Pricing longevity-linked securities in the presence of mortality trend changes

Arne Freimann
Institute for Financial and Actuarial Sciences (ifa), Ulm & Institute of Insurance, Ulm University
Lise-Meitner-Straße 14, 89081 Ulm, Germany
E-mail: a.freimann@ifa-ulm.de; Phone: +49 (731) 20 644-253

Working Paper

April 7, 2020

Abstract

Pricing longevity-linked securities in the absence of a liquid and complete market for longevity risk still raises many challenges. In particular, it requires careful modeling of all relevant aspects of longevity risk in conjunction with a reasonable approach for determining adequate risk premiums. Otherwise, the risks involved in a longevity transaction might be systematically understated which might in turn lead to insufficient risk premiums and to inadequate prices. In fact, historical mortality patterns often reveal structural break points at which the current rate of mortality improvement, the so-called mortality trend, has permanently changed.

In this paper, we present a stochastic modeling framework for the valuation of longevity-linked securities which explicitly considers the risk of random future changes in the long-term mortality trend. We construct a set of meaningful probability distortions which imply equivalent risk-adjusted pricing measures under which the basic model structure is preserved. Inspired by capital requirements for (re)insurers under risk-based solvency regimes, we also establish a cost of capital pricing approach, which then serves as the appropriate reference framework for finding a reasonable range for the market price of longevity risk. Our multi-population setting is applicable to customized as well as to index-based instruments. In a numerical application, we demonstrate that the approaches produce plausible risk loadings and find that in the presence of mortality trend changes most of the risk loading is allocated to longer maturities.

JEL classification: G13; G22
Keywords: Longevity risk; Longevity derivatives; Stochastic mortality; Mortality trend processes

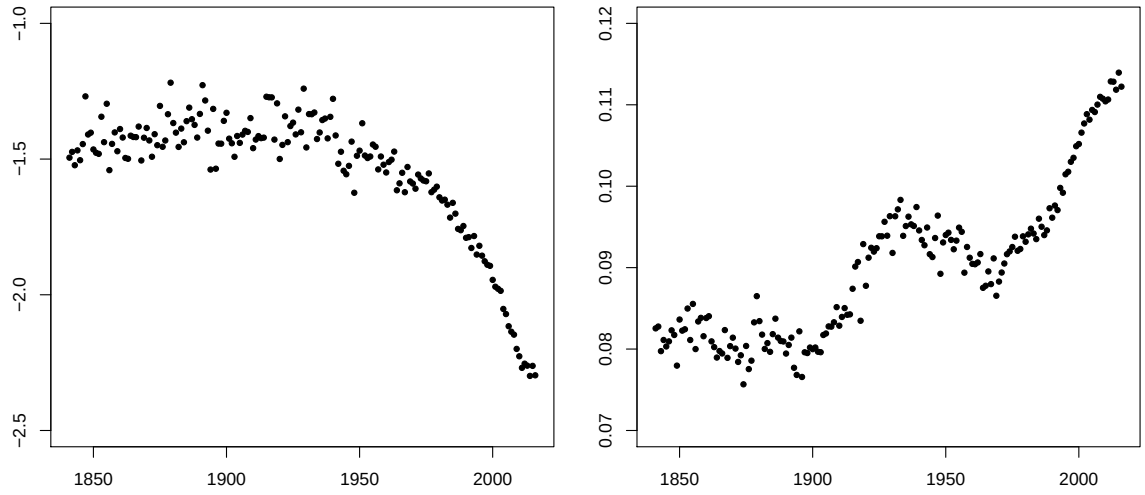

(a) $\kappa_t^{(1)}$ (b) $\kappa_t^{(2)}$

Figure 1: Historical time processes for English and Welsh males in the CBD model

1 Introduction

The risk that future mortality trends imply stronger mortality improvements than anticipated is commonly referred to as longevity risk and constitutes a major risk for life insurers and pension funds. For its modeling, measurement, and management, academics have proposed a variety of stochastic mortality models. The majority of these models, including the widely used Lee-Carter model (cf. Lee and Carter (1992)) and Cairns-Blake-Dowd (CBD) model (cf. Cairns et al. (2006)), fall into the class of parametric mortality models, which decompose mortality across the dimensions age, period, and sometimes cohort. A general procedure for constructing parametric mortality models can for instance be found in Hunt and Blake (2014). For stochastic mortality projections, often a random walk with drift (RWD) is used to project the time-dependent parameters (also called period effects) into the future.

However, several authors have questioned the adequacy of the RWD for long-term mortality projections since it solely captures diffusion risk around the most recent mortality trend without accounting for the risk of potential future trend changes, see e.g. Li et al. (2011). In fact, historical mortality patterns often reveal structural breaks which are highly unlikely to be generated from a RWD. Figure 1 shows, exemplarily, the historical time processes $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ in the CBD model for English and Welsh males (the model calibration is provided in the appendix). The first time process clearly shows a significant improvement in mortality over the last century. However, the annual rate of improvement does not appear to be constant over time. In fact, the historical mortality trends for English and Welsh males appear to be trend stationary around piecewise linear trends. In particular, several break points can be identified for both time processes at which the prevailing mortality trend has changed. Motivated by similar empirical findings for many countries, several authors have

proposed stochastic processes for mortality models that explicitly model changes in mortality trends. For an overview and comparison of different approaches, we refer to Börger and Schupp (2018). Even though allowing for random future trend changes naturally comes at the cost of higher model complexity, it seems indispensable for a proper modeling and assessment of longevity risk, especially over long-term horizons.

In addition to the question of how to model longevity risk, its securitization and the recent development of the global longevity risk transfer market have attracted huge attention among both academics and practitioners. For this emerging market, a great variety of longevity-linked securities, also called longevity hedges or longevity derivatives, have been proposed in the literature, see Blake et al. (2019) for an overview. However, the longevity risk transfer market is still in an early state and illiquid and incomplete. Hence, longevity-linked instruments cannot be priced based on observable market prices, i.e. mark-to-market, and the question arises how to put a 'reasonable' value on them. Since stochastic mortality models are normally calibrated to historical mortality data, they typically provide values on a best estimate basis rather than market prices. In particular, these actuarial models do not include a risk loading beyond best estimate values and further steps in the form of risk adjustments are required to arrive at reasonable and consistent (mark-to-model) prices, which are required under modern risk-based solvency guidelines, see e.g. Wüthrich (2016). To this end, several pricing techniques have been proposed in the literature, and we refer to Bauer et al. (2010) and Leung et al. (2018) for an overview. In this work, we focus on two promising approaches:

- Inspired by risk-neutral valuation of financial derivatives, the *risk-adjusted pricing approach*, often also referred to as *risk-neutral approach*, is widely used, see among others Boyer and Stentoft (2013), Cairns et al. (2006), Chen and Cox (2009), and Leung et al. (2018). The basic idea is to derive an equivalent risk-adjusted measure under which prices are defined as expected values of discounted future cash flows. Compared to the objective measure, this pricing measure assigns higher probability mass to scenarios of stronger mortality improvements reflecting the risk adjustment made by longevity risk takers.

- Börger (2010) suggests to make inferences on prices for longevity-linked securities from capital requirements for (re)insurers under modern risk-based solvency regimes, such as Solvency II or the Swiss Solvency Test. (Re)insurance companies are required to provide adequate regulatory capital, so-called Solvency Capital Requirements (SCRs), for all risks they are exposed to including longevity. The so-called *cost of capital approach* postulates that a market player will only be willing to take longevity risk if the expected return exceeds the expected additional capital charges for taking the risk. Otherwise, the deal is not considered to be economically viable. Since the market at present is still dominated by global reinsurance companies, this approach seems to be of high practical relevance. Levantesi and Menzietti (2017) implement it in the context of a regulatory capital model under Solvency II. However, their framework is not entirely consistent with the one-year view of Solvency II as they determine the relevant 99.5% Value-at-Risk (VaR) over a multi-year horizon. By contrast, Zeddouk and Devolder (2019) establish a setup that is more in line with the one-year view of Solvency II and compare the cost of capital approach with alternative pricing approaches.

Naturally, mark-to-model valuation of longevity-linked securities relies crucially on the suitability of the underlying stochastic mortality model for the given purpose and on the

adequacy of its mortality forecasts. In the field of mortality securitization, Chen and Cox (2009) and Zhou et al. (2015) argue that a model for pricing hedges against catastrophic mortality risk should account for the risk of outliers, such as wars and pandemics. Ignoring important risk drivers might significantly understate the actual amount of risk taken and might consequently lead to insufficient risk premiums and to inadequate prices. The same argument also applies to the risk of permanent mortality trend changes which constitutes a major risk for longevity risk takers. Hence, it should be properly modeled when pricing long-term longevity-linked securities.

