



City Research Online

City, University of London Institutional Repository

Citation: Hunt, A. and Blake, D. ORCID: 0000-0002-2453-2090 (2015). Forward Mortality Rates in Discrete Time I: Calibration and Securities Pricing (PI-1511). Pensions Institute.

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/26072/>

Link to published version: PI-1511

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.



DISCUSSION PAPER PI-1511

Forward Mortality Rates in Discrete Time I:
Calibration and Securities Pricing.

Andrew Hunt and David Blake

September 2015

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University London
106 Bunhill Row
London EC1Y 8TZ
UNITED KINGDOM

<http://www.pensions-institute.org/>

Forward Mortality Rates in Discrete Time I: Calibration and Securities Pricing

Andrew Hunt*

Cass Business School, City University London

Corresponding author: `andrew.hunt.1@cass.city.ac.uk`

David Blake

Pensions Institute, Cass Business School, City University London

September 2015

Abstract

Many users of mortality models are interested in using them to place values on longevity-linked liabilities and securities. Modern regulatory regimes require that the values of liabilities and reserves are consistent with market prices (if available), whilst the gradual emergence of a traded market in longevity risk needs methods for pricing new types of longevity-linked securities quickly and efficiently. In this study, we develop a new forward mortality framework to enable the efficient pricing of longevity-linked liabilities and securities in a market-consistent fashion. This approach starts from the historical data of the observed mortality rates, i.e., the force of mortality. Building on the dynamics of age/period/cohort models of the observed force of mortality, we develop models of forward mortality rates and then use

*Work in this paper was presented at the 49th Actuarial Research Conference in Santa Barbara, USA, in July 2014, the Tenth International Longevity Conference in Santiago, Chile, in September 2014, and the Society of Actuaries Longevity Seminar in Chicago, USA, in February 2015. We are grateful to participants at these conferences for their comments and suggestions, to Andrew Cairns and Pietro Millossovich for their helpful review on an earlier draft of this paper, and to Andrés Villegas for many useful discussions on this and related topics.

a change of measure to incorporate whatever market information is available. The resulting forward mortality rates are then used to value a number of different longevity-linked securities, such as q-forwards, s-forwards and longevity swaps.

JEL Classification: G12

Keywords: Mortality modelling, age/period/cohort models, forward mortality rates, Esscher transform, longevity-linked securities

1 Introduction

Many users of mortality models are interested in using them to place values on longevity-linked liabilities and securities. Modern regulatory regimes require that the values of liabilities and reserves are consistent with market prices (if available), whilst the gradual emergence of a traded market in longevity risk needs methods for pricing new types of longevity-linked securities quickly and efficiently. These needs have spurred the development of increasingly sophisticated models of mortality rates.

Cairns et al. (2006b) pointed out that the majority of mortality models that have been proposed are models of the mortality hazard rate, which is analogous to the short rate of interest. By analogy with interest rate models, Cairns et al. (2006b) developed formally the concept of “mortality forward rates”, which was extended in Miltersen and Persson (2005). However, the idea of forward mortality rates has a long history, indeed Milevsky and Promislow (2001) pointed out that *“the traditional rates used by actuaries are really ‘forward rates’ exactly analogous to a forward interest rate implied by existing bond prices”*.

Such forward mortality rates could be used to price longevity-linked securities, in the same fashion as forward interest rates are used to value cash-flows dependent on future interest rates. Therefore, a number of models for forward mortality rates have been proposed to date which build upon the theory of forward interest rates. These have included the models of Barbarin (2008), Bauer et al. (2008) and Tappe and Weber (2013), which adopted the Heath-Jarrow-Morton framework used for interest rates in continuous time, and the model of Zhu and Bauer (2011a,b, 2014) which adopted a

semi-parametric factor approach in discrete time. An alternative approach, developed in Olivier and Jeffrey (2004), Smith (2005) and Cairns (2007), also works in discrete time but uses gamma-distributed random variables to update a forward mortality surface that is initially assumed.

However, it is important not to over-extend the analogy between interest rates and mortality rates, as the two are fundamentally different processes. Most obviously, the forward interest rate curve at any instant depends only upon term, whilst forward mortality rates will exist across a surface of ages and years. Mortality rates typically also increase exponentially with age, unlike interest rates which are typically bounded as term increases. More fundamentally, the analogy between survivorship under a force of mortality and discounting under a force of interest, whilst mathematically appealing, is not exact, since mortality will affect the actual amount of any cashflow payable (say, in an annuity or life assurance contract) in a way that discounting does not. We therefore do not believe that simply taking existing models which work well for forward interest rates and applying them directly to mortality rates is appropriate.

In addition, we must be able to calibrate a model of forward mortality rates to the small number of longevity-linked securities in existence. This means that models which start by assuming the existence of sufficient market prices to define a forward mortality surface (such as those based on the Heath-Jarrow-Morton framework) and then define the dynamics of this surface are not practical. This approach is inherited from the interest rate markets, where liquid markets in bonds across the whole of the relevant term structure can provide such information. Unfortunately, this simply does not hold for the market in longevity-linked securities, and will not hold for the foreseeable future.

Instead, we propose a new approach, which is described in two studies, of which this is the first. Our approach starts from the historical data on the observed mortality rates, i.e., the observed force of mortality which is analogous to the short rate of interest. Building on the dynamics of models of the observed force of mortality, we can recast them in the form of models of forward mortality rates and then use a change of measure to incorporate whatever market information is available. This approach ensures that the dynamics of the forward mortality surface are consistent with those observed

for the force of mortality, including features such as “cohort effects” which are unique to mortality rate models, and which helps to ensure demographic significance.¹

We begin our analysis in this paper in Section 2.1 with models of the force of mortality from the age/period/cohort (APC) family, which have been specifically constructed in order to capture the dynamics of mortality parsimoniously and with demographic significance. APC mortality models are considered in detail in Hunt and Blake (2015i) and encompass a broad class of existing and popular models of the force of mortality, such as the Lee-Carter (Lee and Carter (1992)), Cairns-Blake-Dowd (Cairns et al. (2006a)) and classic APC (Hobcraft et al. (1982)) models, as well as many of the extensions of these models (see Hunt and Blake (2014) for examples). We then develop the mathematical framework required to convert any APC model of the force of mortality into a model of the forward mortality surface in Section 2.2 and Section 2.3. In Section 2.4, we use the dynamics of the period and cohort parameters observed in the historical data to define a forward surface of mortality rate. This enables consistent modelling of both the short and forward mortality rates, and so avoids any inconsistencies between the two.

Section 3 then builds on this by transforming the forward mortality rate surface, using the Esscher transform, from a measure consistent with the “real-world” process observed in the historical data to one consistent with market prices. These “market-consistent” forward mortality rates are then used to price various longevity-linked securities. Finally, Section 4 concludes.

The approach established in this paper is extended in our second paper, Hunt and Blake (2015d), which analyses how the forward surface of mortality can be updated dynamically. This enables the forward mortality rate framework developed in this paper to be used for managing longevity risk in a life assurance book or in a portfolio of longevity-linked securities.

¹Demographic significance is defined in Hunt and Blake (2015i) as the interpretation of the components of a model in terms of the underlying biological, medical or socio-economic causes of changes in mortality rates which generate them.

2 Forward mortality rates in discrete time

2.1 Age/period/cohort models of the force of mortality

In Hunt and Blake (2015i), we discussed discrete-time mortality models of the form

$$\eta_{x,t} = \alpha_x + \sum_{i=1}^N \beta_x^{(i)} \kappa_t^{(i)} + \gamma_{t-x} \quad (1)$$

where

- we have historical data for ages, x , in the range $[1, X]$ and periods, t , in the range $[1, \tau]$ and therefore observations of cohorts born in years, y , in the range $[1 - X, \tau - 1]$;
- $\eta_{x,t} = \ln(\mu_{x,t})$ is the log-link function which connects the Poisson distributed death counts, $D_{x,t}$, to the proposed predictor structure;
- α_x is a static function of age;
- $\kappa_t^{(i)}$ are period functions governing the evolution of mortality with time;
- $\beta_x^{(i)}$ are age functions modulating the impact of the period function dynamics over the age range;² and
- γ_y is a cohort function describing mortality effects which depend upon a cohort's year of birth and follow that cohort through life as it ages.

Defining $\boldsymbol{\beta}_x = \left(\beta_x^{(1)}, \dots, \beta_x^{(N)} \right)^\top$ and $\boldsymbol{\kappa}_t = \left(\kappa_t^{(1)}, \dots, \kappa_t^{(N)} \right)^\top$, we can re-write Equation 1 as

$$\eta_{x,t} = \alpha_x + \boldsymbol{\beta}_x^\top \boldsymbol{\kappa}_t + \gamma_{t-x} \quad (2)$$

In this paper, we will use the log-link function $\eta_{x,t} = \ln(\mu_{x,t})$. In Hunt and Blake (2015i), we discussed how this is appropriate if the death count

²These can be non-parametric in the sense of being one fitted without imposing any a priori shape for the function across ages, or be parametric in the sense of having a specific functional form, $\beta_x^{(i)} = f^{(i)}(x; \theta^{(i)})$ selected a priori. Potentially, parametric age functions can have free parameters $\theta^{(i)}$ which are set with reference to the data.

at age x and time t is a (conditionally independent) Poisson random variable, $D_{x,t} \sim Po(\mu_{x,t} E_{x,t}^c)$, where $E_{x,t}^c$ are central exposures to risk. This is preferred over the alternative choice of the logit-link function and binomially distributed death counts due to the distributional properties of the forward mortality rates, as discussed in Section 2.3.

This structure defines the class of age/period/cohort (APC) mortality models and is very flexible. Many of the most common mortality models fit into this structure, for instance, the benchmark Lee-Carter (LC) model of Lee and Carter (1992), the cohort extension to this denoted H1 in Haberman and Renshaw (2009), the Cairns-Blake-Dowd (CBD) model of Cairns et al. (2006a) and many of its extensions in Cairns et al. (2009), the Plat model of Plat (2009) and the model of Börger et al. (2013). In Hunt and Blake (2014), we describe a “general procedure” for constructing bespoke models within this class which are tailored to the structure within a given dataset.³ It is, therefore, appropriate to use this class of models of the force of mortality as the starting point for defining the forward mortality surface, as discussed below.

2.2 Defining forward mortality rates

In a discrete-time framework, the force of mortality, $\mu_{x,t}$, at age x and time t is assumed to be constant over each age and year, i.e.,

$$\begin{aligned} \mu_{x+\xi, t+\tau} &= \mu_{x,t} \\ x, t &\in \mathbb{N} \\ \xi, \tau &\in [0, 1) \end{aligned} \tag{3}$$

Therefore, the one-year survival probability from age x at time t to age $x+1$ at time $t+1$, $p_{x,t}$,⁴ is equal to $p_{x,t} = \exp(-\mu_{x,t})$. If we further assume that

³The forward mortality framework described in this study is not significantly affected if the cohort parameters are modulated by an age function, $\beta_x^{(0)}$, as in the model of Renshaw and Haberman (2006). However, for simplicity and the reasons discussed in Hunt and Blake (2015i), we do not consider such models in this study.

⁴ $p_{x,t} = 1 - q_{x,t}$, the one-year probability of death.

survival in each year is conditionally independent, this implies

$${}_t p_{x,\tau} = \prod_{u=1}^t p_{x+u,\tau+u} = \exp \left(- \sum_{u=1}^t \mu_{x+u,\tau+u} \right) \quad (4)$$

where ${}_t p_{x,\tau}$ is the survival probability of an individual from age x at time τ to age $x+t$ at time $\tau+t$.⁵ If $\tau+t$ lies in the future, ${}_t p_{x,\tau}$ will be a random variable, as future values of the force of mortality will be subject to systematic mortality risk.

To define the structure of forward mortality rates, we assume that the fundamental longevity-linked security⁶ of interest, from which all other longevity-linked securities can be constructed, is the “longevity zero”.⁷ A longevity zero is defined in Blake et al. (2006) as a zero-coupon bond which pays out a principal at a future time, dependent on the survivorship of a suitably large cohort (to reduce the idiosyncratic risk in the estimation of survival rates) over the term of the bond.⁸ Therefore, a t -year longevity zero at time τ would have price

$$\text{Price}(t, \tau) = B(\tau, \tau+t) \mathbb{E}_{\tau}^{\mathbb{Q}} {}_t p_{x,\tau}$$

where $B(\tau, \tau+t)$ is the time τ price of a t -year zero coupon bond paying one unit at maturity, and where the expectation is defined under some “market-consistent” measure, \mathbb{Q} (to be discussed in Section 3).⁹

In doing so, we have implicitly assumed that the longevity risk is independent of the other financial risks in the market, such as interest rates and

⁵ ${}_0 p_{x,\tau} = 1$ trivially.

⁶In this paper, we use the term “security” to refer to any tradable financial contract, and so also include derivative securities such as forwards and options in this definition.

⁷Longevity zeros were also used to define forward mortality rates in Barbarin (2008) for use in a Heath-Jarrow-Morton framework and in Cairns (2007) and Alai et al. (2013) to develop extensions of the Olivier-Smith model.

⁸It is important that the security used to define the forward mortality rates depends purely on the systematic component of longevity risk, rather than on the idiosyncratic time of death of any individual lives, in order to avoid the potential for conflicting definitions of the forward rates described in Norberg (2010).

⁹We adopt the convention that the subscript on operators $\mathbb{E}_{\tau}(\cdot)$, $\text{Var}_{\tau}(\cdot)$ or $\text{Cov}_{\tau}(\cdot)$ denotes conditioning on the information available at time τ , i.e., \mathcal{F}_{τ} .

inflation, in both the real-world measure, \mathbb{P} , and the market-consistent measure, \mathbb{Q} . This is in common with the majority of studies, such as Cairns et al. (2006b) and Bauer et al. (2008) and with the available evidence to date, as discussed in Loeys et al. (2007). Although there may be some situations where longevity risk is not independent of other financial risks in the real-world measure, as in the examples of Miltersen and Persson (2005), we believe that these situations are relatively extreme and are better considered by scenario analysis rather than through a stochastic model. Furthermore, Dhaene et al. (2013) show that independence between longevity risk and financial risks in the real-world measure does not automatically ensure independence in the market-consistent measure. However, more complicated models are required in order to allow for any dependence between longevity and investment risks, which require more market information for calibration. Therefore, we believe that the assumption of independence between longevity risk and other financial risks is necessary and justifiable at this early stage of development of the longevity risk market.

