 2006. Affine stochastic mortality. Insurance: Mathematics and Economics, 38:81–97.

Wang, S. , 2002. A Universal Framework for Pricing Financial and Insurance Risks. ASTIN Bulletin, 32: 213–234.

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