In the current paper, we address the valuation of longevity-linked securities in a stochastic modeling framework which explicitly considers the risk of random future mortality trend changes. At the core of our framework is a simulation model for the underlying long-term mortality trend of a suitable reference population that builds on previous works of Börger and Schupp (2018) and Schupp (2019). Furthermore, we allow mortality in selected subpopulations to systematically differ from the underlying long-term trend of the refernce population. The model is therefore applicable to customized as well as to more standardized instruments that are linked to the mortality indices of the reference population. Within this framework, we establish the above-introduced approaches for determining adequate longevity risk premiums: First, we apply the technique of multivariate normalized exponential tilting, which is based on the widely used Wang distortion operator, to construct a set of stochastic discount factors, also called deflators, which imply equivalent risk-adjusted pricing measures under which the basic model structure is preserved. In a second step, we rely on the cost of capital approach to find a reasonable range for the market price of longevity risk in the current incomplete market environment. For deriving longevity SCRs consistently with regulatory requirements under Solvency II, we follow Börger et al. (2019b) and clearly distinguish between two mortality trends: the actual (but unobservable) mortality trend (AMT) prevailing at a certain point in time and the estimated mortality trend (EMT) an observer would estimate given the observed mortality patterns in previous years. This allows to model the impact of new mortality information on best estimate mortality assumptions over a one-year horizon which essentially determines the SCRs. We discuss and compare both pricing approaches with regard to their economic justification, practical applicability, and calibration. In a numerical application, we calibrate the presented approaches and apply them to price longevity swaps as examples for widely used longevity-linked instruments.

Our work contributes to the literature on longevity risk pricing in the following ways: First and foremost, we explicitly include the major risk of random future mortality trend changes in the mark-to-model valuation of longevity hedges. In light of the potential threat posed by an unprecedented change in the long-term mortality trend to the solvency of longevity risk bearers, our modeling framework constitutes a valuable alternative to existing approaches. Second, we contribute to the ongoing debate of how to find reasonable market prices of longevity risk in an incomplete market by relying on regulatory capital charges for longevity risk as an appropriate reference framework. In particular, we derive (future) capital charges for longevity risk consistently with regulatory capital requirements under Solvency II and properly account for its stochastic nature. As the longevity risk transfer market is currently dominated by reinsurance companies, this approach has the advantage of giving a reasonable range for the market price of longevity risk without the need for market price data, which are currently not publicly available.

The remainder of this paper is organized as follows. Section 2 introduces the underlying stochastic mortality trend model. In particular, we describe how to simulate the AMT and

how to derive the EMT, i.e., current best estimate mortality assumptions, based on most recent observed mortality. A calibration of the model to the historical mortality experience of English and Welsh males is provided in Appendix A. In Section 3, we construct a set of equivalent pricing measures and establish the cost of capital approach in the context of a regulatory capital model under Solvency II. In a numerical application in Section 4, we calibrate both pricing approaches and derive risk-adjusted forward rates for customized and index-based longevity swaps. Finally, Section 5 concludes.

2 Stochastic mortality trend model

As outlined in the introduction, we now describe a stochastic modeling framework which explicitly accounts for random future changes in the long-term mortality trend as well as for portfolio-specific mortality characteristics. After introducing the underlying multi-population model structure in Section 2.1, we present the AMT model component for generating paths of future mortality in Section 2.2. Finally, Section 2.3 presents the EMT component for pathwise derivations of best estimate mortality assumptions at future points in time. The model as a whole can be seen as a multi-population extension of the combined AMT/EMT setup of Börger et al. (2019b).

2.1 Multi-population model structure

Our modeling framework builds on a multi-population extension of the CBD model structure as originally proposed by Cairns et al. (2006). We assume a *reference population* for which the logit of the one-year death probabilities is parametrized as

$$\text{logit}\left(q_{x,t}^{[R]}\right) := \log\left(\frac{q_{x,t}^{[R]}}{1-q_{x,t}^{[R]}}\right) = \kappa_t^{(1)[R]} + (x-\bar{x})\,\kappa_t^{(2)[R]},$$

where $\bar{x}$ denotes the average age of the calibration age range. The first period effect $\kappa_t^{(1)[R]}$ captures the general level of mortality over time while the second time-dependent parameter $\kappa_t^{(2)[R]}$ describes the slope of the mortality line in the logit plot.

Relative to the reference population, we model mortality differentials for $N_{Sub} \geq 1$ sub-populations, e.g. distinct subgroups of different socioeconomic status, also in logit-bilinear form, i.e.,

$$\text{logit}\left(q_{x,t}^{[p]}\right) - \text{logit}\left(q_{x,t}^{[R]}\right) = \kappa_t^{(1)[p]} + (x-\bar{x})\,\kappa_t^{(2)[p]}, \qquad p \in \{1,\ldots,N_{Sub}\},$$

where $\bar{x}$ is chosen analogously to the reference population. In this setting, the time-dependent parameters for the reference population represent the common long-term mortality trend of all (sub-)populations while the subpopulation-specific time processes account for mortality differentials over time. Relative modeling approaches of this kind are widely used in the field of multi-population mortality modeling, see e.g. Villegas et al. (2017).

We would like to point out two important features of this multi-population CBD model structure. First, it possesses the convenient 'new-data-invariant property' (see Tan et al. (2014)), which will be an essential feature for the EMT component of our model. Second, it allows for a natural logit-extrapolation of mortality for high ages as outlined in Appendix A.

2.2 AMT model component

For stochastic mortality projections, we clearly distinguish between the following components of longevity risk: the long-term mortality trend risk of the overall population, systematic subpopulation-specific mortality differentials, and unsystematic small sample risk in portfolios of limited size. This rigorous differentiation between the three risk drivers will later allow to assign a risk premium to each longevity risk component, which in turn allows to price index-based instruments that solely cover the risk originating from the reference population.

2.2.1 Long-term mortality trend risk

For the projection of the time-dependent parameters of the reference population, we rely on the stochastic trend process which was originally proposed by Börger and Schupp (2018) and further refined by Schupp (2019). They interpret the observable time processes as random noise around an underlying unobservable piecewise linear mortality process, i.e.,

$$\kappa_t^{(i)[R]} = \hat{\kappa}_t^{(i)[R]} + \epsilon_t^{(i)[R]}, \qquad i = 1, 2,$$

where the noise vector $\epsilon_t^{[R]} := (\epsilon_t^{(1)[R]}, \epsilon_t^{(2)[R]})'$ accounts for yearly mortality fluctuations and is modeled as a two-dimensional normal distribution with mean zero and constant covariance matrix $\Sigma^{[R]}$. The actual underlying mortality processes are then projected linearly as

$$\hat{\kappa}_t^{(i)[R]} = \hat{\kappa}_{t-1}^{(i)[R]} + \hat{d}_t^{(i)[R]}, \qquad i = 1, 2,$$

where $\hat{d}_t^{(i)[R]}, i = 1, 2$ is interpreted as the unobservable AMT prevailing at time t. We follow Börger and Schupp (2018) and update the AMTs independently of each other as follows:

$$\hat{d}_t^{(i)[R]} = \hat{d}_{t-1}^{(i)[R]} + O_t^{(i)[R]} S_t^{(i)[R]} M_t^{(i)[R]}, \qquad i = 1, 2,$$

where

- The Bernoulli-distributed random variable $O_t^{(i)[R]} \in \{0, 1\}$ indicates whether a trend change occurs between $t-1$ and t or not. The trend change probability is denoted as $p^{(i)[R]} > 0$.
- Given a trend change occurs, $S_t^{(i)[R]} \in \{-1, 1\}$ determines its sign, where we assume positive or negative trend changes to be equally likely under the objective measure.
- Finally, $M_t^{(i)[R]} > 0$ denotes the absolute trend change magnitude, which is modeled by a lognormal distribution with parameters $\mu_M^{(i)[R]}$ and $\sigma_M^{(i)[R]^2}$.

As argued by Börger and Schupp (2018), the above decomposition of the trend change intensities and the choice of lognormally distributed trend change magnitudes offer some desirable properties. First, the symmetric distribution for the trend change intensities assures that the prevailing AMT (even though unobservable) always represents the best estimate trend for any future point in time. Second, the heavy-tailed lognormal distribution with only very little mass around zero produces rather significant trend changes and especially allows for strong trend changes, which can be observed for some countries in the past.