We define

$$\begin{aligned} {}_tP_{x,\tau}^{\mathbb{Q}}(\tau) &= \mathbb{E}_{\tau}^{\mathbb{Q}} {}_tp_{x,\tau} \\ &= \mathbb{E}_{\tau}^{\mathbb{Q}} \exp \left(- \sum_{u=1}^t \mu_{x+u,\tau+u} \right) \end{aligned} \tag{5}$$

In this, ${}_tP_{x,\tau}^{\mathbb{Q}}(\tau)$ are the market-consistent forward survival probabilities, i.e., the “*market’s best view*” (in the words of Miltersen and Persson (2005)) at τ of the probability of an individual aged x at τ surviving a further t years. Mathematically, we can see that these factors are analogous to discount factors based on the prices of zero-coupon bonds. It is this analogy which has motivated much of the development of forward mortality rate models to date, which have been mainly adapted from widely used interest rate models. In continuous-time forward rate models, such as in Bauer et al. (2008), forward mortality rates are defined from Equation 5 as

$$\nu_{x,t}^{\mathbb{Q}}(\tau) \equiv - \frac{\partial}{\partial t} \ln ({}_tP_{x-t,\tau}^{\mathbb{Q}}(\tau))$$

via the analogy with forward interest rates. In a discrete time model, we

modify this to define forward mortality rates as

$$\nu_{x,t}^{\mathbb{Q}}(\tau) \equiv -\ln \left(\frac{{}_{t-\tau+1}P_{x-t+\tau,\tau}^{\mathbb{Q}}(\tau)}{{}_{t-\tau}P_{x-t+\tau,\tau}^{\mathbb{Q}}(\tau)} \right) \quad (6)$$

Existing forward mortality models, such as those in Cairns (2007) and Zhu and Bauer (2011b, 2014) use similar definitions, but these studies are interested in the dynamics of the forward surface of mortality and so are interested in the behaviour of $\nu_{x,t}(\tau+1)/\nu_{x,t}(\tau)$, rather than the forward mortality rates at τ themselves (which are assumed a priori in these studies). We discuss these dynamics in Hunt and Blake (2015d). In contrast, this paper is interested in the connection between the force of mortality and forward mortality rates, and so we use the definition above to give

$${}_tP_{x,\tau}^{\mathbb{Q}}(\tau) = \exp \left(- \sum_{u=1}^t \nu_{x+u,\tau+u}^{\mathbb{Q}}(\tau) \right) \quad (7)$$

Comparing Equations 4 and 7, we see

$$\exp \left(- \sum_{u=1}^t \nu_{x+u,\tau+u}^{\mathbb{Q}} \right) = \mathbb{E}_{\tau}^{\mathbb{Q}} \exp \left(- \sum_{u=1}^t \mu_{x+u,\tau+u} \right) \quad (8)$$

which shows the connection between the market-consistent forward rates and the expectations of the force of mortality in the market-consistent measure.

By Jensen's inequality

$$\mathbb{E}_{\tau}^{\mathbb{Q}} \exp \left(- \sum_{u=1}^t \mu_{x+u,\tau+u} \right) \geq \exp \left(- \sum_{u=1}^t \mathbb{E}_{\tau}^{\mathbb{Q}} \mu_{x+u,\tau+u} \right) \quad (9)$$

In practice, the variation in $\mu_{x,t}$ is sufficiently small that Equation 9 holds approximately as an equality over almost all ages and years.¹⁰ We therefore make the assumption that

$$\exp \left(- \sum_{u=1}^t \nu_{x+u,\tau+u}^{\mathbb{Q}}(\tau) \right) = \exp \left(- \sum_{u=1}^t \mathbb{E}_{\tau}^{\mathbb{Q}} \mu_{x+u,t+u} \right) \quad (10)$$

¹⁰This approximation is tested numerically in Appendix B.

and define the forward mortality rates as

$$\nu_{x,t}^{\mathbb{Q}}(\tau) = \mathbb{E}_{\tau}^{\mathbb{Q}} \mu_{x,t} \quad (11)$$

Thus, the forward mortality rate at age x and year t is assumed to be equal to the expectation under the market-consistent measure of the force of mortality at the same age and year, conditional on information observed at time τ . Thus, if we can specify the dynamics of the force of mortality (in the market-consistent measure), we are able to find the forward mortality rates directly.

We define the “forward mortality surface” as the collection of forward mortality rates, $\nu_{x,t}^{\mathbb{Q}}(\tau)$ over all ages, x , and future years, t , at a given point in time, τ . In most cases, it is more natural to consider the forward mortality surface as a single object, since the individual forward mortality rates are expected to vary smoothly across ages and across future years. However, it is important to realise that the forward mortality surface is three-dimensional, defined by x , t and τ . In this paper we shall consider its structure across the dimensions of x and t and how this can be determined at the observation time, τ , which is assumed to be fixed. This contrasts with Hunt and Blake (2015d), where we discuss how the surface varies dynamically with τ .

In defining the forward mortality surface, we assume that all longevity-linked securities can be constructed from a portfolio of longevity zeros. We shall see in Section 3.3 that this is trivially true in the case of longevity swaps.¹¹ We extend this by assuming that the value of any other longevity-linked security at time τ can be replicated as a portfolio of longevity zeros and, therefore, written as a function of the $\nu_{x,t}^{\mathbb{Q}}(\tau)$. Hence, the forward surface of mortality can be used to give consistent prices for all longevity-linked liabilities and securities.

Unfortunately, however, it is currently impossible to reliably specify the dynamics of short or forward mortality rates in the market-consistent measure, since an actively-traded market in longevity-linked securities does not currently exist. Indeed, the absence of genuine market information on the prices for any longevity-linked securities is a critical problem for all studies

¹¹It is also true for the valuation of annuities for reserving purposes, since idiosyncratic risk is not allowed for in this context.

that seek to value the few longevity-linked securities which do exist. There have been a number of different methods proposed to overcome this and calibrate the market-consistent measure. For instance, Bauer et al. (2008) proposed using generational life tables (i.e., those which allow mortality rates to depend upon an individual’s year of birth) in order to provide a forward mortality surface. However, these are updated infrequently and are not based on market information (and when used to price financial contracts, typically have margins for risk aversion added to them). Alternatively, Miltersen and Persson (2005) and Bayraktar and Young (2007) have suggested using the market for endowment assurances for calibration purposes, since these have a similar price structure to longevity zeros. Unfortunately, Norberg (2010) showed how using securities dependent on the idiosyncratic risk of individual lives, such as endowment assurances, can lead to inconsistent definitions of the forward mortality rates and so this approach is not feasible.

Instead, we propose to use the historical data to model the dynamics of the force of mortality in the “historical” or “real-world” measure, \mathbb{P} , using relatively simple APC mortality models, as described in Section 2.1. These real-world dynamics of the force of mortality can then be used to generate the forward surface of mortality in the real-world measure by using Equation 11. Then, in Section 3.1, we show how to change from the real-world to a market-consistent measure, \mathbb{Q} , using the Esscher transform which is calibrated using whatever (limited) market information for longevity risk is available. Thus, real-world data on historical mortality rates is used to supplement the limited market data we have, and increasing volumes of market information can be incorporated into the forward mortality surface as the market for longevity-linked securities develops.

2.3 Forward APC mortality models

Combining Equations 2 and 11, we define forward mortality rates in the real-world measure, \mathbb{P} , as

$$\nu_{x,t}^{\mathbb{P}}(\tau) = \mathbb{E}_{\tau}^{\mathbb{P}} \exp(\alpha_x + \beta_x^{\top} \boldsymbol{\kappa}_t + \gamma_{t-x}) \quad (12)$$

We assume that the age functions are known with certainty at time τ and therefore the uncertainty in future mortality rates comes from the projection of $\boldsymbol{\kappa}_t$ and γ_{t-x} , i.e., the forward mortality surface only allows for process risk

from the projection of the period and cohort functions, in the terminology of Cairns (2000), but not parameter uncertainty or model risk. In the real-world measure, we first obtain fitted values of κ_t and γ_y by fitting the APC model to the historical data. We then estimate the dynamics of the time series processes for κ_t and γ_y from these fitted values.

If we further assume that our projected κ_t and γ_y are normally distributed, then $\eta_{x,t}$ is also normally distributed and consequently $\mu_{x,t}$ follows a log-normal distribution.¹² Therefore

$$\nu_{x,t}^{\mathbb{P}}(\tau) = \exp \left(\alpha_x + \beta_x^{\top} \mathbb{E}_{\tau}^{\mathbb{P}} \kappa_t + \frac{1}{2} \beta_x^{\top} \text{Var}_{\tau}^{\mathbb{P}}(\kappa_t) \beta_x + \mathbb{E}_{\tau}^{\mathbb{P}} \gamma_{t-x} + \frac{1}{2} \text{Var}_{\tau}^{\mathbb{P}}(\gamma_{t-x}) \right) \quad (13)$$

The assumption that projected period and cohort parameters are normally distributed is in line with the majority of studies, which use standard ARIMA methods to project these parameters. If the projected period and cohort parameters are not normally distributed, however, it is unlikely that the resulting forward mortality framework would be analytically tractable. This is because the distribution of $\mu_{x,t}$ would not have the finite moments required. A number of studies have used alternative methods and distributions to make projections. These include models which allow for regime changes (Milidonis et al. (2011) and Lemoine (2014)) or trend changes (Sweeting (2011) and Hunt and Blake (2015c)) in the processes used to project the parameters. Another approach has been to use other distributions for the innovations in the time series processes for the period or cohort functions (such as the t-distribution, the variance-gamma and the normal-inverse-gamma, which were used to model the innovations for κ_t in the Lee-Carter model in Wang et al. (2011)). In some of these cases, it may be possible to extend the forward mortality rate framework to allow for the non-Gaussian distributions. However, we do not consider alternative distributions for the projected period or cohort functions further within this study.

¹²Note that, if we were using $\eta_{x,t} = \text{logit}(q_{x,t})$ in conjunction with a binomial model for the death count, then $q_{x,t}$ would follow a “logit-normal” distribution (see Frederic and Lad (2008)). Unfortunately, this is not analytically tractable and does not possess closed form expressions for the expectation. Therefore, we are unable to define a forward mortality framework in the logit-link function / binomial death count model as we can in the log-link function / Poisson death count model.

2.4 Projecting the APC model

2.4.1 Period functions

Since Lee and Carter (1992), the most common method used to project the period functions in an APC mortality model has been the random walk with drift. This was also used for the CBD model in Cairns et al. (2006a), the period functions in various mortality models in Cairns et al. (2011) and Haberman and Renshaw (2011), and the first (dominant) period function in Plat (2009).

The random walk model is attractive as it allows the period functions to be non-stationary with a variability that increases with time, giving biologically reasonable¹³ projections of the force of mortality.

In Hunt and Blake (2015f,g), we discuss how projected mortality rates should not depend upon the identifiability constraints used when fitting the model to data, and therefore that we should use “well-identified” projection methods which achieve this. In the context of the random walk with drift model, this means we should project the period functions using

$$\kappa_t = \mu X_t + \kappa_{t-1} + \epsilon_t \quad (14)$$

where X_t is a set of deterministic functions (“trends”) chosen to ensure identifiability and μ are the corresponding “drifts”.¹⁴ For example, the classic random walk with drift process has a constant trend, $X_t = 1$, with the “drift”, μ , found by regressing $\Delta\kappa_t$ on this trend. Similarly, the random walk with linear drift introduced in Hunt and Blake (2015g) and Hunt and Blake (2015c) has constant and linear trends, $X_t = (1, t)^\top$, with the drifts found by regressing $\Delta\kappa_t$ against X_t in a similar fashion.

¹³Introduced in Cairns et al. (2006b) and defined as “*a method of reasoning used to establish a causal association (or relationship) between two factors that is consistent with existing medical knowledge*”.

¹⁴Note, we assume that the drifts μ are known at time τ and will not be re-estimated on the basis of new information arising in the future. Therefore, the forward mortality framework described in this paper and in Hunt and Blake (2015d) does not allow for “recalibration” risk as defined in Cairns (2013), i.e., the risk caused by the uncertainty in the drift. This risk is potentially substantial, as discussed in Li et al. (2004) and Li (2014). However, we leave the inclusion of recalibration risk to future work.

The random drift model in Equation 14 is solved to give

$$\boldsymbol{\kappa}_t = \boldsymbol{\kappa}_\tau + \mu\chi_{\tau,t} + \sum_{s=\tau+1}^t \boldsymbol{\epsilon}_s \quad (15)$$

where $\chi_{\tau,t} = \sum_{s=\tau+1}^t X_s$. Note that, in the simplest case where we use a classic random walk with drift to project the period functions, $X_t = 1$ and hence $\chi_{\tau,t} = t - \tau$. We assume

$$\begin{aligned} \mathbb{E}_\tau \boldsymbol{\epsilon}_t &= \mathbf{0} \\ \text{Cov}_\tau(\boldsymbol{\epsilon}_t, \boldsymbol{\epsilon}_s) &= \Sigma \mathbf{I}_{t-s} \end{aligned}$$

where \mathbf{I}_{t-s} is an indicator variable taking a value of unity if $t = s$ and zero otherwise. This means that the innovations have zero mean and are independent across different periods, i.e., they are white noise. In addition, we assume that the innovations are normally distributed for the reasons discussed above. From Equation 15, we find

$$\mathbb{E}_\tau \boldsymbol{\kappa}_t = \boldsymbol{\kappa}_\tau + \mu\chi_{\tau,t} \quad (16)$$

$$\text{Var}_\tau(\boldsymbol{\kappa}_t) = (t - \tau)\Sigma \quad (17)$$

In an age/period mortality model without a cohort term, such as the Lee-Carter or CBD model, allowing for the uncertainty in the period functions is sufficient in conjunction with Equation 13, to define forward mortality rates in the real-world measure. However, more sophisticated mortality models often include cohort terms, whose analysis is considerably more complicated, as we now see.