For practical applications, the model parameters need to be estimated from historical mortality data. As argued by Börger et al. (2019a), this typically involves a considerable

amount of uncertainty. We explicitly account for parameter uncertainty in the starting values of the simulation as well as in the trend change parameters by sampling them at the start of each simulation from suitable distributions, we refer to the appendix for details.

2.2.2 Subpopulation-specific mortality differentials

Next, also the vectors of subpopulation-specific period effects $\kappa_t^{[p]} := (\kappa_t^{(1)[p]}, \kappa_t^{(2)[p]})', p \in \{1, \dots, N_{Sub}\}$ need to be projected. We follow Villegas et al. (2017) and model them independently of the stochastic trend processes for the reference population. The latter serve as the underlying long-term mortality trend for all subpopulations, which, by construction, implies a reasonable correlation between the subpopulations.

To this end, often stationary processes are used to guarantee 'coherent', i.e., non-diverging, mortality rates between closely related populations in the long run. However, enforcing mean-reversion actually constitutes a strong modeling assumption, see e.g. Li et al. (2017). For this reason, we follow Villegas and Haberman (2014) and consider instead a multivariate RWD of the form

$$\kappa_{t+1}^{[p]} = \mu^{[p]} + \kappa_t^{[p]} + \epsilon_t^{[p]}, \qquad p \in \{1, \dots, N_{Sub}\},$$

where $\mu^{[p]} \in \mathbb{R}^2$ denotes the drift and $\epsilon_t^{[p]} \in \mathbb{R}^2$ denotes the annual random innovations for subpopulation $p \in \{1, \dots, N_{Sub}\}$. The joint vector $\epsilon_t^{[Sub]} := (\epsilon_t^{[1]}, \dots, \epsilon_t^{[N_{Sub}]})' \in \mathbb{R}^{2N_{Sub}}$ of annual innovations follows a multivariate normal distribution with mean zero and constant covariance matrix $\Sigma^{[Sub]} \in \mathbb{R}^{(2N_{Sub})\times(2N_{Sub})}$.

We would like to stress that the individual model components can be specified independently of each other. Hence, the RWD can easily be replaced by a mean-reverting process (for instance a first-order vector autoregressive (VAR) process) if desired.

2.2.3 Unsystematic small sample risk

Finally, idiosyncratic risk in a portfolio of limited size is accounted for by sampling survivors over time from a Binomial distribution. In each model path, we draw realizations for survivors aged $x+1$ at time $t+1$ coming from subpopulation p from the following Binomial distribution:

$$B_{x+1,t+1}^{[p]} \sim \text{Binom}\left(B_{x,t}^{[p]}, 1 - q_{x,t+1}^{[p]}\right), \qquad p \in \{1, \dots, N_{Sub}\},$$

given $B_{x,t}^{[p]}$ survivors aged x still alive at time t.

2.3 EMT model component

For deriving longevity SCRs consistently with regulatory requirements under Solvency II, the framework needs to model the impact of new mortality information on best estimate mortality assumptions at future points in time (see Section 3.2 for details). As argued by Börger et al. (2019b), an observer at some future point in time T is generally not able to clearly distinguish between a recent change in the underlying long-term mortality trend and a 'normal' random fluctuation around it. For a best estimate mortality projection starting at time T, an actuary would calibrate a mortality model to the most recent available mortality data. For simplicity, we assume recalibration of the same CBD model structure to sufficiently

large populations under consideration. Its 'data-invariant-property' ensures that this recalibration procedure results exactly in the same time-dependent parameters as simulated by the underlying simulation model up to time T. Hence, instead of refitting the whole model within each simulation path, the updated best estimate mortality assumptions can directly be derived based on the realized time processes, which makes the modeling framework highly efficient in practical applications. Note that for other models, for instance for the model of Lee and Carter (1992) or for models with cohort effects, a full model recalibration would be required which comes at a high computational cost in a Monte Carlo simulation.

In our multi-population setting, this pathwise derivation of best estimate mortality consists of two steps: estimating the AMT for the overall population followed by an adjustment for differing mortality levels and trends for the subpopulations.

2.3.1 Reference population

For a best estimate mortality projection for the reference population beyond time T, we follow a deterministic central path of mortality under the CBD model, i.e.,

$$\operatorname{logit}\left(\tilde{q}_{x,t}^{[R]}(T)\right) = \tilde{\kappa}_T^{(1)[R]} + (t-T)\,\tilde{d}_T^{(1)[R]} + (x-\bar{x})\left(\tilde{\kappa}_T^{(2)[R]} + (t-T)\,\tilde{d}_T^{(2)[R]}\right), \qquad t > T,$$

where $\tilde{\kappa}_T^{(i)[R]}, i = 1,2$ denotes the prevailing mortality level and $\tilde{d}_T^{(i)[R]}, i = 1,2$ denotes the current EMT. These quantities need to be estimated based on the observed mortality patterns in previous years. To this end, we follow Börger et al. (2019b) by applying a weighted linear regression on the most recent data points based on the following exponentially decaying weights:

$$w^{(i)}(t,T) := \frac{1}{\left(1+1/\psi^{(i)}\right)^{T-t}}, \qquad i = 1,2, \qquad t \leq T.$$

The weighting parameters $\psi^{(i)} \geq 0, i = 1,2$ need to be specified. As argued by Börger et al. (2019b), it is important to find a reasonable trade-off between the following objectives:

- Avoiding an overreaction of the EMT to 'normal' random fluctuations around the AMT.
- Ensuring a prompt detection of a recent mortality trend change.

While the first objective can be achieved by including enough data points, the second objective requires sufficient weight on the most recent data points, which naturally contain the most relevant information. We follow the recommendation of Börger et al. (2019b) and numerically determine the optimal weighting in a Monte Carlo simulation by minimizing the mean squared errors between the unobservable AMT and the derived EMT, see Appendix A. We would like to stress that exponential weighting guarantees that the updated EMT will remain unchanged if the new data point realizes exactly as expected under the previous EMT.

2.3.2 Subpopulations

Given best estimate mortality for the overall population, subpopulation-specific mortality rates are projected beyond time T by means of the following adjusted central path of mortality:

$$\tilde{q}_{x,t}^{[p]}(T) := \operatorname{logit}^{-1}\left(\operatorname{logit}\left(\tilde{q}_{x,t}^{[R]}(T)\right) + \tilde{\kappa}_T^{(1)[p]} + (t-T)\,\tilde{\nu}_T^{(1)[p]} + (x-\bar{x})\left(\tilde{\kappa}_T^{(2)[p]} + (t-T)\,\tilde{\nu}_T^{(2)[p]}\right)\right)$$

for each subpopulation $p \in \{1,\ldots,N_{Sub}\}$, where the subpopulation-specific adjustment terms

- $\tilde{\kappa}_T^{(i)[p]}, i = 1, 2$ account for differing mortality levels and
- $\tilde{\nu}_T^{(i)[p]}, i = 1, 2$ capture differing mortality trends relative to the reference population.

These parameters also have to be estimated based on the observed mortality patterns in previous years. For the sake of consistency, we derive these quantities by applying the same weighted linear regression (including the same weights) as for the reference population. Due to the properties of the weighted linear regression, this ensures 'order-invariance' in the sense that the above two-step estimation approach produces the same best estimate mortality projection as directly applying a weighted linear regression to the overall subpopulation-specific period effects $\kappa_t^{(i)[R]} + \kappa_t^{(i)[p]}, i = 1, 2$.

3 Pricing approaches

We consider a market participant, referred to as the *hedge provider* or *risk taker*, who is willing to take on longevity risk by issuing longevity-linked securities at time zero. To focus on longevity risk, we do not deal with operational risk or counterparty credit risk for all longevity-linked transactions and assume independence between mortality and interest rates. The time zero random present value of all future instrument cash flows to the counterparty, also referred to as the *hedger*, is denoted as

$$H_0 := \sum_{t>0}^{\tau} B_t^{-1} h_t,$$

where τ denotes the contract maturity, $(B_t)_{t\geq 0}$ denotes the value process of a risk-free bank account with $B_0 := 1$, and $h_t, 0 < t \leq \tau$ denotes the cash flow at time t, where a positive cash flow represents a payment from the risk taker to the hedger.

In exchange, the risk taker demands the instrument's objective best estimate value plus a risk premium, which needs to be specified. Depending on the instrument's payout structure, either a single contract premium, denoted as P_0, is charged at inception or the risk premium is directly included in the instrument's payoff structure by means of risk-adjusted forward rates so that no payment exchanges hands at inception. In either case, the risk premium demanded on top of the objective best estimate value needs to be determined.

We now establish two pricing methods for deriving reasonable risk premiums in the previously introduced modeling framework: the risk-adjusted pricing approach based on a pricing measure in Section 3.1 and the cost of capital approach in Section 3.2. In Section 3.3 we briefly discuss and qualitatively compare both approaches. By way of illustration, we will price longevity swaps in Section 4.1 assuming constant deterministic interest rates, i.e., $B_t := (1 + r)^t, t \geq 1$ for some constant interest rate r, although we avoid any loss of generality in the approaches presented in this section.

3.1 Risk-adjusted measure

A widely used approach for pricing longevity-linked instruments is to derive a risk-adjusted measure $\mathbb{Q}$, which is equivalent to the objective (or real-world) measure $\mathbb{P}$, under which prices are defined as expected values of discounted future cash flows, i.e.,

$$P_0 := \mathbb{E}^{\mathbb{Q}}(H_0).$$

For forward-type instruments, the forward rates are typically determined at time zero so that no payment exchanges hands at inception, i.e., $\mathbb{E}^{\mathbb{Q}}(H_0) = 0$. Unlike many financial assets, the underlying longevity risk drivers are not continuously tradeable. Hence, longevity-linked payments cannot be replicated and a unique risk-adjusted measure cannot be deduced based on the principle of no-arbitrage. Instead, we rely on the concept of *deflators*, also called *state price densities*, to construct an equivalent risk-adjusted measure. In our incomplete market setting in discrete time, prices can be represented in the following equivalent form:

$$P_0 = \mathbb{E}^{\mathbb{Q}}\left(\sum_{s>0}^{\tau} B_s^{-1} h_s\right) = \mathbb{E}^{\mathbb{P}}\left(\sum_{s>0}^{\tau}\prod_{t=1}^{s}\varphi_t h_s\right),$$

where the random variables $\varphi_t \gg 0, t \geq 1$, commonly referred to as *span-deflators*, play the role of stochastic discount factors that assign a value at time $t-1$ to random cash flows that occur at time t. Hence, there is a one-to-one correspondence between risk-adjusted valuation and stochastic discounting and the task of constructing an appropriate pricing measure is equivalent to specifying reasonable span-deflators. For the theoretical foundations of market-consistent actuarial valuation, we refer to Section 2 in Wüthrich (2016).

In our stochastic mortality trend model, longevity-linked payments depend on multiple risk drivers simultaneously, namely the most prominent risks of unpredictable occurrence (O), sign (S), and magnitude (M) of future trend changes, as well as annual noise (ϵ) around the AMT, and subpopulation-specific mortality differentials (Sub). Given stochastically independent risk drivers, we define the following multiplicative span-deflators:

$$\varphi_t := \varphi_t^r \varphi_t^O \varphi_t^S \varphi_t^M \varphi_t^\epsilon \varphi_t^{Sub}, \qquad t \geq 1,$$

where $\varphi_t^r, t \geq 1$ plays the role of interest rate discounting, e.g. $\varphi_t^r := (1+r)^{-1}$ in a constant interest rate environment, and $\varphi_t^{Risk}, t \geq 1$ denotes the span-deflators for the longevity risk driver $Risk \in \{O, S, M, \epsilon, Sub\}$, which need to be carefully specified in order to obtain a reasonable model, cf. Section 2.6.2 in Wüthrich (2016). Loosely speaking, meaningful span-deflators for longevity risk pricing should take high values in states that are less favorable for longevity risk takers and low values in states of higher mortality. Therefore, they imply an equivalent pricing measure that assigns, compared to the objective measure, higher probability mass to scenarios of stronger mortality improvements, which typically trigger higher (or more) payouts of longevity-linked instruments. In line with the meaning of the period effects in the CBD model, the risk adjustment should aim at shifting the distributions of both time processes downwards, particularly also the second period effect since longevity at high ages generally presents the greatest risk. Thereby, a risk loading beyond best estimate values is induced for pricing purposes. To avoid ambiguity, the span-deflators should satisfy $\mathbb{E}\left(\varphi_t^{Risk}\right) = 1$ for all $t \geq 1$ and each longevity risk driver $Risk \in \{O, S, M, \epsilon, Sub\}$.

To construct a class of meaningful span-deflators (or equivalent pricing measures respectively), we follow Chen and Cox (2009) and Boyer and Stentoft (2013) and apply the technique of multivariate normalized exponential tilting, which offers several appealing features (cf. Wang (2007)): First, given independent risks, it is equivalent to applying the widely used (multivariate) Wang transform to each longevity risk driver individually, which in turn ensures consistency between the risk drivers, we refer to Wang (2002, 2007) for details. At this point, we would like to stress that our distortions do *not* act directly on the distribution of the instrument one seeks to price, which would be in the original spirit of Wang (2002), but on the distributions of the underlying longevity risk drivers. As pointed out by Chen and Cox

(2009), this has the advantage of giving a risk-adjusted measure that properly accounts for correlations of mortality over time and is applicable to a wide range of instruments. Second, the (multivariate) Wang transform preserves the Bernoulli, normal, and lognormal distribution. Hence, the basic model structure is preserved under $\mathbb{Q}$ which allows to directly simulate the risk-adjusted distribution, which is of particular importance for practical applications. In detail, we apply the (multivariate) Wang transform based on the distortion parameters $\lambda_{Risk}^{(i)} > 0, i = 1, 2$ to each risk driver $Risk \in \{O, S, M, \epsilon, Sub\}$ individually and obtain the following span-deflators and implied risk-adjusted dynamics for each year $t \geq 1$:

- **Trend change occurrence**
  We define $\varphi_t^O := \prod_{i=1}^2 \varphi_t^{(i)O}$ with span-deflators for time process $i = 1, 2$ of

$$\varphi_t^{(i)O} := \begin{cases} \frac{\Phi\left(\Phi^{-1}\left(p^{(i)[R]}\right)+\lambda_O^{(i)}\right)}{p^{(i)[R]}} \geq 1, & \text{if } O_t^{(i)[R]} = 1 \text{ (trend change occured)} \\ \frac{1-\Phi\left(\Phi^{-1}\left(p^{(i)[R]}\right)+\lambda_O^{(i)}\right)}{1-p^{(i)[R]}} \leq 1, & \text{if } O_t^{(i)[R]} = 0 \text{ (no trend change occured),} \end{cases}$$

  where Φ denotes the standard normal cumulative distribution function (CDF). Under the risk-adjusted measure, the random variables $O_t^{(i)[R]}, i = 1, 2$ indicating the occurrence of a trend change between two consecutive years are again Bernoulli-distributed with risk-adjusted trend change probabilities $\Phi\left(\Phi^{-1}\left(p^{(i)[R]}\right) + \lambda_O^{(i)}\right) \geq p^{(i)[R]}, i = 1, 2$ (see Definition 2.3 in Wang (2007)). Hence, the risk-adjusted measure assigns a higher probability to scenarios with more trend changes.

- **Trend change sign**
  Regarding the signs of future trend changes, we define $\varphi_t^S := \prod_{i=1}^2 \varphi_t^{(i)S}$, where the span-deflators for time process $i = 1, 2$ are given by

$$\varphi_t^{(i)S} := \begin{cases} \frac{\Phi\left(\Phi^{-1}(0.5)+\lambda_S^{(i)}\right)}{0.5} \geq 1, & \text{if } O_t^{(i)[R]} = 1, S_t^{(i)[R]} = -1 \text{ (negative trend change)} \\ \frac{1-\Phi\left(\Phi^{-1}(0.5)+\lambda_S^{(i)}\right)}{1-0.5} \leq 1, & \text{if } O_t^{(i)[R]} = 1, S_t^{(i)[R]} = +1 \text{ (positive trend change)} \\ 1, & \text{else (no trend change occured).} \end{cases}$$

  Hence, the deflators take low values in scenarios of positive trend changes in favor of paths with negative trend changes. Under $\mathbb{Q}$, negative trend changes then have a probability of $\mathbb{Q}\left(S_t^{(i)[R]} = -1\right) = \Phi\left(\Phi^{-1}(0.5) + \lambda_S^{(i)}\right) \geq 0.5, i = 1, 2.$

- **Trend change magnitude**
  For the stochastic trend change magnitudes $M_t^{(i)[R]}$, we choose $\varphi_t^M := \prod_{i=1}^2 \varphi_t^{(i)M}$ based on the following span-deflators for time process $i = 1, 2$:

$$\varphi_t^{(i)M} := \begin{cases} \exp\left(\lambda_M^{(i)} \frac{\log\left(M_t^{(i)[R]}\right)-\mu_M^{(i)[R]}}{\sigma_M^{(i)[R]}} - \frac{1}{2}{\lambda_M^{(i)}}^2\right), & \text{if } O_t^{(i)[R]} = 1 \text{ (trend change occured)} \\ 1, & \text{else (no trend change occured).} \end{cases}$$

  Under $\mathbb{Q}$, the trend change magnitudes $M_t^{(i)[R]}, i = 1, 2$ are again lognormally distributed with risk-adjusted parameters $\mu_M^{(i)[R]} + \lambda_M^{(i)}\sigma_M^{(i)[R]} \geq \mu_M^{(i)[R]}, i = 1, 2$ and same volatility as under $\mathbb{P}$ (see Definition 2.3 in Wang (2007)).