2.4.2 Cohort function

Most common techniques for projecting the cohort function use standard ARIMA processes, which assume that there is a clear distinction between those cohort parameters which are estimated from historical data, which are assumed to be known, and those cohort parameters which are projected using some time series process. In the forward mortality rate framework, we can see that this would lead to a sharp discontinuity in the forward mortality surface. For many purposes, such as the valuation of longevity-linked securities and liabilities, such a discontinuity is clearly undesirable.

To illustrate this problem, consider the case where a (well-identified) AR(1) process is used to project the cohort parameters

$$\gamma_y - \beta \tilde{X}_y = \rho(\gamma_{y-1} - \beta \tilde{X}_{y-1}) + \varepsilon_y$$

where \tilde{X}_y are deterministic functions corresponding to the unidentifiable trends in the cohort parameters,¹⁵ and β are the corresponding regression coefficients (see Hunt and Blake (2015g)). Such a process would be solved to give

$$\gamma_y = \rho^{y-Y}(\gamma_Y - \beta \tilde{X}_Y) + \beta \tilde{X}_y + \sum_{s=Y+1}^y \rho^{y-s} \varepsilon_s$$

for $y \geq Y$, the year of birth of the last fitted cohort parameter.¹⁶ The variance of this process is

$$\mathbb{V}ar_{\tau}(\gamma_y) = \begin{cases} 0 & \text{if } y \leq Y \\ \frac{1-\rho^{2(y-Y)}}{1-\rho^2} \sigma^2 & \text{if } y > Y \end{cases}$$

From Equation 13, we see that this would give a discontinuity in the forward mortality surface at the interface between the fitted and projected cohort parameters. Such a discontinuity would give rise to pricing anomalies and therefore cannot be permitted in a well-designed forward mortality framework. Consequently, we must use alternative processes to project the cohort parameters for use with forward mortality models.

In Hunt and Blake (2015a), we developed a Bayesian approach to overcome this issue. This assumes that all cohort parameters, γ_y , are random variables that are not fully observed until cohort y is fully extinct at time $y + X$. For observation times $\tau < y + X$, we have partial information based on observations of the cohort to date. This information is summarised in the estimated cohort parameters, $\bar{\gamma}_y(\tau)$, found by fitting the APC mortality model to data to time τ . From the analysis in Hunt and Blake (2015a), we have

$$\gamma_y | \mathcal{F}_{\tau} \sim N(M(y, \tau), V(y, \tau)) \quad (18)$$

¹⁵In general, these have a similar form to the deterministic functions for the period parameters, X_t , in Section 2.4.1.

¹⁶Typically, cohort parameters for the last few years of birth are not estimated due to the lack of data, for instance, see Renshaw and Haberman (2006).

where

$$\mathbb{P}_{\tau-y,s} \equiv \prod_{r=0}^{s-1} (1 - D_{\tau-y+r}) \quad (19)$$

$$\begin{aligned} \mathbb{E}_{\tau}^{\mathbb{P}} \gamma_y &\equiv M(y, \tau) \\ &= \sum_{s=0}^{\infty} \mathbb{P}_{\tau-y,s} \rho^s \left[D_{\tau-y} \bar{\gamma}_y(\tau) + (1 - D_{\tau-y+s}) \beta (\tilde{X}_{y-s} - \rho \tilde{X}_{y-s-1}) \right] \end{aligned} \quad (20)$$

$$\begin{aligned} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_y) &\equiv V(y, \tau) \\ &= \sum_{s=0}^{\infty} \mathbb{P}_{\tau-y,s}^2 (1 - D_{\tau-y+s}) \rho^{2s} \sigma^2 \end{aligned} \quad (21)$$

for $y \leq Y$, where

$$M(y, \tau) = \rho^{y-Y} \left(M(Y, \tau) - \beta \tilde{X}_Y \right) + \beta \tilde{X}_y \quad (22)$$

$$V(y, \tau) = \frac{1 - \rho^{2(y-Y)}}{1 - \rho^2} \sigma^2 + \rho^{2(y-Y)} V(Y, \tau) \quad (23)$$

for $y > Y$. In this,

- D_x is the proportion of a cohort assumed to still be alive by age x ;
- ρ and σ^2 are the autocorrelation and variance of the AR(1) process assumed to be driving the evolution of the cohort parameters;
- \tilde{X}_y and β are the trends and drifts for the cohort parameters as defined above;¹⁷
- $\bar{\gamma}_y(\tau)$ are the estimates of the cohort parameters, fitted by the mortality model at time τ ; and
- \mathcal{F}_{τ} is the total information available at time τ , including observations of the cohort parameters up to year of birth y , i.e., $\{\bar{\gamma}_v(\tau) \mid v \leq y\}$.

¹⁷Note that the drifts, β , depend upon the arbitrary identifiability constraints chosen. In practice, we therefore impose a set of identifiability constraints such that $\beta = 0$ to simplify matters considerably.

In Hunt and Blake (2015a), it was shown that this framework allows the historical and projected cohort parameters to be treated consistently, without any sharp discontinuities in the uncertainty between them. It was also shown that these projections are well-identified, in the sense that they do not depend upon the arbitrary identifiability constraints made when fitting the model. In addition, it is shown in Hunt and Blake (2015d) that the Bayesian framework allows us to update estimates of the cohort parameters over a one-year period to proxy for the impact that new data would have on our parameter estimates, which is essential for risk management purposes. The Bayesian framework is therefore well adapted for use in a forward mortality context, and we will use it for all APC mortality models which include cohort parameters.

2.5 Estimation and projection

The framework described in Sections 2.3 and 2.4 is very general and can be used in conjunction with any APC mortality model for the force of mortality. To see this in practice, we consider estimating the forward mortality rates on male data for the UK for the period 1950 to 2011 and ages 50 to 100 from the Human Mortality Database (2014) for five different APC models:

1. the Lee-Carter (“LC”) model of Lee and Carter (1992);
2. the “CBDX” model discussed in Hunt and Blake (2015f), which extends the Cairns-Blake-Dowd model of Cairns et al. (2006a) with a static age function and uses a log-link function;
3. the “classic APC” model of Hobcraft et al. (1982) and others;
4. the “reduced Plat” (“RP”) model of Plat (2009) discussed in Hunt and Blake (2015g);¹⁸ and
5. the model produced by the “general procedure” (“GP”) in Hunt and Blake (2015b) for the data described above.

¹⁸That is, the simplification of the main model discussed in Plat (2009) without the third, high-age term or, equivalently, an extension of the CBDX model with a cohort term.

These models have the forms¹⁹

$$\ln(\mu_{x,t}) = \alpha_x^{(LC)} + \beta_x^{(LC)} \kappa_t^{(LC)} \quad (24)$$

$$\ln(\mu_{x,t}) = \alpha_x^{(CBDX)} + \kappa_t^{(CBDX,1)} + (x - \bar{x}) \kappa_t^{(CBDX,2)} \quad (25)$$

$$\ln(\mu_{x,t}) = \alpha_x^{(APC)} + \kappa_t^{(APC)} + \gamma_{t-x}^{(APC)} \quad (26)$$

$$\ln(\mu_{x,t}) = \alpha_x^{(RP)} + \kappa_t^{(RP,1)} + (x - \bar{x}) \kappa_t^{(RP,2)} + \gamma_{t-x}^{(RP)} \quad (27)$$

$$\ln(\mu_{x,t}) = \alpha_x^{(GP)} + \sum_{i=1}^4 f^{(GP,i)}(x) \kappa_t^{(GP,i)} + \gamma_{t-x}^{(GP)} \quad (28)$$

The parameters in these models have been estimated by fitting the model to the UK population data described above. These fitted parameters have, in turn, been used to estimate the parameters of the time series processes discussed in Sections 2.4.1 and 2.4.2 for κ_t and γ_y (if applicable). Using these parameter estimates, we can calculate forward mortality rate surfaces in the real-world measure using Equation 12.

These models have been chosen to give a reasonable cross section of the different APC mortality models which could be used in practice. One of the advantages of the forward mortality rate framework described in this paper is that it allows for consistency between the model of the force of mortality and the forward mortality surface. Consequently, as a check, we compare these forward surfaces of mortality for each model to the mean mortality rates calculated using Monte Carlo simulations (shown in Figure 1 for the GP model) and find that the small difference between the two is explained by sampling error in the simulations.

¹⁹See Hunt and Blake (2015b) for full details of the construction of the GP model. For all models, we also select age functions which are normalised so that $\sum_x |\beta_x| = \sum_x |f(x)| = 1$. This involves either including normalisation constants or choosing age functions which are “self-normalising” in the sense of Hunt and Blake (2015f). However, for clarity, these are not shown, although they are taken into account in the fitting algorithms.

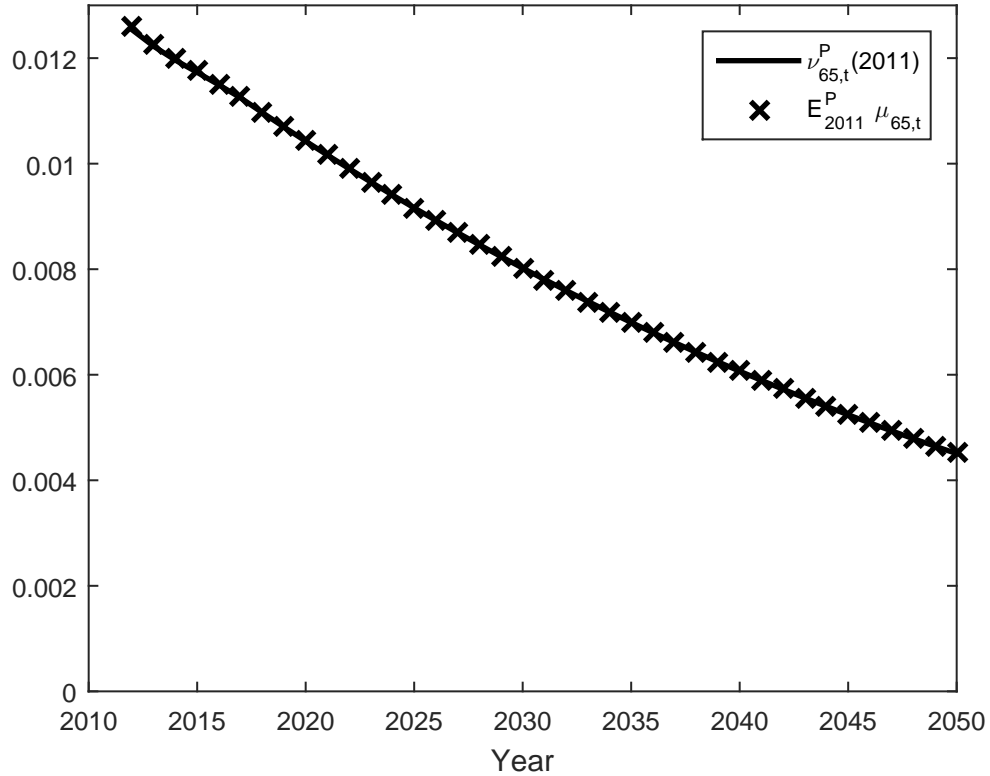


Figure 1: Difference between forward mortality rates and those obtained from Monte Carlo simulations using the GP model

3 Pricing securities and the market price of longevity risk

3.1 The market-consistent measure

In Section 2.4, we calculated mortality forward rates using the time series processes estimated from the fitted parameters. This means that the expectations in Equation 13 were calculated in the historical, real-world measure, \mathbb{P} .

It is obviously important that longevity-linked securities prices are consistent across different types of security in order to limit the potential for

pricing anomalies and arbitrage opportunities in the market. In addition, modern solvency regimes require that liability values and technical provisions for pension schemes and insurers must also be consistent with market prices. Identifying a suitable market-consistent measure, \mathbb{Q} , is therefore a critical component of the forward mortality framework.

The starting point of modern financial theory is to assume that the financial markets are “complete” in the sense that every financial claim in them can be hedged perfectly using tradable assets. In complete markets, the market-consistent measure exists and is unique. Derivative securities in complete markets can be perfectly replicated using these underlying securities without risk (and hence these measures are also referred to as “risk-neutral”) and the costs of these hedging strategies give the derivatives their unique prices. Complete markets are also free from arbitrage, since all prices can be derived using these underlying hedging strategies and any deviation from these prices will be arbitrated away by informed investors. The assumption of market completeness is a reasonable one in many contexts, such as developed markets for equities and interest rates in large and advanced economies.

However, the market for longevity risk is not complete. Not only are there insufficient tradable longevity-linked securities to fully replicate all financial claims, there are almost no longevity-linked securities being actively traded, full stop. Therefore, defining a market-consistent measure for longevity risk is a major problem for all mortality models which seek to price longevity-linked securities.

Some studies, for instance Schrager (2006), assume a priori that any market will be risk-neutral with respect to longevity risk and therefore that the historical and market-consistent measures are equal. We believe this is unlikely, given that any market in longevity risk is likely to be dominated by parties that suffer financially from rising life expectancy (see Loeys et al. (2007)) and therefore will be generally seeking to hedge the risk of future improvements in mortality rates.

In light of this absence of information, Barrieu et al. (2012, p. 224) suggested that the real-world measure must play a key role in the definition of any market-consistent measure:

What will be a good pricing measure for longevity? It is expected that the historical probability measure will play a key role, due to the reliable data associated with it. Therefore, it seems natural to look for a pricing probability measure equivalent to the historical probability measure. Important factors to consider are that a relevant pricing measure must be: robust with respect to the statistical data, and also compatible with the prices of the liquid assets quoted in the market. Therefore, a relevant probability measure should make the link between the historical vision and the market vision. Once the subsets of all such probability measures that capture the desired information are specified, a search can commence for the optimal example by maximising the likelihood or the entropic criterion.

We agree with this analysis, and use the Esscher transform to define a market-consistent measure that is equivalent to the real-world measure and that satisfies many of these desirable properties. This transformation is relatively parsimonious, with a small number of free parameters which can be calibrated using any market information we possess. Below, we further show that the Esscher transform gives us closed form expressions for the market-consistent forward mortality rates as shown below, and therefore is relatively straightforward to implement and robust to calibrate to data.

The Esscher transform has often been used in securities pricing in imperfect markets since the work of Gerber and Shiu (1994). As discussed in Kijima (2005), it is related to other widely used distortion methods for adjusting to a risk-neutral measure, such as the Wang transform (developed in Wang (2000, 2002) and Cox et al. (2006), and used in Denuit et al. (2007) for example), and the Sharpe ratio in modern financial theory (used in Milevsky et al. (2005) and Loeys et al. (2007)). It is also consistent with pricing in the real-world measure for an individual with an exponential utility function, as discussed in Milidonis et al. (2011).

For a risk $X_{x,t}$ in the \mathbb{P} measure, the general Esscher transform to the \mathbb{Q} measure can be defined by

$$\mathbb{E}^{\mathbb{Q}} X_{x,t} = \frac{\mathbb{E}^{\mathbb{P}} [X_{x,t} \exp(-Z_{x,t})]}{\mathbb{E}^{\mathbb{P}} \exp(-Z_{x,t})} \quad (29)$$

where $Z_{x,t}$ is a random variable containing the parameters defining the market-consistent measure.

In the context of mortality forward rates, we choose $X_{x,t} = \mu_{x,t} = \exp(\eta_{x,t})$ and correspondingly define

$$Z_{x,t} = \boldsymbol{\lambda}^\top \boldsymbol{\kappa}_t + \lambda^\gamma \gamma_{t-x} \quad (30)$$

where $\boldsymbol{\lambda}$ is an $(N + 1)$ column vector. Hence, there are $N + 1$ parameters (which we refer to collectively as $\lambda^{(j)}, j \in \{1, \dots, N, \gamma\}$), which correspond to the N age/period terms (in the vector $\boldsymbol{\lambda}$), and the cohort term (with single parameter $\lambda^{(\gamma)}$) in the general APC mortality model in Equation 2. It is important to note that the values found for these parameters will depend upon the specifics of the underlying model, and so are not comparable between different models.

Due to the paucity of genuine market information to price longevity risk, one might have a natural inclination to prefer simpler models, such as the LC model (which has only one free parameter for the Esscher transform). Such models could be felt to be more parsimonious, having fewer market prices for longevity risk and therefore requiring fewer market prices for longevity-linked securities in order to calibrate the market-consistent measure. For example, calibrating the LC model would require only one market price in order to calibrate the market-consistent measure, whilst calibrating the GP model in Section 2.5 requires four market prices. Using overly simple models, however, would be a mistake which can lead to unreasonable prices for other longevity-linked securities as shown in Section 3.3.

Using the Esscher transform with Equation 11 and this definition for $Z_{x,t}$

gives

$$\begin{aligned}
\nu_{x,t}^{\mathbb{Q}}(\tau) &= \mathbb{E}_{\tau}^{\mathbb{Q}} \mu_{x,t} \\
&= \mathbb{E}_{\tau}^{\mathbb{Q}} \exp(\eta_{x,t}) \\
&= \frac{\mathbb{E}_{\tau}^{\mathbb{P}} \exp(-Z_{x,t}\eta_{x,t})}{\mathbb{E}_{\tau}^{\mathbb{P}} \exp(-Z_{x,t})} \\
&= \frac{\mathbb{E}_{\tau}^{\mathbb{P}} \exp(\alpha_x + (\boldsymbol{\beta}_x - \boldsymbol{\lambda})^{\top} \boldsymbol{\kappa}_t + (1 - \lambda^{\gamma})\gamma_{t-x})}{\mathbb{E}_{\tau}^{\mathbb{P}} \exp(-\boldsymbol{\lambda}^{\top} \boldsymbol{\kappa}_t - \lambda^{\gamma}\gamma_{t-x})} \\
&= \exp \left(\alpha_x + \boldsymbol{\beta}_x^{\top} \mathbb{E}_{\tau}^{\mathbb{P}} \boldsymbol{\kappa}_t + \frac{1}{2} \boldsymbol{\beta}_x^{\top} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t) \boldsymbol{\beta}_x + \mathbb{E}_{\tau}^{\mathbb{P}} \gamma_{t-x} \right. \\
&\quad \left. + \frac{1}{2} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_{t-x}) - \frac{1}{2} \boldsymbol{\beta}_x^{\top} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t) \boldsymbol{\lambda} - \frac{1}{2} \boldsymbol{\lambda}^{\top} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t) \boldsymbol{\beta}_x - \lambda^{\gamma} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_{t-x}) \right) \\
&= \exp \left(-\boldsymbol{\beta}_x^{\top} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t) \boldsymbol{\lambda} - \lambda^{\gamma} \mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_{t-x}) \right) \nu_{x,t}^{\mathbb{P}}(\tau) \tag{31}
\end{aligned}$$

due to the symmetry of $\mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t)$.

This gives us closed-form expressions which allow us to adjust the forward mortality rates in the real-world measure to a market-consistent measure. The existence of closed-form expressions is why we argued that the Esscher transform neatly complements the forward mortality framework: these results could not have been achieved with alternative transformations to the market-consistent measure. Since we have already found expressions for $\mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t)$ and $\mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_y)$, transforming the forward mortality surface in the real-world measure into a market-consistent measure is simply a matter of finding the values of free parameters of the Esscher transform. This can be done if we have sufficient prices for longevity-linked securities, as discussed in Section 3.2 below.

Through the analogy with utility pricing and the Sharpe ratio, we refer to the parameters of the Esscher transform as the “market prices of longevity risk” associated with each of the age/period and cohort terms. For this analogy to be reasonable, we would anticipate that the parameters, $\lambda^{(j)}$, should be positive. However, this is not necessarily the case in the forward mortality framework, for the following reasons.

As discussed in Loeys et al. (2007), we anticipate that the marginal participant in the market for longevity-linked securities will be a life insurer

seeking to hedge longevity risk. Such a life insurer will be averse to longevity risk, and so, we would expect the market-consistent forward mortality rates to be lower than those in the real-world measure

$$\nu_{x,t}^{\mathbb{Q}}(\tau) \leq \nu_{x,t}^{\mathbb{P}}(\tau)$$

In order for this to be true,

$$\begin{aligned} \exp\left(-\boldsymbol{\beta}_x^\top \mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t)\boldsymbol{\lambda} - \lambda^\gamma \mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_{t-x})\right) &\leq 1 \\ \Rightarrow \boldsymbol{\beta}_x^\top \mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t)\boldsymbol{\lambda} + \lambda^\gamma \mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_{t-x}) &\geq 0 \end{aligned}$$

Since $\mathbb{V}ar_{\tau}^{\mathbb{P}}(\boldsymbol{\kappa}_t)$ is a positive definite matrix and $\mathbb{V}ar_{\tau}^{\mathbb{P}}(\gamma_y) \geq 0$, this will certainly be true if $\lambda^\gamma > 0$ and the elements of $\boldsymbol{\lambda}$ are also positive. However, individual market prices of longevity risk can be negative, whilst still ensuring that hedgers pay a positive price to transfer longevity risk overall. Since some market prices can be negative, the term “market prices” might be considered misleading. Although we shall refer to these parameters as market prices in this paper and in Hunt and Blake (2015d), it should be borne in mind that they are probably best thought of as simply parameters in the Esscher transform in Equation 29 rather than true market prices of longevity risk based on an expected utility approach (such as that discussed in Zhou et al. (2015)).

The Esscher transform approach has some other practical advantages, beyond the existence of closed-form expressions for the forward mortality rates. The forward mortality surface in the real-world measure will be updated only infrequently, typically once every year when new mortality data is released. However, market information will need to be updated far more frequently, especially as the market for longevity-linked securities develops. It is desirable in practice to be able to take the (infrequently changing) \mathbb{P} -measure forward mortality surface and make relatively simple adjustments to this to reflect changing market information, rather than having to re-estimate the model completely every time the pricing information changes.

However, a limitation of the forward mortality framework outlined in this study is that it is currently unable to price longevity-linked securities with optionality, for example, a call option on mortality rates. In order to do this, the dynamics of mortality rates in the market-consistent measure would need to be specified, in addition to simply the expectation, $\mathbb{E}_{\tau}^{\mathbb{Q}} \mu_{x,t}$. We leave the

extension of the forward mortality framework to the inclusion of longevity-linked options to future work.

We also note that, looking solely at the age/period terms, Equations 16 and 17 imply

$$\begin{aligned}\beta_x^\top \mathbb{E}^\mathbb{P}_\tau \kappa_t + \beta_x^\top \text{Var}_\tau(\kappa_t) \lambda &= \beta_x^\top [\kappa_\tau + \mu \chi_{\tau,t} + (t - \tau) \Sigma \lambda] \\ &= \beta_x^\top [\kappa_\tau + \hat{\mu} \chi_{\tau,t}]\end{aligned}$$

since $t - \tau$ is always one of the deterministic functions in $\chi_{\tau,t}$. Hence, we see that for an age/period model such as the LC and CBDX models, the Esscher transform to the market-consistent measure is equivalent to making an adjustment to the drift of the random walk in Equation 14. This approach is developed further in Hunt and Blake (2015e). In this form, the use of the Esscher transform can be compared with some of the other approaches that have been suggested in previous studies. For instance, Loeys et al. (2007) suggested that the price of a q-forward should be calculated as

$$q^f = (1 - (t - \tau) \tilde{\lambda} \sigma^2) q^e$$

where σ^2 is defined as the annual volatility of the mortality rate, i.e., $\sigma^2 = \text{Var}^\mathbb{P}(\ln q)$. We can compare this pricing formula to what our forward mortality framework would give were we to use the LC model as the underlying mortality model. This has one period function, κ_t , with one associated market price of risk, λ . From Equation 31 applied to the LC model, we find

$$\nu_{x,t}^\mathbb{Q}(\tau) = \exp(-(t - \tau) \beta_x \Sigma \lambda) \nu_{x,t}^\mathbb{P}(\tau)$$

We can therefore see that the pricing formula in Loeys et al. (2007) is similar in form to Equation 31, although based on forward contracts on probabilities of death, $q_{x,t}$, rather than the longevity-zeros which are used as the underlying securities in this study.

Cairns et al. (2006a) adjusted the drift of the random walk used to project the period functions directly, in order to incorporate market prices for longevity risk without recourse to the Esscher transform

$$\mu^\mathbb{Q} = \mu^\mathbb{P} - C \tilde{\lambda}$$

where $CC^\top = \Sigma$ and $\boldsymbol{\lambda}$ is a vector of the market prices of risk. If such an approach were to be used for the CBDX model in a forward mortality rates framework such as above, we would find market-consistent forward mortality rates

$$\nu_{x,t}^{\mathbb{Q}}(\tau) = \exp\left(-(t-\tau)\boldsymbol{\beta}_x^\top C\tilde{\boldsymbol{\lambda}}\right)\nu_{x,t}^{\mathbb{P}}(\tau)$$

Therefore, we see that the approach used in Cairns et al. (2006a) is equivalent to that used in this study, except using $C\tilde{\boldsymbol{\lambda}}$ instead of $\Sigma\boldsymbol{\lambda}$. Equating these gives

$$\begin{aligned} C\tilde{\boldsymbol{\lambda}} &= \Sigma\boldsymbol{\lambda} \\ \tilde{\boldsymbol{\lambda}} &= C^\top\boldsymbol{\lambda} \end{aligned}$$

Hence, the more rigorous forward mortality framework defined in this study achieves results which are consistent with those of Cairns et al. (2006a), but is also able to justify the otherwise ad hoc adjustments to the drift made in that study.

3.2 Calibration of the market-consistent measure

As has been mentioned previously, a major problem with forward mortality models is the lack of market information to specify the market-consistent measure. An advantage of using the forward mortality framework described in this study is that, rather than requiring sufficient market prices to define the full forward mortality surface, we require only $N + 1$ prices to uniquely specify the market prices of longevity risk used in the Esscher transform. This substantially reduces the market information required.

However, even this simplification is unlikely to be adequate at present, given the paucity of traded longevity-linked securities. Many of those which do exist, such as the extreme mortality bonds listed in Lane (2011), are not suitable as they involve options on mortality rates which cannot be priced using the forward mortality framework as proposed here.²⁰ For illustrative purposes, we will demonstrate how the forward mortality rate framework could be calibrated with respect to the sort of information which is available

²⁰We extend the forward mortality framework to allow for the valuation of longevity-linked options in Hunt and Blake (2015e).

currently or is likely to be available in the foreseeable future, and how this “external” market in longevity risk could be supplemented by use of an “internal” market for longevity risk based on the assumptions used to value and reserve for longevity risk within a life insurer.²¹

3.2.1 External market

A number of “external” markets exist for products which depend upon longevity, for instance the markets for endowment assurances and individual annuities. These were used to provide market information for pricing longevity risk in Bayraktar and Young (2007) and Bauer et al. (2008). However, both of these products are sold to individuals, and therefore are subject to idiosyncratic mortality risk as well as systematic longevity risk, which makes them unsuitable for use in a forward mortality rate framework, as discussed by Norberg (2010). Furthermore, insurers will include loadings for expenses and other risks, in addition to longevity risk when pricing these products, which makes using them to calibrate a forward mortality model problematic.

Instead, any forward mortality model will need to be calibrated using securities dependent on aggregate mortality rates (preferably from national populations) rather than those that are sold to individuals. Such securities are also more likely to be traded, thereby giving informed and responsive market prices. The problem remains, however, that there is currently no actively-traded market in such securities which can be used to provide the pricing information required to calibrate the market-consistent measure.

To date, probably the most active market in longevity-linked securities has been that for bespoke longevity swaps (see Hunt and Blake (2015k)). A longevity swap is an agreement between two parties to swap a series of cashflows - a fixed leg based on the best estimate of the survivorship of a cohort but then increased by a constant percentage (the swap margin) and a floating leg based on the actual survivorship observed for the cohort. A bespoke longevity swap is one which is tailored to the characteristics of a spe-

²¹In a sense, the difference between the external and internal markets for longevity risk could be compared to the difference between using mark-to-market and mark-to-model valuation methods when valuing securities in company accounts, depending upon whether deep and liquid markets exist for them.

cific population such as a pension scheme. As such, bespoke longevity swaps are unlikely to be widely traded, and act more as customised reinsurance contracts than standardised longevity-linked securities which could form the basis for a market in longevity risk. In contrast, an index-based swap, such as that described in Dowd et al. (2006), is one where the cohort in question is from a national population. Although index-based longevity swaps have not yet been widely traded, the development of the bespoke longevity swap market to date implies that, if a market in longevity risk does develop in the near future, it is likely that index-based swaps will form a key component of it.