Risk driver	Distribution under $\mathbb{P}$	Distribution under $\mathbb{Q}$
Trend change occurrence	Bernoulli $\left(p^{(i)[R]}\right)$	Bernoulli $\left(\Phi\left(\Phi^{-1}\left(p^{(i)[R]}\right)+\lambda_O^{(i)}\right)\right)$
Trend change sign	$\mathbb{P}\left(S_t^{(i)[R]}=-1\right)=0.5$	$\mathbb{Q}\left(S_t^{(i)[R]}=-1\right)=\Phi\left(\Phi^{-1}\left(0.5\right)+\lambda_S^{(i)}\right)$
Trend change magnitude	$\mathcal{LN}\left(\mu_M^{(i)[R]},{\sigma_M^{(i)[R]}}^2\right)$	$\mathcal{LN}\left(\mu_M^{(i)[R]}+\lambda_M^{(i)}\sigma_M^{(i)[R]},{\sigma_M^{(i)[R]}}^2\right)$
Noise around AMT	$\mathcal{N}\left(0,\Sigma^{[R]}\right)$	$\mathcal{N}\left(-D^{\frac{1}{2}}_{\Sigma^{[R]}}\Sigma_\rho^{[R]}\Lambda_\epsilon,\Sigma^{[R]}\right)$
Mortality differentials	$\mathcal{N}\left(0,\Sigma^{[Sub]}\right)$	$\mathcal{N}\left(-D^{\frac{1}{2}}_{\Sigma^{[Sub]}}\Sigma_\rho^{[Sub]}\Lambda_{Sub},\Sigma^{[Sub]}\right)$

Table 1: Objective and risk-adjusted model dynamics.

- **Annual noise around the AMT**
  Considering the normally-distributed annual noise $\epsilon_t^{[R]}$ around the AMT as risk driver, we obtain span-deflators of the form

$$\varphi_t^\epsilon := \prod_{i=1}^{2}\exp\left(-\lambda_\epsilon^{(i)}\frac{\epsilon_t^{(i)[R]}-0}{\sqrt{\Sigma_{i,i}^{[R]}}}-\frac{1}{2}{\lambda_\epsilon^{(i)}}^2\right)\frac{1}{c_\epsilon}=\exp\left(-\Lambda_\epsilon^{'}D^{-\frac{1}{2}}_{\Sigma^{[R]}}\epsilon_t^{[R]}-\frac{1}{2}\Lambda_\epsilon^{'}\Lambda_\epsilon\right)\frac{1}{c_\epsilon},$$

  where $D_{\Sigma^{[R]}}$ denotes a 2-by-2 matrix with entries $\Sigma_{i,i}^{[R]}, i=1,2$ on the diagonal and $\Lambda_\epsilon := (\lambda_\epsilon^{(1)},\lambda_\epsilon^{(2)})'$ is the relevant market price of risk vector. The normalizing coefficient $c_\epsilon := \exp\left(-\frac{1}{2}\Lambda_\epsilon^{'}\Lambda_\epsilon+\frac{1}{2}\Lambda_\epsilon^{'}\Sigma_\rho^{[R]}\Lambda_\epsilon\right)$ with $\Sigma_\rho^{[R]}$ being the correlation matrix of $\epsilon_t^{[R]}$ ensures that $\mathbb{E}\left(\varphi_t^\epsilon\right)=1$ by the properties of the multivariate lognormal distribution. Under $\mathbb{Q}$, annual fluctuations around the AMT are again normally distributed with risk-adjusted mean $-D^{\frac{1}{2}}_{\Sigma^{[R]}}\Sigma_\rho^{[R]}\Lambda_\epsilon$ and same covariance matrix $\Sigma^{[R]}$ (see Theorem 5.1 in Wang (2007)).

- **Subpopulation-specific mortality differentials**
  Analogously, we obtain for the Gaussian subpopulation-specific mortality differentials

$$\varphi_t^{Sub} := \exp\left(-\Lambda_{Sub}^{'}D^{-\frac{1}{2}}_{\Sigma^{[Sub]}}\epsilon_t^{[Sub]}-\frac{1}{2}\Lambda_{Sub}^{'}\Lambda_{Sub}\right)\frac{1}{c_{Sub}},$$

  where $D_{\Sigma^{[Sub]}}$ denotes a diagonal matrix with entries $\Sigma_{i,i}^{[Sub]}, i=1,\ldots,2N_{Sub}$. For the market price of risk vector $\Lambda_{Sub} := \frac{1}{N_{Sub}}(\lambda_{Sub}^{(1)},\lambda_{Sub}^{(2)},\ldots,\lambda_{Sub}^{(1)},\lambda_{Sub}^{(2)})' \in \mathbb{R}^{2N_{Sub}}$ we assume a common market price of risk for subpopulation-specific mortality differentials $\lambda_{Sub}^{(i)}, i=1,2$ which is evenly distributed among the subpopulations. Of course, this assumption might be relaxed when fitting the model to actual market prices, which is at present not possible due to the lack of publicly available pricing information (see Section 3.3 for a discussion). Again, a normalization coefficient of the form $c_{Sub} := \exp\left(-\frac{1}{2}\Lambda_{Sub}^{'}\Lambda_{Sub}+\frac{1}{2}\Lambda_{Sub}^{'}\Sigma_\rho^{[Sub]}\Lambda_{Sub}\right)$ is required with $\Sigma_\rho^{[Sub]}$ representing the correlation matrix. Under $\mathbb{Q}$, the Gaussian model structure is preserved and subpopulation-specific mortality differentials are again driven by a multivariate normal distribution with risk-adjusted mean $-D^{\frac{1}{2}}_{\Sigma^{[Sub]}}\Sigma_\rho^{[Sub]}\Lambda_{Sub}$ and same covariance matrix as under $\mathbb{P}$ (see Theorem 5.1 in Wang (2007)).

Besides, we omit a further adjustment for diversifiable small sample risk and we follow most of the existing literature and refrain from pricing parameter uncertainty to avoid over-complexity. Nevertheless, we account for parameter uncertainty in selected model parameters

by sampling them at the start of each simulation path from suitable distributions as outlined in Appendix A. The change of measure for pricing purposes is then conducted conditional on the drawn parameter set. As summarized in Table 1, the basis model structure is retained under this change of measure.

However, the question of how to calibrate the market prices $\lambda^{(i)}_{Risk} > 0, i = 1, 2$ assigned to the individual risk drivers $Risk \in \{O, S, M, \epsilon, Sub\}$, i.e., which span-deflators to choose from the derived set of reasonable candidates, remains open. As long as the market lacks completeness, a unique risk-adjusted measure cannot be deduced from market data. We address the calibration of the market price of longevity risk in Section 3.3.

3.2 Cost of capital approach

Under modern risk-based solvency regimes, insurance companies are required to back their business with adequate economic capital. From the perspective of a longevity hedge provider, the expected return needs to be higher than the additional expected costs of providing economic capital for assuming the risk. Hence, the expected cost of capital for taking on the risk can be interpreted as a minimum risk loading on top of the objective best estimate value.

Under Solvency II, the SCR is defined as the 99.5% VaR of the basic own funds over a one-year horizon, where the basic own funds correspond to the difference between the market value of assets and the market value of liabilities. Loosely speaking, the SCR corresponds to the capital required to cover all losses which may occur over the following year at a confidence level of at least 99.5%. Alternatively, the Solvency II directive allows companies to apply a standard formula for longevity risk based on a simplified one-off shock approach. However, we focus on a risk-based (partial) internal model since the standard formula's adequacy for pricing purposes is questionable due to structural shortcomings, see Börger (2010). In what follows, we exclusively focus on the longevity SCR and neglect any further potential diversification effects with an existing business mix (see Section 3.3 for a discussion of diversification effects in the context of this pricing approach).

From the risk taker's perspective, longevity risk over a one-year horizon consists of two components (cf. Börger (2010)):

- First, the risk that next year's realized mortality will be lower than anticipated, e.g. due to a mild winter, triggering higher (or more) payments than anticipated.
- Second, the risk that the company will have to depreciate its position at the end of the year due to revised long-term mortality assumptions for the time beyond.

Clearly, both components of longevity risk might have an impact on the value of a longevity-linked security. Following Börger (2010) and Börger et al. (2019b), we assume that the remaining assets do not contribute to the longevity SCR. Moreover, we assume that there is no loss-absorbing capacity of technical provisions and that the value of the instrument does not include a risk margin when computing the SCR. Under these assumptions, the SCR in year T for issuing a longevity-linked instrument H is defined as the 99.5% quantile of

$$P\left(T, T+1\right)\left(\widetilde{H}_{T+1} + h_{T+1}\right) - \widetilde{H}_T,$$

where

- $\widetilde{H}_T$ denotes the time-T best estimate value of all future cash flows,

- h_{T+1} denotes the random cash flows that might occur over the year,
- $\widetilde{H}_{T+1}$ denotes the best estimate value of all future cash flows at time $T+1$ according to potentially revised mortality assumptions, and
- the discount factor $P(T, T+1)$ represents the time-T price of a zero coupon bond with maturity $T+1$.

The instrument's valuation based on most recent best estimate mortality assumptions is conducted by using a pragmatic best estimate central path of mortality derived under the EMT valuation model and the prevailing risk-free interest rate term structure.

The time zero random present value of the cost of capital is then given by

$$CoC := r_{CoC} \sum_{t \geq 0} P(0, t+1) SCR_t,$$

where r_{CoC} denotes the minimum rate of return in excess of the risk-free rate that investors expect for providing capital. Obviously, the SCR at a future point in time depends on realized mortality up to that point in time and hence is a random variable. Consequently, the accumulated cost of capital is also a random variable and its exact computation requires the determination of each year's SCR conditional on realized mortality. We deal with this additional layer of complexity within a two-level nested Monte Carlo simulation: Given an outer simulation path for realized mortality, the SCR at a certain point in time is computed by simulating mortality over a one-year horizon, reevaluating the position using the EMT model component, and empirically deriving the 99.5% quantile. Note that the use of a pragmatic central path of mortality in the EMT model component avoids an additional layer of nested simulations in practical applications and constitutes a valid approach for symmetric payout structures since it provides a reasonable approximation for the mean. Nevertheless, the valuation of option-type contracts over a multi-year horizon would call for a third level of nested simulations (or for a more tractable stochastic mortality model). Overall, the entire distribution of the time zero random present value of future capital charges can be derived.

The risk taker demands at least the instrument's objective best estimate value plus the expected cost of capital giving a lower bound for the contract premium of

$$P_0 \geq \sum_{t>0}^{\tau} P(0,t) \mathbb{E}^{\mathbb{P}}(h_t) + \mathbb{E}^{\mathbb{P}}(CoC).$$

Assuming a cost of capital rate of 6%, the expected cost of capital corresponds to the *risk margin* under Solvency II, which has to be reserved in addition to the best estimate liabilities. The risk margin serves as a risk loading for non-hedgeable risk, which has to be provided in order to guarantee a proper run-off. In practice, the risk margin is usually approximated to avoid nested simulations. For the required projection of SCRs, several approximations have been proposed in the literature, see e.g. Börger (2010). Typically, future SCRs are estimated by assuming a future mortality evolution according to its current best estimate, see e.g. Zeddouk and Devolder (2019). In contrast, our method allows to empirically derive an entire distribution for future capital charges from which the risk margin can be derived.

3.3 Discussion and qualitative comparison

In this section, we briefly compare and comment on the presented pricing approaches with regard to their economic justification, practical applicability, and calibration.

3.3.1 Economic justification

As argued by Cairns et al. (2006), pricing longevity-linked securities under a risk-adjusted measure would only be fully justified in a complete market in which all longevity-linked payments can be replicated. Hence, given the current stage of the market, the risk-adjusted measure is not unique and different market players might rely on different choices. However, in light of the potential size of the global longevity risk transfer market, further market development towards higher liquidity can be expected (cf. Blake et al. (2019)). Once a deep and liquid market for longevity risk develops, observable market prices might ultimately imply a unique risk-adjusted measure. Even if risk-adjusted valuation under a pricing measure is not yet fully economically justified for longevity risk pricing, it provides at least a theoretically valid approach for deriving consistent values which might eventually prevail in a fully developed market.

The cost of capital approach on the other side is inspired by capital requirements under modern risk-based solvency regimes for (re)insurers. It is in line with the interpretation of the risk margin under Solvency II as a loading for non-hedgeable risk that other (re)insurance companies would be expected to require to take over the obligations. Of course, this argument only applies to (re)insurers that are regulated by a solvency regime. Entities for which these capital requirements are not legally binding might rely on alternative economic capital models and might in turn arrive at different prices. Nevertheless, since the global longevity risk transfer market is at present still dominated by reinsurance companies, this approach seems to be of high practical relevance and economically justified.

3.3.2 Practical applicability

From a practical point of view, risk-adjusted valuation under a pricing measure offers several desirable features. As previously outlined, the applied change of measure preserves the basic model structure which allows to directly simulate the risk-adjusted distribution and to efficiently price various types of longevity-linked securities, including option-type contracts, in a Monte Carlo simulation. In particular, a set of risk-adjusted forward rates can be derived for a class of small building blocks, such as q-forwards or s-forwards. More complex instruments which can be written as a linear combination of these smaller building blocks can then be priced based on these forward rates due to the linearity of the pricing operator.

By contrast, the cost of capital approach is more complex to implement and (at least in the presented form) limited to instruments with symmetric payout structures. Furthermore, it merely offers a valid starting point for a reasonable risk premium rather than a definite value. While the expected cost of capital provides a lower bound for the risk premium, prices might also depend on potential diversification benefits with the risk taker's existing business mix. In practice, SCRs are typically derived for each submodule separately and then aggregated by explicitly taking into account potential diversification benefits between the submodules. In order to evaluate the contribution of the longevity-linked instrument to the overall SCR (and in turn also on the expected cost of capital), a reallocation principle (such as the Euler Principle or the Covariance Principle) can be applied. This would allow to quantify the expected additional cost of capital for issuing a longevity-linked security net of diversification benefits. Consequently, different companies might have different appetites for longevity risk and the pricing operator is therefore not necessarily linear. As already stated in the previous section, we do not explicitly deal with such diversification effects.

3.3.3 Calibration

The price at which a party is willing to offer a longevity-linked instrument when applying the proposed approaches obviously depends on the assumed distortion parameters, i.e. the market prices of the longevity risk drivers, and the company's cost of capital rate respectively.

There is an ongoing debate on how to find reasonable market prices of longevity risk. If a rich set of price data was publicly available, the optimal distortion parameters could be numerically calibrated to fit market prices in terms of a predefined target function, see e.g. Chen and Cox (2009). However, this approach would require estimation of up to ten distortion parameters in our setup. Unfortunately, sufficient pricing information on long-term longevity-linked securities is currently not publicly available. As argued by Wang (2007) and Chen and Cox (2009), multivariate normalized exponential tilting ensures a consistent interpretation of the tilting parameters among different risks via percentile mapping to standard normal variables. This justifies the widely used assumption of equal market prices for all risk drivers and allows to interpret the common distortion parameter, denoted by $\lambda > 0$, as market price of longevity risk, see e.g. Cairns et al. (2006), Chen and Cox (2009), or Leung et al. (2018). Under this restriction, the calibration narrows down to the specification of a single parameter. Nevertheless, even though several ideas have been proposed on how to find a reasonable market price of longevity risk in the absence of reliable market data, no consensus has been reached. For instance, Lin and Cox (2005) rely on market annuity quotes, Loeys et al. (2007) attempt to deduce a reasonable market price from other asset classes like stocks, or Xu et al. (2019) determine implied market prices in the BlackRock CoRI Retirement Indexes.