For illustrative purposes, we therefore assume the existence of a single index-based longevity swap, which we believe might be typical of the sort of security which may be traded during the early stages of the development of an external market in longevity-linked securities. We assume that this index-based longevity swap has been written on a standard cohort of men in the UK aged 65 in 2011 and has a term of 35 years (i.e., until the cohort is aged 100). The floating leg of this swap will therefore have the value

$$\sum_{t=1}^{35} {}_tP_{65,\tau}^{\mathbb{Q}}(\tau)B(\tau, \tau + t)$$

i.e., the same price as a series of the longevity zeros discussed in Section 2.2. The fixed-leg cashflows will reflect a typical “best estimate” agreed between the contracting parties when the swap is initiated. For illustrative purposes, we assume these cashflows are set by calculating the survivorship of the reference cohort using the fitted mortality rates in $\tau = 2011$ projected using the “CMI Projection Model” (Continuous Mortality Investigation (2009a,b, 2013)) with a “long-term rate of improvement” assumption of 1.5% p.a..²² We denote the survival probabilities of the reference cohort from time τ to $\tau + t$ using this assumption as ${}_t\tilde{P}_{65,\tau}(\tau)$. While there is currently no active market in index-based swaps, this assumption is typical of those used to define the fixed leg of bespoke longevity swaps in our experience. These cashflows are then increased by a swap premium of 4%, which is a typical level on bespoke

²²The use of the CMI Projection Model in this context is purely illustrative and should not imply that we believe that this is the best model to use for pricing longevity-linked securities, although it is typical of what has been used in practice in our experience.

swaps in our experience.

The price of the swap is therefore

$$\sum_{t=1}^{35} \left({}_tP_{65,\tau}^{\mathbb{Q}}(\tau) - 1.04 {}_t\tilde{P}_{65,\tau}(\tau) \right) B(\tau, \tau + t) \quad (32)$$

and will be zero at time τ . We therefore calibrate the market prices of risk to impose this using standard numerical optimisation algorithms. In these calculations, we assume a flat real yield of 1.0% p.a. for the zero-coupon bond prices, $B(\tau, \tau + t)$

For models with only one source of risk (for instance, the LC model), this single, external price is sufficient to specify the single market price of longevity risk uniquely. For more complicated models, with multiple risk sources, we require additional prices in order to specify the market prices of longevity risk.

3.2.2 Internal market

We observe that, while genuine market information is in scarce supply, many insurance companies will effectively have an internal market for longevity risk due to the cross-subsidies between different lines of business with different exposures to longevity risk. For instance, an insurer which writes both annuity and life assurance lines of business has, de facto, established an internal market for longevity risk due to the presence of natural hedging between the two lines of business, as discussed in Cox and Lin (2007). The “price” of longevity risk in this internal market will find expression in the mortality improvement assumptions used in the pricing and reserving for these different lines of business. It is therefore natural to use these “internal” market signals to supplement those coming from the genuine external market if there are insufficient traded longevity-linked securities to define the market-consistent measure.

Alternatively, an insurer may develop an “internal” price for longevity risk by analysing the cost of longevity reinsurance via bespoke longevity swaps. Although these contracts do not solely transfer longevity risk - they also

transfer basis and idiosyncratic risks - they could still give some indication of a price for the systematic longevity risk present, and so be used to calibrate the market-consistent measure.

For example, we assume that the forward mortality framework is being used by an organisation with an internal, deterministic assumption that constitutes their “house view” of mortality improvements. This house view would then feed through into the assumptions used in pricing and reserving, and inform those assumptions that are used for accounting and regulatory purposes if there is sufficient flexibility in how these are set. The existence of such a house view would therefore determine the organisation’s appetite for longevity risk across multiple lines of business and so underpin the “internal” market for longevity risk.

To illustrate the sort of internal market that might be considered typical, we assume a house view that mortality rates improve in line with the projections from the CMI Projection Model with a long-term rate of improvement of 1.75%.²³ Again, this is in line with the sort of assumptions used to reserve for and price annuity business in the UK in our experience. In order to translate this house view into the market prices of longevity risk in our forward mortality framework, we try to minimise the (weighted) relative distance between the surface of probabilities of dying given by the internal assumption, $\tilde{q}_{x,t}$, and those given the forward mortality surface in the \mathbb{Q} -measure

$$\mathcal{Q}_{x,t}(\tau) = 1 - \exp(-\nu_{x,t}^{\mathbb{Q}}(\tau))$$

at certain key ages, subject to the swap also being priced fairly at time, τ ,

²³This value of 1.75% can be compared with the assumption of a long-term rate of improvement of 1.5% used for the fixed leg of the index-based longevity swap above. The long term rate of improvement is likely to be higher on an annuity reserving basis than for valuing a longevity swap, since it is common practice, in our experiences, for annuity providers to include an implicit margin for prudence in their mortality projection. In contrast, the assumption used in a longevity swap typically reflects a best estimate of future mortality improvements and risk is explicitly allow for via the swap premium rather than an implicit margin in the mortality assumption.

i.e.,

$$\min_{\lambda} \sum_{t,x \in \mathcal{X}} B(\tau, \tau + t) \frac{(\tilde{q}_{x,t} - Q_{x,t})^2}{\tilde{q}_{x,t}}$$

subject to Equation 32 = 0

where $\mathcal{X} = \{50, 55, 60, 65, 70, 75, 80\}$. This procedure is equivalent to determining the market-consistent measure by reference to an external market in q-forwards, as proposed in Coughlan et al. (2007) and discussed in Section 3.3.2 below, if such a market existed. We consider these key ages partly to ensure that the forward mortality surface in the market-consistent measure is biologically reasonable over a wide age range and because, if a market in q-forwards does emerge, it is at these ages where the market is likely to be most liquid (see Li and Luo (2012)). Therefore, the use of the internal market for longevity risk is simply a proxy for information from an external market for longevity risk, and will be supplanted should a genuine external market develop.

We use these assumptions for the external and internal markets for longevity risk in order to calibrate the parameters of the Esscher transform for all five models described in Section 2.5. These parameters, along with the forward mortality surfaces obtained in Section 2.5, allow us to construct the forward mortality surface in the market-consistent measure, which can then be used to value other longevity-linked liabilities and securities in a market-consistent fashion.

3.3 Pricing longevity-linked securities

The forward mortality framework described above provides a single surface of forward mortality rates, calibrated from all the available information on longevity-linked securities. It can, therefore, be used to value any other longevity-linked securities and give prices consistent with those observed. We demonstrate this for a range of different longevity-linked securities below.

3.3.1 Survivor derivatives

Longevity zeros and s-forwards

In Section 2.2, we defined the forward mortality rates assuming the existence of a market in longevity zeros. These were used as they are the fundamental securities dependent upon the survivorship of a cohort of individuals, and can be used to construct more complicated survivor securities such as annuities and longevity swaps, as discussed below. Related to longevity zeros are “s-forwards”, as proposed in Dowd (2003), Blake et al. (2006) and the Life and Longevity Markets Association,²⁴ which are forward contracts defined on a longevity zero (and hence are more capital efficient).

From Equation 7, we can see that

$$\mathcal{S}_{x,t}(\tau) = {}_tP_{x,\tau}^{\mathbb{Q}} = \exp \left(- \sum_{u=1}^t \nu_{x+u,\tau+u}^{\mathbb{Q}}(\tau) \right)$$

where $\mathcal{S}_{x,t}(\tau)$ is the forward price of an s-forward at time τ , defined on a cohort aged x at τ , with a maturity of t years. Figure 2 shows s-forward prices defined on the cohort of individuals aged 65 in 2011 with different maturities.

As can be seen, most of the models give broadly comparable s-forward prices, especially those calibrated using the internal market information. We note that the LC model gives s-forward prices which are slightly different from these models, with higher probabilities of survival over the first few decades followed by a period of higher mortality rates (and hence a steeper gradient for the curve), but these are still biologically reasonable.

Annuities

The most relevant longevity-linked instruments for many life insurance companies are annuities. For the reasons discussed in Section 3.1 and Norberg (2010), individual annuities cannot be used to calibrate the forward mortality surface in the market-consistent measure, since the cashflows of these instruments are explicitly linked to the survivorship of a named individual and, hence, their prices include an allowance for individual mortality risk. In addition, they are not traded, and, therefore, cannot provide timely information on their values. However, when a life insurer reserves for a book

²⁴<http://www.llma.org/>

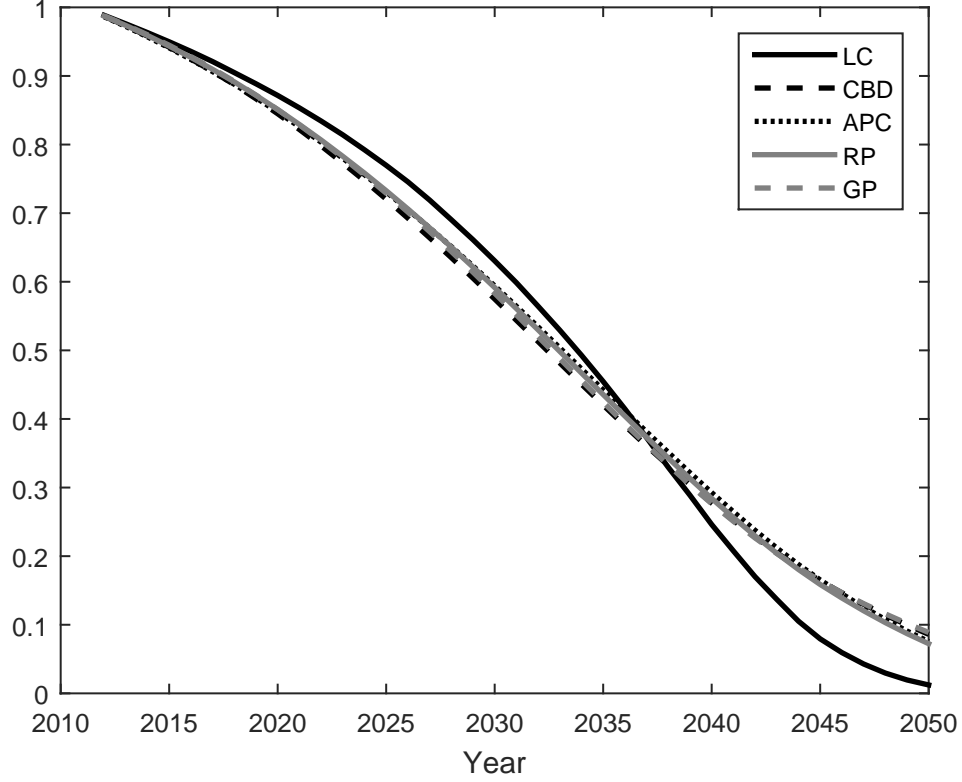


Figure 2: S-forward prices for five different mortality models

of annuities, the idiosyncratic mortality risks are diversifiable and so are not included in the value of any specific annuity but through the additional capital required for the book.²⁵ In addition, modern solvency regimes, such as Solvency II, require the best estimate of the liabilities in respect of annuity policies to be calculated using market-consistent assumptions. Therefore, the market-consistent forward framework could, potentially, be used as the basis for an insurer’s “internal model” under Solvency II, as discussed in EIOPA (2014).²⁶

²⁵There will therefore be a distinction between the price an annuity is sold to the public for and the amount it is reserved for by the life insurer, with the additional margin for idiosyncratic mortality risk charged to the individual forming part of the profit margin of the product.

²⁶This is discussed further in Hunt and Blake (2015j).

The value of an annuity can be directly constructed from a portfolio of longevity zeros using

$$a_x(\tau) = \sum_{t=0}^{\infty} {}_tP_{x,\tau}^{\mathbb{Q}}(\tau) B(\tau, \tau + t) \quad (33)$$

To calculate the values of longevity zeros beyond the maximum age in our data, we use the topping out procedure of Denuit and Goderniaux (2005). We therefore see that annuity values are very closely related to the swap price given in Equation 32. We calculate annuity prices²⁷ for men at different ages in 2011 using the five different models, and the results are shown in Figure 3.

We can see from this that the different models give broadly similar annuity values. This is not surprising given that they all use the same external market information (i.e., the swap price) in order to calibrate the market-consistent measure. Indeed, all the models give exactly the same value for an annuity at age 65, since this is determined by the swap price we have assumed and an annuity is equivalent to the floating leg of a longevity swap. However, the annuity values given by different models diverge slightly as we move away from this fixed reference point, with the LC model giving lower annuity values at higher ages than the other models.

Index-based longevity swaps

We can also use these results to investigate the potential pricing of index-based longevity swaps at different ages. Extending the definition of the swap value in Equation 32 for different ages to

$$0 = \sum_{t=1}^{35} \left({}_tP_{x,\tau}^{\mathbb{Q}}(\tau) - (1 + \pi) {}_t\tilde{P}_{x,\tau}(\tau) \right) B(\tau, \tau + t) \quad (34)$$

we can use the same “best estimate” assumption based on the CMI Projection Model for the fixed legs of the swaps, to calculate the implied swap

²⁷Annuities are valued using a real discount rate of 1% p.a..

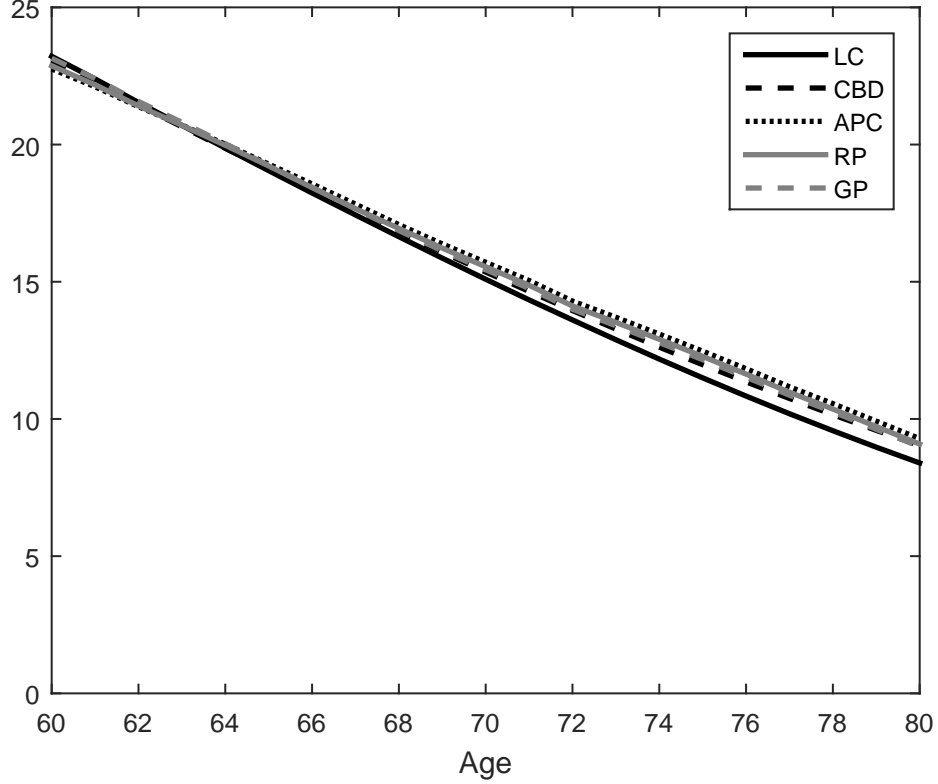


Figure 3: Annuity values for five different mortality models

premium, π , on index-based longevity swaps at different ages. The implied swap premiums are shown in Figure 4.