Naturally, there is also a large uncertainty regarding the 'correct' cost of capital rate. In practice, the anticipated rate of return at which shareholders are willing to provide equity might differ among different (re)insurers depending on their solvency and credit rating. For the risk margin computation under Solvency II, the cost of capital rate is currently set to 6%. Since this rate is legally binding for (re)insurers which are subject to Solvency II, Börger (2010) argues that it offers a valid starting point for finding reasonable market prices of longevity risk: The market price of longevity risk can be calibrated to the capital charges for a given longevity-linked instrument so that the risk-adjusted measure reproduces this risk loading (cf. Börger (2010), Zeddouk and Devolder (2019)). Given the interpretation of the risk margin under Solvency II, this is also in the spirit of Wüthrich and Merz (2013) who argue that risk-adjusted valuation based on probability distortions should provide values at which *"liabilities could be transferred between two knowledgeable and willing parties in an arm's length transaction."* In light of the current market environment, we believe that the presented cost of capital approach provides an appropriate reference framework for finding a reasonable market price of longevity risk. Obviously, the results typically depend on the considered instrument and vary between companies. While (re)insurers will typically target a return on equity above the regulatory cost of capital rate on the one hand, they may, on the other hand, also accept lower risk premiums due to diversification benefits. Overall, we believe that a calibration based on a cost of capital rate of 6% yields a reasonable range for the market price of longevity risk. We implement this approach in the next section.

4 Numerical illustrations

For our numerical illustrations, we calibrate the mortality model to the historical mortality experience of the male population of England and Wales. For the reference population,

we consider the national population and we rely on $N_{Sub} = 5$ socioeconomic subpopulations (ordered from the most to the least deprived areas) based on the Index of Multiple Deprivation (IMD) for England. We extrapolate mortality up to a limiting age of 130. For more details on the model calibration, including the underlying data set and the resulting model parameters, we refer to Appendix A. Furthermore, we assume a risk-free interest rate of $r = 2\%$ and a cost of capital rate of $r_{CoC} = 6\%$. We perform a two-level nested Monte Carlo simulation with 10,000 outer sample paths and additional 10,000 inner one-year scenarios for the SCR computations.

In this section, we apply the proposed pricing methods in a simplified setting. By way of illustration, we consider a market participant wishing to hedge the longevity risk arising from a portfolio of immediate (or deferred) life annuities. Starting at the retirement age of $x_R = 65$, the beneficiaries receive a lifelong payment of one unit of currency paid at the beginning of each year. The portfolio consists of $N_{Book} = 10,000$ policyholders belonging to the same cohort with starting age x_0 at time zero. We consider three exemplary portfolios with different starting ages $x_0 \in \{50, 65, 80\}$. Assuming an initial socioeconomic book composition of $\eta = (0, 0, 30\%, 30\%, 40\%)$, the number of initial policyholders coming from subpopulation $p \in \{1, \dots, N_{Sub}\}$ is given by $B^{[p]}_{x_0,0} := \eta_p N_{Book} \in \mathbb{N}$.

In Section 4.1, we introduce longevity swaps as suitable hedging instruments for this setting. As outlined in the previous section, we then derive a range of reasonable market prices of longevity risk in Section 4.2. Relying on these calibrated market prices of longevity risk, we finally derive risk-adjusted forward rates in Section 4.3.

4.1 Longevity swaps

At present, the global longevity risk transfer market is dominated by so-called longevity swaps (cf. Blake et al. (2019)). These contracts can either be customized, i.e., tailored to the individual characteristics of the hedger's own book population, or index-based, i.e., linked to more general mortality indices of a reference population.

4.1.1 Customized longevity swaps

In a customized longevity swap, the risk taker pays a floating series of cash flows corresponding to the realized annuity payments in exchange for a set of pre-agreed payments resulting in annual net payments of

$$h_t := \mathbb{1}_{\{x_0+t \geq x_R\}} \left(S_{x_0+t,t} - F_{x_0+t,t} \right), \qquad 0 \leq t \leq \tau,$$

where the floating lag $S_{x_0+t,t} := \sum_{p=1}^{N_{Sub}} B^{[p]}_{x_0+t,t}$ corresponds to the sum of all annuity payments in year t for which the annuity provider is liable. The fixed forward rates $F_{x_0+t,t}$ need to be specified at inception through the pricing approach. By construction, an unlimited ($\tau = \infty$) customized longevity swap with payments until the last policyholder has died provides perfect protection against longevity risk.

4.1.2 Index-based longevity swaps

Alternatively, the instrument can be linked to the mortality indices of the reference population resulting in annual payments of

$$h_t := \mathbb{1}_{\{x_0+t \geq x_R\}} \left(S^{[R]}_{x_0+t,t} - F^{[R]}_{x_0+t,t} \right), \qquad 0 \leq t \leq \tau,$$

where

$$S^{[R]}_{x_0+t,t} := \sum_{p=1}^{N_{Sub}} B^{[p]}_{x_0,0} \prod_{s=0}^{t-1} \left(1 - q^{[R]}_{x_0+s,s+1}\right), \qquad t \geq 0$$

denotes the ex-post survival probability for an individual from the reference population aged x_0 at time zero and $F^{[R]}_{x_0+t,t}$ denotes the index-based forward rates, which need to be specified. From the hedger's perspective, this design gives rise to population basis risk since the evolution of the book population might deviate from that captured by the reference population due to systematic socioeconomic mortality differentials or due to unsystematic fluctuations. For discussions on population basis risk in longevity hedging, we refer to Li et al. (2019) or Villegas et al. (2017). On the other side, an index-based design might be more appealing to institutional investors since it does not require any knowledge on the characteristics of the insured population.

4.2 Capital charges and the market price of longevity risk

We start by analyzing capital charges for issuing unlimited customized longevity swaps. Figure 2 shows histograms for the cost of capital and the corresponding risk margins (as the expected capital charges $\mathbb{E}^{\mathbb{P}}(CoC)$) for the three considered portfolios of different starting age.

As anticipated, the expected capital charges are higher for lower starting ages since the company has to provide regulatory capital over a longer contract duration. Moreover, the company's exposure to longevity risk is more pronounced for lower starting ages since longevity risk has more time to accumulate, which in turn requires higher SCRs. Interestingly, the histograms reveal a considerable uncertainty regarding future capital charges. In fact, their standard deviations, which are listed in Table 2, result in considerable coefficients of variation between 15% and 23%. We conclude that regulatory capital charges differ significantly between the simulation paths. The reason for this is twofold. First, SCRs over time naturally depend on realized mortality, in particular on the occurrence of unanticipated trend changes. While negative trend changes lead to higher SCRs in the following years, positive trend changes imply a capital relief. Obviously, scenarios of strong mortality improvements require higher SCRs against longevity risk over a longer time horizon compared to scenarios of higher mortality. Second, even before a trend change occurs, the overall level of SCRs generally differs between the simulation paths since we account for parameter uncertainty in the trend change parameters as outlined in Appendix A.

In a second step, we fix the obtained risk margins and numerically derive the corresponding market prices of longevity risk λ so that the risk-adjusted pricing approach produces the same risk premium as the cost of capital approach. Table 2 shows the risk margins and the corresponding market prices of longevity risk for the considered starting ages. As expected, the calibrated market prices of risk are positive since the demand for de-risking is coming from an annuity provider. As pointed out by Cairns et al. (2006), λ might actually be negative if the market was dominated by hedgers of mortality risk. Furthermore, we observe that for lower starting ages higher values of λ are required to reproduce the risk loading derived with the cost of capital approach. This phenomenon is due to the fact that long-term liabilities that are exposed to longevity risk are rather capital intensive under risk-based internal models for longevity risk which implies slightly higher values for the market price of longevity risk.

For comparison, values between 0.15 and 0.4 are commonly reported in the literature for the market price of longevity risk. For instance, Leung et al. (2018) report market prices of

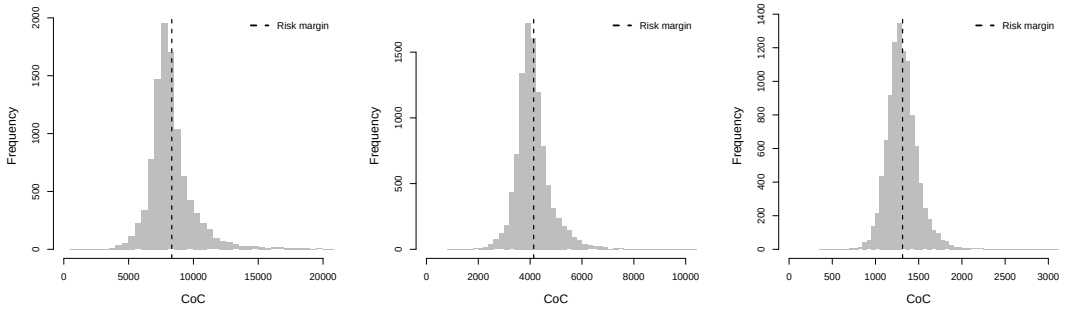

(a) $x_0 = 50$ (b) $x_0 = 65$ (c) $x_0 = 80$

Figure 2: Histograms for the random present value of the cost of capital and corresponding risk margins for issuing unlimited customized longevity swaps for different starting ages.