As can be seen, the behaviour of the swap premium depends strongly upon the model being used. For the classic APC, RP and GP models, which include a cohort term, the swap premium slightly increases with age, from around 4% at age 65 to around 6% between ages 75 and 80 (note that a value of 4% was assumed at age 65). Swap premiums for the CBDX model decrease slowly with age, to around 3% at age 75. However, for all of these models, the swap premium remains positive and do not appear unreasonable at any age.

In contrast, the LC model gives swap premiums which decrease rapidly

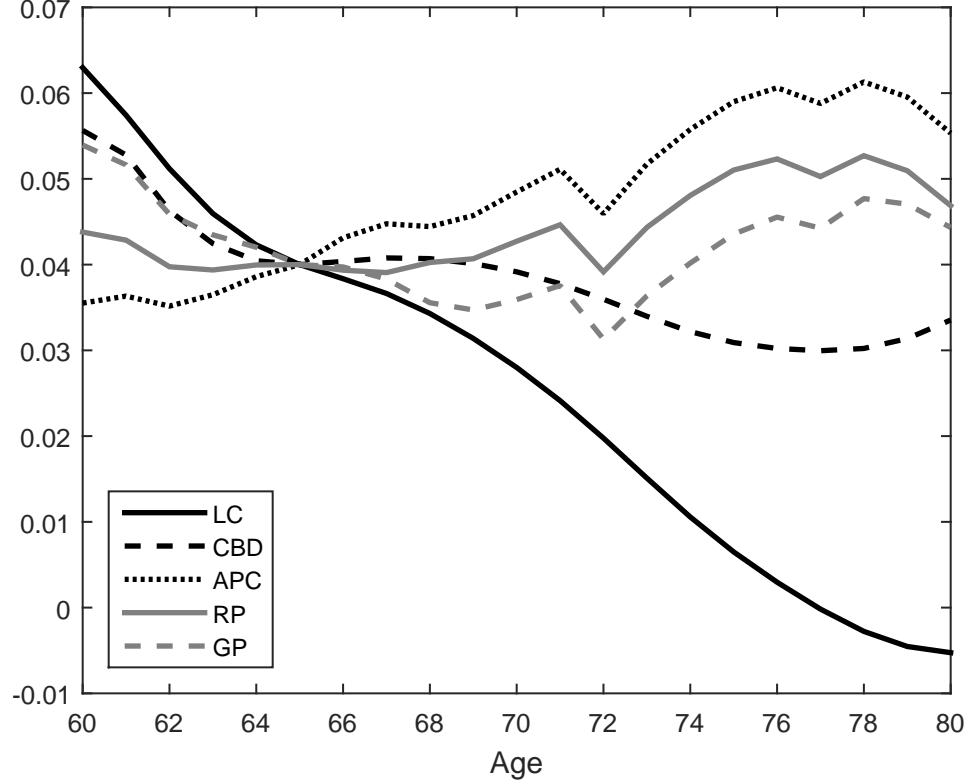


Figure 4: Swap premiums for five different mortality models

with age, giving negative swap premiums at higher ages (i.e., a premium would be paid to receive the floating payments on the swap) which does not appear reasonable. This is because the LC model gives relatively low values for annuities at higher ages - lower than would be found using the deterministic CMI Projection Model. We therefore see that there is a trade-off. On the one hand, we would like to use simple models which have relatively few free parameters and so are simple to calibrate from sparse data (and, in particular, would avoid the use of an internal market for longevity risk). On the other hand, we also need to obtain plausible prices for different longevity-linked liabilities and securities and across a wide range of ages.

3.3.2 Other longevity-linked securities

A number of other longevity-derivatives not based on the survivorship of a cohort have been proposed, and these can also be valued using the forward mortality framework proposed here. A number of these are illustrated below. However, the important point to note is that any security which does not have a non-linear payoff (i.e., which is not an option) can be valued using the forward mortality framework proposed in this paper.

q-forwards

Forward contracts on future probabilities of death, known as “q-forwards”, were introduced in Coughlan et al. (2007) and represent another, distinct, family of potential longevity-linked securities. There have been a number of hedging transactions using q-forwards, as discussed in Blake et al. (2013), and so q-forwards are one of the major contenders to form the basis of a traded market for longevity risk if it develops. In addition, the internal market assumption, used in Section 3.2 to calibrate all of the models other than the LC model, implicitly makes use of a market for q-forwards, albeit one that is internal to the life insurer rather than an externally traded market.

Values for q-forwards at age 75 and different maturities, calculated using the forward mortality models, are shown in Figure 5, along with the $q_{x,t}$ values projected using the CMI Projection Model. For the models which used the internal market assumption to calibrate the market-consistent measure, we see that the q-forward values are broadly consistent with those from the CMI Projection Model. However, they are not identical, since the calibration process also has to match the swap price exactly and minimise the difference in q-forward prices at ages other than 75. However, because the GP model has more market prices of risk to calibrate, it achieves a slightly closer fit to the internal market assumption than the other models, including the cohort effect observed around 2025 (i.e., for cohorts born around 1950).

In contrast, the LC model gives q-forward values which are very different from those of the other models, with implausibly rapid decreases in q-forward values. Again, this is because, with a single market price for longevity risk, the LC model has to severely distort the forward mortality surface in the

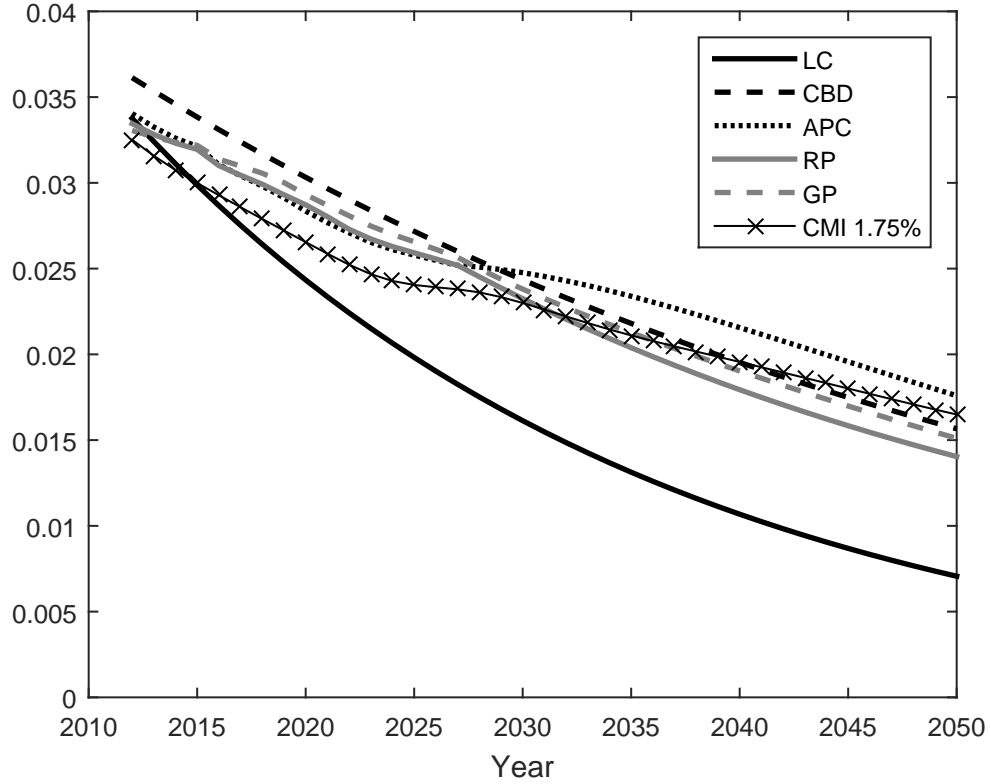


Figure 5: q-forward prices at age 75 for five different mortality models

real-world \mathbb{P} -measure in order to price the longevity swap. It cannot ensure that mortality rates across a wide range of other ages and years behave in a plausible fashion in the market-consistent measure. We therefore see that more sophisticated underlying APC mortality models, as well as being able to incorporate pricing information from a wider range of sources, will also tend to give more biologically-reasonable forward surfaces for mortality in the market-consistent measure.

e-forwards

Period life expectancy is a very commonly used aggregate measure of mortality rates, since it can be calculated easily from observed data and can

be compared across different populations. It is, therefore, natural to consider its use as an index for longevity risk transfer, based on the suggestion of Denuit (2009). In particular, we consider a market in forwards on period life expectancy, which we refer to as “e-forwards” (from the demographic symbol for period life expectancy). Using the forward mortality framework, we calculate forward period life expectancies as

$$\mathcal{E}_{65,t}(\tau) = 0.5 + \sum_{u=1}^{\infty} \exp \left(- \sum_{v=1}^u \nu_{65+v,t}^{\mathbb{Q}}(\tau) \right)$$

Figure 6 shows the forward period life expectancies at age 65 from each of the five models in the market-consistent measure.

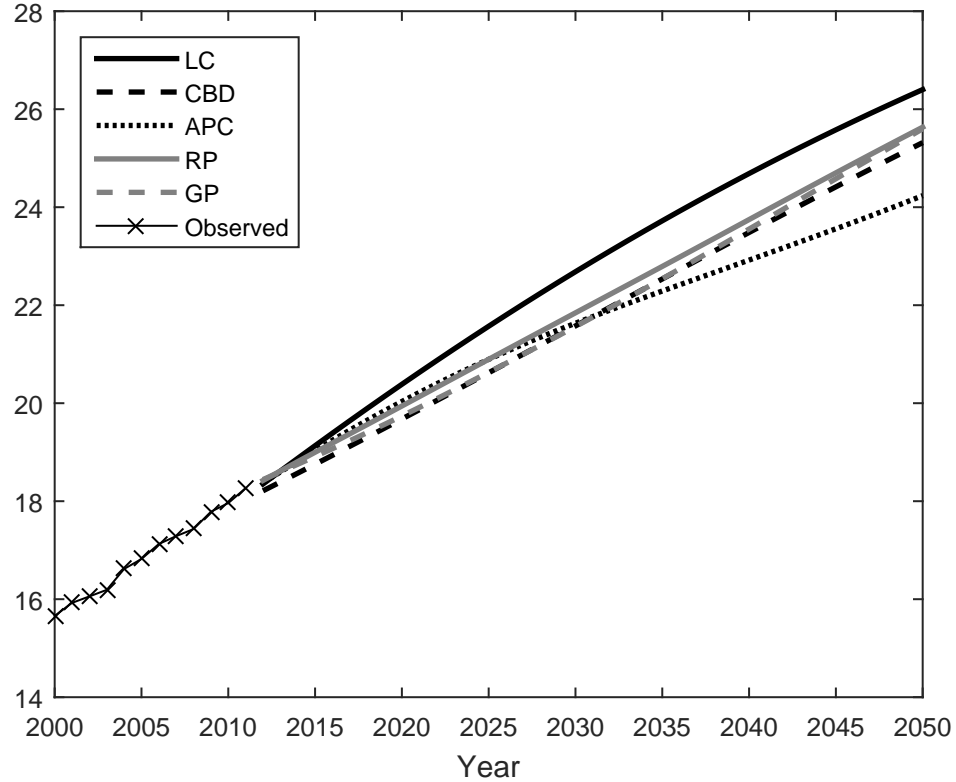


Figure 6: Period life expectancies at age 65 for five different mortality models

We note that all of the models give forward period life expectancies which can be considered biologically reasonable and consistent with the findings of Oeppen and Vaupel (2002), i.e., that they increase roughly linearly. Life expectancies from the LC model increase slightly faster than the other models, which otherwise give broadly consistent forward values. This is because of the use of the internal market to calibrate these other models, ensuring greater consistency between their forward mortality surfaces.

k-forwards

In Hunt and Blake (2015h), we discussed how the indices based on the observed rates of improvement in mortality rates, such as the indices which were defined in the construction of the Swiss Re Kortis bond, could potentially form the basis for a market in longevity risk. Improvement rates may be a natural basis for a market in longevity, as they are often used by actuaries to express long term assumptions regarding the evolution of mortality rates. Building on this, we also consider the forward value of the index for men in the UK defined by

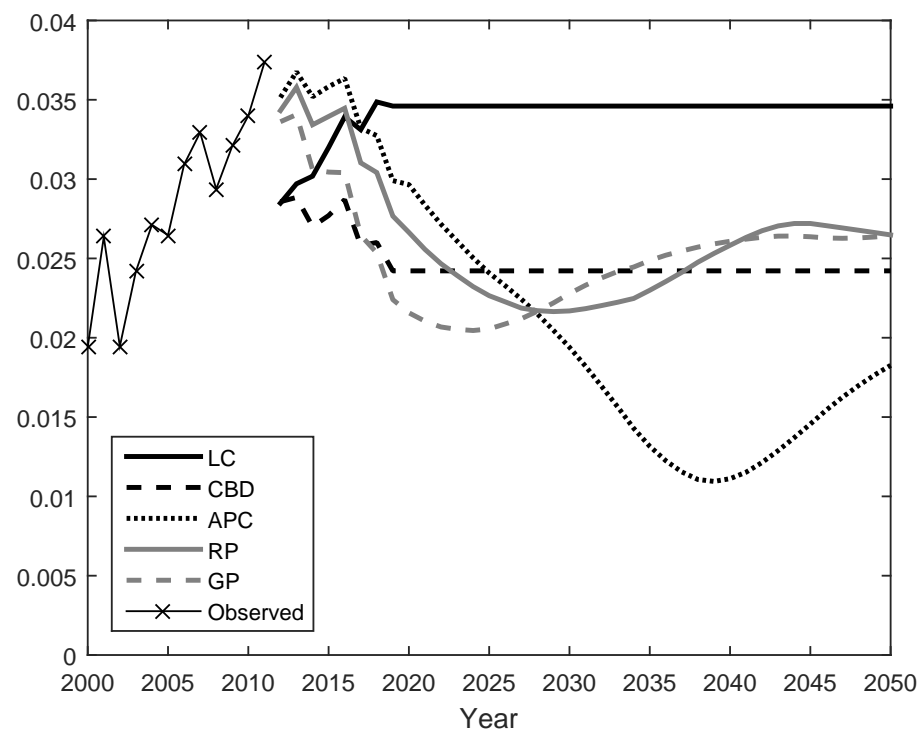
$$\mathcal{K}_t(\tau) = \frac{1}{11} \sum_{x=75}^{85} \left(1 - \left[\frac{\nu_{x,t}^{\mathbb{Q}}(\tau)}{\nu_{x,t-8}^{\mathbb{Q}}(\tau)} \right]^{\frac{1}{8}} \right)$$

This index was constructed to measure the average rate of improvement in mortality rates between ages 75 and 85 for men in the UK and so could be used for hedging or transferring longevity risk in a portfolio of annuities. Unlike the Kortis bond, however, we only consider an index constructed for a single population (i.e., men in the UK) rather than the difference between two populations, and only consider pricing the index rather than an option on the index.²⁸

In Hunt and Blake (2015h) it was suggested that forward contracts based on this Kortis index could form the basis of a market in longevity risk. We refer to such contracts as “k-forwards” in the same manner at q-, s- and e-forwards discussed above. Figure 7 shows the projected k-forward values in the market-consistent measure. As discussed in Hunt and Blake (2015h), the

²⁸See Hunt and Blake (2015h) for a further discussion of the Swiss Re Kortis bond and its construction.

Figure 7: Kortis index values for five different mortality models



Kortis index is designed to be very sensitive to the rates of improvement in longevity, which are determined by the drift, μ , of the random walk used for the period parameters. Indeed, for models which lack a cohort term, the drift in the random walk exactly determines the projected index values, and hence they are constant beyond 2020.²⁹ For the models which include cohort parameters, the value of the index in the short term depends strongly upon the cohort parameters fitted by the model, as discussed in Hunt and Blake (2015h), resulting in a distinctive curved pattern. In general, the models containing a cohort term give market-consistent assumptions for the rate of improvement in longevity which decrease from its currently observed level of around 3.5% to around 2% in 20 years' time. This is not surprising given this is broadly in line with the assumptions used to calibrate the market-consistent measure, i.e., the CMI Mortality Projection Model with a long term rate of improvement of either 1.5% or 1.75%.

As in the case of the q-forwards, the index values for the LC model show a very different evolution due to the limited ability of this model to both price the market information and give a biologically reasonable forward surface of mortality. However, the alternative models appear to give index values which are biologically reasonable and consistent with the historical, realised values for the k-forwards, which potentially means that forwards on the index could form a viable basis for a market in longevity risk.

Other longevity-linked securities

The forward mortality surface could also be used to value life assurance policies in the same manner. In conjunction with the results of Hunt and Blake (2015d), the forward mortality framework could therefore be used as a standard model for both the valuation of a life insurer's technical provisions and the assessment of longevity risk within them, in accordance with the Solvency II regulatory regime described in EIOPA (2014). We describe how this can be accomplished in Hunt and Blake (2015j). In addition, for life insurers writing both annuity and assurance policies, it may be desirable to value these consistently in the technical provisions, in order to achieve the

²⁹Before 2020, the Kortis index is based partly on projected and partly on observed mortality rates, and hence exhibits more variability than after 2020.

benefits from natural hedging discussed in Cox and Lin (2007).

Beyond the examples discussed above, the forward mortality framework could be used to value any longevity-linked security with a linear payoff in the underlying index. Hence, although the market for longevity-linked securities is in the early stage of development currently and it is unclear which form of securities will ultimately come to be traded, we believe that the framework described in this paper is flexible enough to be able to price any of them in a manner consistent with any other prices for longevity-linked liabilities and securities which are available.

As discussed previously, one disadvantage of any forward mortality rate framework as described in this study is that it cannot be used to value longevity-linked options, since it only looks at the expected mortality rates in the market-consistent measure. For example, it could not be used directly to value mortality catastrophe bonds, such as the Swiss Re Vita bond (discussed in Bauer and Kramer (2007)), Longevity Experience Options (described in Fétiqueau and Jia (2014)), bespoke index-based solutions (described in Michaelson and Mulholland (2014)), a guaranteed annuity option (discussed in Pelsser (2003) and Ballotta and Haberman (2006)) or a bond similar to the Kortis bond with the principal being a non-linear function of the index value. At the present time, we do not think that this is a fatal limitation of the forward mortality rate framework discussed here, as currently the market for longevity-linked securities is not sufficiently developed to allow a full calibration of the forward mortality rate surface, let alone the dynamics of the force of mortality in the market-consistent measure, which is required to model longevity-linked options. We extend the forward mortality framework developed here to be able to value longevity-linked options in Hunt and Blake (2015e).

4 Conclusion

The valuation of longevity-linked liabilities and securities requires us to predict future rates of mortality. Modern solvency regulations and the gradual emergence of a market in longevity-linked securities require these predictions to incorporate market information, in order to give prices for different securities which are consistent with those observed in the marketplace. As many

previous studies have shown, forward mortality models are ideally placed to achieve this.

We therefore believe that the answer to the titular question raised in Norberg (2010) - are forward mortality rates the way forward? - is yes. Nevertheless, it is important to take on board the criticisms of Norberg (2010) and to develop a framework specifically to model mortality rates, rather than borrow a pre-existing framework developed for interest rates and to define this framework using securities which do not depend on the idiosyncratic timing of individual deaths. This is because, with a properly developed framework, we can derive a model which is capable of capturing the complex dynamics of mortality rates, and so obtain consistency between models of the force of mortality and the forward mortality rates.

In this study, we have developed such a framework for forward mortality rates which is based upon the dynamics of the force of mortality given by the class of age/period/cohort mortality models. This framework has the advantage of being easier to estimate from historical data than existing models, with market information being incorporated via a relatively parsimonious transformation of the forward mortality rates in the real-world measure. The framework is also very flexible, as it can be used in conjunction with many of the most popular models of the force of mortality, such as those proposed in Lee and Carter (1992) and Cairns et al. (2006a).

We have shown how market information can be incorporated into the model and used the resulting forward mortality surface to value a range of existing and proposed longevity-linked securities. All of the prices calculated from the same model are consistent with each other, as they are derived from the same forward surface of mortality. This allows for a unified approach to the valuation of a wide range of liabilities and longevity-linked securities.

Finally, we note that the main virtue of forward mortality models is their ability to specify the dynamics of the forward mortality surface and, hence, their applicability to the assessment and management of longevity risk. We develop these themes in the second part of this study, in Hunt and Blake (2015d). Together, these two studies show that the framework proposed can provide an integrated solution to many of the valuation and risk management problems in respect of longevity risk that are faced by life insurance

companies.

A Identifiability and mortality forward rates

In Hunt and Blake (2015f) and Hunt and Blake (2015g), we discuss the identifiability issues in AP and APC mortality models, respectively. In particular, we find that almost all APC mortality models possess “invariant” transformations, i.e., transformations of the parameters of the model which leave the fitted mortality rates unchanged. In order to find a unique set of parameters, we impose a set of identifiability constraints on them. Typically, these are chosen to give a particular demographic significance to each term in the model. However, since any interpretation of demographic significance is subjective, it is important that our choice of identifiability constraints does not have any impact on any conclusions we draw about historical or projected mortality rates. For instance, we discuss in Hunt and Blake (2015f,g) how to ensure that projected force of mortality is independent of the choice of identifiability constraint.

It is also important that the forward mortality rate framework described in this study is independent of the choice of identifiability constraints used when fitting the underlying APC model to historical data. However, due to our definitions of the forward mortality rates in Equation 11, we see that $\nu_{x,t}^{\mathbb{P}}(\tau)$ in the real-world measure is automatically independent of the identifiability constraints if the distribution of $\mu_{x,\tau}$ is also independent of the identifiability constraints. We therefore do not need to do any additional work to ensure identifiability in the forward rates once the methods used to project the force of mortality are well-identified.

We also need to ensure that the forward mortality surface in the market-consistent measure is also independent of the choice of arbitrary identifiability constraints. This is mostly straightforward, as we see that Equation 31 depends upon the forward mortality rates in the real-world measure (which should be independent of the identifiability constraints for the reasons discussed above), the variances of the period and cohort functions (which are independent of the allocation of any levels and linear trends if the projection methods are well-identified, as discussed in Hunt and Blake (2015g))

and the market prices of longevity risk. However, we note that if the model transformed using

$$\{\hat{\beta}_x, \hat{\kappa}_t\} = \{(A^{-1})^\top \beta_x, A\kappa_t\}$$

then the market prices of risk are also transformed in the model to $\hat{\lambda} = (A^{-1})^\top \lambda$. Hence we see that, not only are the values of the market prices of risk dependent upon the underlying APC model used for the force of mortality, they will also depend upon the normalisation scheme and specification of the age function in the model, and so are not the same across all models which give the same fitted mortality rates.

B Impact of Jensen's inequality

In Section 2.2, it was argued that

$$\begin{aligned} {}_tP_{x,\tau} &= \mathbb{E}_\tau \left[\exp \left(- \sum_{u=1}^t \mu_{x+u,\tau+u} \right) \right] \\ &\approx \exp \left(- \sum_{u=1}^t \mathbb{E}_\tau \mu_{x+u,\tau+u} \right) \end{aligned} \quad (35)$$

due to the relatively low degree of variability in $\mu_{x,t}$, and hence it was shown in Section 2.2 that

$$\nu_{x,t}(\tau) \approx \mathbb{E}_\tau \mu_{x,t}$$

This assumption can be tested numerically, as follows.

For simplicity, we consider $P_{x,t} = \mathbb{E}_\tau \exp(-\mu_{x,t})$. Therefore

$$P_{x,t} = \mathbb{E}_\tau \exp(-\exp(\eta_{x,t}))$$

In Section 2.3, we assume that

$$\eta_{x,t} \sim N(\mathcal{M}_{x,t}, \mathcal{V}_{x,t})$$

and therefore

$$\mathbb{E}_\tau \exp(-\mu_{x,t}) \approx \exp(-\mathbb{E}_\tau \mu_{x,t}) = \exp(-\exp(\mathcal{M}_{x,t} + 0.5\mathcal{V}_{x,t})) \quad (36)$$

Holland and Ahsanullah (1989) discussed the log-log distribution, where X is such that

$$\ln(-\ln(X)) \sim N(\mathcal{M}, \mathcal{V})$$

We therefore see that $P_{x,\tau}(\tau)$ is given by the mean of the log-log distribution if $\eta_{x,t}$ is normally distributed. However, the moments of this distribution do not have a closed form solution. Holland and Ahsanullah (1989) showed that the r^{th} raw moment of the distribution is given by

$$\mathbb{E}X^r = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-0.5x^2 - r \exp[\mathcal{M} + x\sqrt{\mathcal{V}}]\right) dx$$

which can be computed numerically.

From Section 2.3, we see

$$\begin{aligned}\mathcal{M}_{x,t} &= \alpha_x + \boldsymbol{\beta}_x^\top \mathbb{E}_\tau \boldsymbol{\kappa}_t + \mathbb{E}_\tau \gamma_{t-x} \\ \mathcal{V}_{x,t} &= \boldsymbol{\beta}_x^\top \mathbb{V}ar_\tau(\boldsymbol{\kappa}_t) \boldsymbol{\beta}_x + \mathbb{V}ar_\tau(\gamma_{t-x})\end{aligned}$$

Hence we can use the results of Holland and Ahsanullah (1989) to compute $P_{x,t}$ numerically, without recourse to the approximation in Equation 36. Using this, we calculate

$$\begin{aligned}P_{x,t} &= \mathbb{E}_\tau \exp(-\mu_{x,t}) \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-0.5z^2 - \exp[\mathcal{M}_{x,t} + x\sqrt{\mathcal{V}_{x,t}}]\right) dz\end{aligned}\quad (37)$$

numerically and compare it with the values assumed in Equation 36. This gives us a check on the accuracy of the approximation in Equation 36, which underpins the forward mortality framework.

Figure 8 shows the ratio of the numerical value of $P_{x,t}$ calculated using Equation 37 and the approximate value calculated using Equation 36 for the five mortality models considered in this paper (in the real-world measure). We can see that in the vast majority of cases, the difference that the assumption makes is less than 0.2% (i.e., ratios less than 1.002) and for no ages and years does the approximation make more than a 1.5% difference to the forward mortality rates. This is consistent with the projected mortality rates

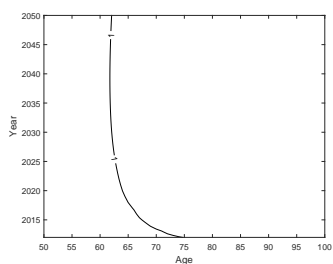
found in Figure 1, which also showed that forward mortality rates (using the approximation) were very close to those calculated using Monte Carlo simulations.

The mortality rates which are most affected by the approximation are those at the highest ages and the years of projection furthest into the future, which makes sense as these are the mortality rates with the greatest levels of uncertainty attached to them. However, they are also the least economically important, since any cashflows that would be affected by these mortality rates would be in respect of individuals who are very old (and so there is very little survivorship to these ages) and far into the future (which means that the present value of the affected cashflows would be very small due to discounting). This gives us reassurance that the approximation in Equation 35 does not systematically distort the results found using the forward mortality framework derived in this paper, compared with those which could be found using an exact but considerably more complicated framework which does not make this assumption.

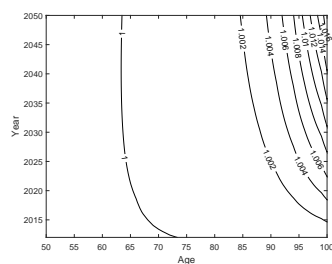
References

- Alai, D. H., Ignatieva, K., Sherris, M., 2013. Modelling longevity risk: Generalizations of the Olivier-Smith Model. Tech. rep., University of New South Wales.
- Ballotta, L., Haberman, S., 2006. The fair valuation problem of guaranteed annuity options: The stochastic mortality environment case. *Insurance: Mathematics and Economics* 38 (1), 195–214.
- Barbarin, J., 2008. Heath-Jarrow-Morton modelling of longevity bonds and the risk minimization of life insurance portfolios. *Insurance: Mathematics and Economics* 43 (1), 41–55.
- Barrieu, P. M., Bensusan, H., Karoui, N. E., Hillairet, C., Loisel, S., Ravanelli, C., 2012. Understanding, modelling and managing longevity risk: Key issues and main challenges. *Scandinavian Actuarial Journal* 3, 203–231.
- Bauer, D., Börger, M., Ruß, J., Zwiesler, H., 2008. The volatility of mortality. *Asia-Pacific Journal of Risk and Insurance* 3 (10), 2153–3792.

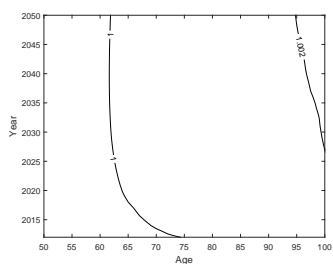
Figure 8: Impact of Jensen's inequality



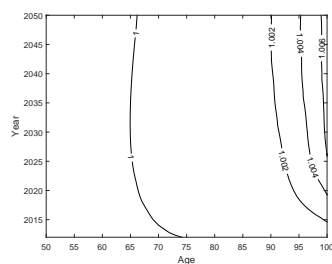
(a) Lee-Carter model



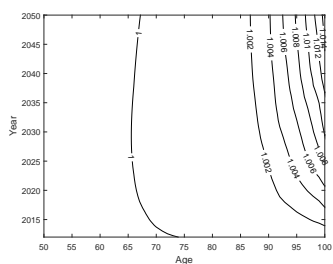
(b) CBDX model



(c) Classic APC model



(d) Reduced Plat model



(e) General procedure model

- Bauer, D., Kramer, F., 2007. Risk and valuation of mortality contingent catastrophe bonds. Tech. rep., University of Ulm.
- Bayraktar, E., Young, V. R., 2007. Hedging life insurance with pure endowments. *Insurance: Mathematics and Economics* 40 (3), 435–444.
- Blake, D., Cairns, A. J., Coughlan, G. D., Dowd, K., MacMinn, R., 2013. The new life market. *Journal of Risk and Insurance* 80 (3), 501–558.
- Blake, D., Cairns, A. J., Dowd, K., 2006. Living with mortality: Longevity bonds and other mortality-linked securities. *British Actuarial Journal* 12 (1), 153–197.
- Börger, M., Fleischer, D., Kuksin, N., 2013. Modeling the mortality trend under modern solvency regimes. *ASTIN Bulletin* 44 (1), 1–38.
- Cairns, A. J., 2000. A discussion of parameter and model uncertainty in insurance. *Insurance: Mathematics and Economics* 27 (3), 313–330.
- Cairns, A. J., 2007. A multifactor generalisation of the Olivier-Smith model for stochastic mortality. Tech. rep., Heriot-Watt University, Edinburgh.
- Cairns, A. J., 2013. Robust hedging of longevity risk. *Journal of Risk and Insurance* 80 (3), 621–648.
- Cairns, A. J., Blake, D., Dowd, K., 2006a. A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk and Insurance* 73 (4), 687–718.
- Cairns, A. J., Blake, D., Dowd, K., 2006b. Pricing death: Frameworks for the valuation and securitization of mortality risk. *ASTIN Bulletin* 36 (1), 79–120.
- Cairns, A. J., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., Khalaf-Allah, M., 2011. Mortality density forecasts: An analysis of six stochastic mortality models. *Insurance: Mathematics and Economics* 48 (3), 355–367.
- Cairns, A. J., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., Ong, A., Balevich, I., 2009. A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. *North American Actuarial Journal* 13 (1), 1–35.

- Continuous Mortality Investigation, 2009a. Working Paper 38 - A prototype mortality projection model: Part one - an outline of the proposed approach.
URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-papers>
- Continuous Mortality Investigation, 2009b. Working Paper 39 - A prototype mortality projections model: Part two - detailed analysis.
URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-papers>
- Continuous Mortality Investigation, 2013. Working Paper 69 - The CMI mortality projections model, CMI_2013, and feedback on the consultation on the future of the CMI Library of Mortality Projections and the CMI Mortality Projections Model.
URL <http://www.actuaries.org.uk/research-and-resources/pages/cmi-working-paper->
- Coughlan, G. D., Epstein, D., Sinha, A., Honig, P., 2007. q-forwards: Derivatives for transferring longevity and mortality risks. JPMorgan Pension Advisory Group.
- Cox, S. H., Lin, Y., 2007. Natural hedging of life and annuity mortality risks. North American Actuarial Journal 11 (3), 1–15.
- Cox, S. H., Lin, Y., Wang, S., 2006. Multivariate exponential tilting and pricing implications for mortality securitization. Journal of Risk and Insurance 73 (4), 719–736.
- Denuit, M. M., 2009. An index for longevity risk transfer. Journal of Computational and Applied Mathematics 230 (2), 411–417.
- Denuit, M. M., Devolder, P., Goderniaux, A.-M., 2007. Securitization of longevity risk: Pricing survivor bonds with Wang transform in the Lee-Carter framework. Journal of Risk and Insurance 74 (1), 87–113.
- Denuit, M. M., Goderniaux, A.-M., 2005. Closing and projecting life tables using log-linear models. Bulletin of the Swiss Association of Actuaries 1, 29–49.
- Dhaene, J., Kukush, A., Luciano, E., Schoutens, W., Stassen, B., 2013. On the (in-)dependence between financial and actuarial risks. Insurance: Mathematics and Economics 52 (3), 522–531.

- Dowd, K., 2003. Survivor bonds: A comment on Blake and Burrows. *Journal of Risk and Insurance* 70 (2), 339–348.
- Dowd, K., Cairns, A. J., Blake, D., 2006. Mortality-dependent financial risk measures. *Insurance: Mathematics and Economics* 38 (3), 427–440.
- EIOPA, 2014. Technical specification for the preparatory phase (Part I). Tech. rep., European Insurance and Occupational Pensions Authority, Frankfurt.
- Fetiveau, C., Jia, C., 2014. Longevity risk hedging with population based index solution - A study of basis risk based on England & Wales population. Tech. rep., Deutsche Bank.
- Frederic, P., Lad, F., 2008. Two moments of the logitnormal distribution. *Communications in Statistics - Simulation and Computation* 37 (7), 1263–1269.
- Gerber, H., Shiu, E., 1994. Option pricing by Esscher transforms. *Transactions of the Society of Actuaries* 46, 99–191.
- Haberman, S., Renshaw, A., 2009. On age-period-cohort parametric mortality rate projections. *Insurance: Mathematics and Economics* 45 (2), 255–270.
- Haberman, S., Renshaw, A., 2011. A comparative study of parametric mortality projection models. *Insurance: Mathematics and Economics* 48 (1), 35–55.
- Hobcraft, J., Menken, J., Preston, S. H., 1982. Age, period and cohort effects in demography: A review. *Population Index* 48 (1), 4–43.
- Holland, B., Ahsanullah, M., 1989. Further results on a distribution of Meinhof and Singpurwalla. *American Statistical Association* 43 (4), 216–219.
- Human Mortality Database, 2014. Human Mortality Database. Tech. rep., University of California, Berkeley and Max Planck Institute for Demographic Research.
URL www.mortality.org
- Hunt, A., Blake, D., 2014. A general procedure for constructing mortality models. *North American Actuarial Journal* 18 (1), 116–138.

- Hunt, A., Blake, D., 2015a. A Bayesian approach to modelling and projecting cohort effects. Work in Progress.
- Hunt, A., Blake, D., 2015b. Basis risk and pension schemes: A relative modelling approach. Work in Progress.
- Hunt, A., Blake, D., 2015c. Consistent mortality projections allowing for trend changes. Work in Progress.
- Hunt, A., Blake, D., 2015d. Forward mortality rates in discrete time II: Longevity risk and hedging strategies. Work in Progress.
- Hunt, A., Blake, D., 2015e. Forward mortality rates in discrete time III: Longevity-linked options. Work in Progress.
- Hunt, A., Blake, D., 2015f. Identifiability in age/period mortality models. Work in Progress.
- Hunt, A., Blake, D., 2015g. Identifiability in age/period/cohort mortality models. Work in Progress.
- Hunt, A., Blake, D., 2015h. Modelling longevity bonds: Analysing the Swiss Re Kortis Bond. Insurance: Mathematics and Economics 63, 12–29.
- Hunt, A., Blake, D., 2015i. On the structure and classification of mortality models. Tech. rep., Pensions Institute PI 15-06, Cass Business School, City University London.
- Hunt, A., Blake, D., 2015j. Solvency II and longevity risk: An application of forward mortality rates. Work in Progress.
- Hunt, A., Blake, D., 2015k. Transferring mortality and longevity risks in pension schemes via bespoke longevity swaps. Work in Progress.
- Kijima, M., 2005. A multivariate extension of equilibrium pricing transforms: The multivariate Esscher and Wang transforms for pricing financial and insurance risks. Tech. rep., Kyoto University.
- Lane, M., 2011. Longevity risk from the perspective of the ILS markets. The Geneva Papers on Risk and Insurance Issues and Practice 36 (4), 501–515.

- Lee, R. D., Carter, L. R., 1992. Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association* 87 (419), 659–671.
- Lemoine, K., 2014. Mortality regimes and longevity risk in a life annuity portfolio. *Scandinavian Actuarial Journal*, Forthcoming.
- Li, J., 2014. A quantitative comparison of simulation strategies for mortality projection. *Annals of Actuarial Science* 8 (02), 281–297.
- Li, J. S.-H., Luo, A., 2012. Key q-duration: A framework for hedging longevity risk. *ASTIN Bulletin* 42 (2), 413–452.
- Li, N., Lee, R. D., Tuljapurkar, S., 2004. Using the Lee-Carter method to forecast mortality for populations with limited data. *International Statistical Review* 72 (1), 19–36.
- Loeys, J., Panigirtzoglou, N., Ribeiro, R., 2007. Longevity: A market in the making. JPMorgan Pension Advisory Group.
- Michaelson, A., Mulholland, J., 2014. Strategy for increasing the global capacity for longevity risk transfer: Developing transactions that attract capital markets investors. *Journal of Alternative Investments* 17 (1), 18–27.
- Milevsky, M., Promislow, S., 2001. Mortality derivatives and the option to annuitise. *Insurance: Mathematics and Economics* 29, 299–318.
- Milevsky, M., Promislow, S., Young, V. R., Bayraktar, E., 2005. Financial valuation of mortality risk via the instantaneous Sharpe ratio. *Journal of Economic Dynamics and Control* 33 (3), 676–691.
- Milidonis, A., Lin, Y., Cox, S. H., 2011. Mortality regimes and pricing. *North American Actuarial Journal* 15 (2), 266–289.
- Miltersen, K. R., Persson, S.-A., 2005. Is mortality dead? Stochastic forward force of mortality rate determined by no arbitrage. Tech. rep., University of Ulm.
- Norberg, R., 2010. Forward mortality and other vital rates - Are they the way forward? *Insurance: Mathematics and Economics* 47 (2), 105–112.

- Oeppen, J., Vaupel, J., 2002. Broken limits to life expectancy. *Science* 296 (5570), 1029–1031.
- Olivier, P., Jeffrey, T., 2004. Stochastic mortality models.
URL http://www.actuaries.ie/EventsandPapers/Events2004/2004-06-01_PensionerMort
- Pelsser, A., 2003. Pricing and hedging guaranteed annuity options via static option replication. *Insurance: Mathematics and Economics* 33 (2), 283–296.
- Plat, R., 2009. On stochastic mortality modeling. *Insurance: Mathematics and Economics* 45 (3), 393–404.
- Renshaw, A., Haberman, S., 2006. A cohort-based extension to the Lee-Carter model for mortality reduction factors. *Insurance: Mathematics and Economics* 38 (3), 556–570.
- Schrager, D., 2006. Affine stochastic mortality. *Insurance: Mathematics and Economics* 38 (1), 81–97.
- Smith, A., 2005. Stochastic mortality modelling.
URL http://www.icms.org.uk/archive/meetings/2005/quantfinance/sci_prog.html
- Sweeting, P. J., 2011. A trend-change extension of the Cairns-Blake-Dowd model. *Annals of Actuarial Science* 5 (02), 143–162.
- Tappe, S., Weber, S., 2013. Stochastic mortality models: An infinite-dimensional approach. *Finance and Stochastics* 18 (1), 209–248.
- Wang, C.-W., Huang, H., Liu, I.-C., 2011. A quantitative comparison of the Lee-Carter model under different types of non-Gaussian innovations. *The Geneva Papers on Risk and Insurance Issues and Practice* 36 (4), 675–696.
- Wang, S., 2000. A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance* 67 (1), 15–36.
- Wang, S., 2002. A universal framework for pricing financial and insurance risks. *ASTIN Bulletin* 32 (2), 213–234.
- Zhou, R., Li, J. S.-H., Tan, K. S., 2015. Economic pricing of mortality-linked securities: A tatonnement approach. *Journal of Risk and Insurance* 82 (1), 65–95.

- Zhu, N., Bauer, D., 2011a. Applications of forward mortality factor models in life insurance practice. *Geneva Papers on Risk and Insurance Issues and Practice* 36, 567–594.
- Zhu, N., Bauer, D., 2011b. Coherent modeling of the risk in mortality projections: A semi parametric approach. Tech. rep., Georgia State University.
- Zhu, N., Bauer, D., 2014. A cautionary note on natural hedging of longevity risk. *North American Actuarial Journal* 18 (1), 104–115.