	$x_0 = 50$	$x_0 = 65$	$x_0 = 80$
$\mathbb{E}^{\mathbb{P}}(CoC)$	8,330	4,140	1,310
$\mathrm{SD}^{\mathbb{P}}(CoC)$	1,920	750	200
corresponding λ	0.315	0.300	0.225

Table 2: Mean and standard deviation of the cost of capital for issuing unlimited customized longevity swaps for different starting ages along with corresponding market prices of longevity risk.

longevity risk between 0.15 and 0.35 for the risk-adjusted pricing approach in their comparative study of different pricing techniques based on the CBD model. Levantesi and Menzietti (2017) arrive at slightly higher values of about 0.41 for the first time-dependent parameter and of around 0.35 for the second period effect when matching the risk-adjusted CBD model with the cost of capital approach. However, it should be kept in mind that reported market prices of longevity risk are typically model specific and not directly comparable due to differences in the intrinsic model setups. Nevertheless, our calibrated market prices of longevity risk seem to be of reasonable magnitude.

4.3 Risk-adjusted forward rates

Relying on the specified set of market prices of longevity risk, we apply the risk-adjusted pricing approach to derive forward rates for customized and index-based longevity swaps. Figure 3 shows quantile plots for the T-year survival probabilities along with their best estimate values (black solid lines) and their risk-adjusted counterparts (red dashed lines) for the considered starting ages, where the upper row shows the customized and the lower row the index-based survival curves.

We first look at the customized designs. When comparing the prediction bands for different starting ages, we observe that the uncertainty in future survival probabilities is much more pronounced for lower starting ages. This is due to the fact that there is more time for potential future trend changes to unfold. Furthermore, we observe that the applied stochastic trend process produces prediction bands which are rather narrow in the short-term and widen over time reflecting the nature of longevity risk as a typically slowly accumulating demographic risk. With regard to this increasing degree of uncertainty over time, the risk-adjusted forward

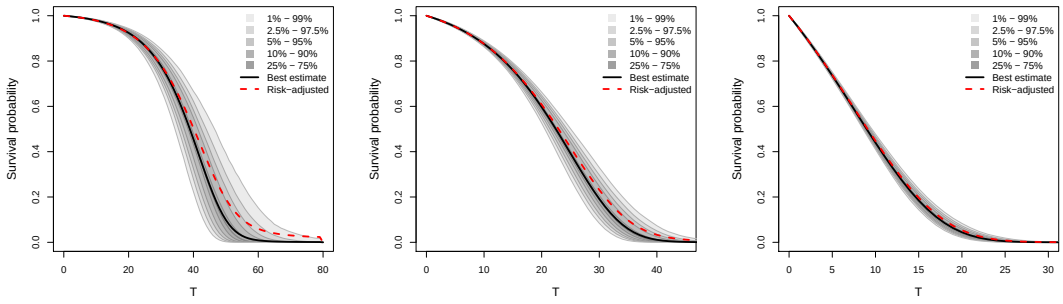

(a) $x_0 = 50$, customized (b) $x_0 = 65$, customized (c) $x_0 = 80$, customized

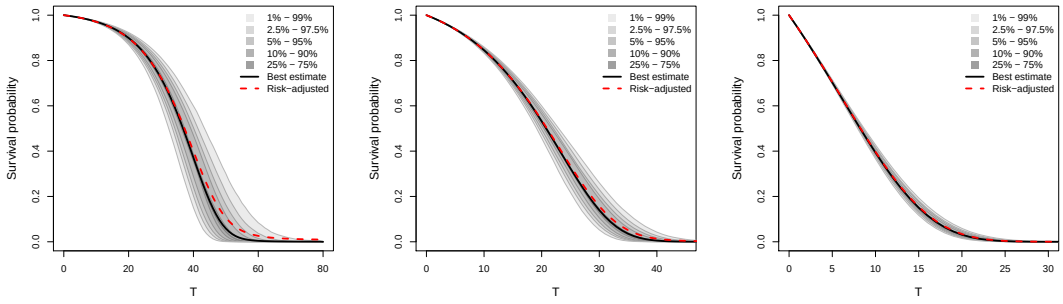

(d) $x_0 = 50$, index-based (e) $x_0 = 65$, index-based (f) $x_0 = 80$, index-based

Figure 3: Quantile plots, best estimate, and risk-adjusted survival probabilities for customized (upper row) and index-based (lower row) longevity swaps for different starting ages.

rates show a reasonable pattern compared to their best estimate counterparts: The forward rates are typically higher and the distances between both curves, which indicate the implied risk loading, increase over time in line with the growing uncertainty regarding future survival rates. Obviously, most of the risk premium is allocated to longer maturities since these cash flows are naturally most affected by potential future changes in the long-term mortality trend.

When comparing the index-based (lower row) with the customized survival curves (upper row), we observe that the index-based prediction bands are slightly narrower. While long-term mortality trend risk constitutes the only source of uncertainty for the index-based rates, the customized rates are also affected by socioeconomic mortality differentials as well as by small sample risk. Due to the assumed existence of a market price for socioeconomic mortality differentials, the implied risk loading is notably higher for the customized designs than for the index-based variants. However, besides being narrower, the prediction bands are similar in shape to those for the customized designs. Hence, it can be concluded that most of the uncertainty in future mortality originates from the long-term mortality trend risk of the reference population. This underlines the importance to properly model and price the risk of unexpected future changes in the long-term mortality trend.

Overall, we conclude that risk-adjusted valuation under our derived pricing measure produces highly plausible forward rates in terms of intrinsic risk loading. Finally, we would like to stress that also longevity swaps with a finite time to maturity (of for instance 30 years) might be priced based on the derived forward rates due to the linearity of the pricing operator.

5 Conclusion

The pricing of longevity-linked securities in the absence of a liquid and complete market for longevity risk has recently been addressed by many authors. Most of the proposed pricing approaches are based on stochastic mortality models in which one or more time-dependent parameters are projected into the future. As argued by several authors in the field of longevity trend modeling, long-term mortality projections should account for the risk of future mortality trend changes. Otherwise, the risks involved in a longevity transaction will be systematically understated which might lead to insufficient risk premiums and to inadequate prices.

We contribute to the literature on longevity risk pricing by addressing the determination of adequate longevity risk premiums within a stochastic modeling framework which explicitly models the risk of random future changes in the long-term mortality trend. We construct a set of meaningful stochastic discount factors based on suitable probability distortions which imply equivalent risk-adjusted pricing measures under which the basic model structure is preserved. These deflators help to understand how the risk-adjusted measure relates to its objective counterpart, under which the model is calibrated. To find a reasonable range for the market price of longevity risk in the current incomplete market environment, we follow Börger (2010) and rely on a cost of capital approach in the context of a regulatory capital model under Solvency II. Unlike previous studies, we explicitly derive entire distributions for future capital charges for longevity risk from which adequate risk premiums can be deduced. Moreover, we also allow mortality in selected subpopulations to differ from the underlying trend of the overall population. Hence, our framework is applicable to customized as well as to index-based instruments. We discuss the presented approaches with regard to their economic justification, practical applicability, and calibration.

In a numerical illustration, we apply the proposed methods to price longevity swaps. For different contract designs, we first derive the required risk margin under the cost of capital approach and then construct the risk-adjusted measure so that both approaches imply the same risk premium. We show that this results in a reasonable range of market prices for longevity risk. Furthermore, we demonstrate that the risk-adjusted pricing approach produces highly plausible forward rates in terms of intrinsic risk loading. We find that in the presence of mortality trend changes most of the risk premium is allocated to longer maturities. This seems reasonable given the longer time horizon for longevity trend changes to unfold. In light of the potentially devastating impact of an unanticipated mortality trend change on the solvency of longevity risk bearers, we conclude that our modeling framework constitutes a valuable alternative to existing approaches for longevity risk pricing.

Interestingly, we find that future capital charges for longevity risk derived under a risk-based internal model are subject to a considerable degree of uncertainty. Since commonly used approximations for the risk margin only provide a point estimate for future capital charges rather than an entire distribution, they might systematically underestimate longevity risk. In particular in scenarios of strong mortality improvements, annuity providers might not only be liable for more payments to policyholders but they also need to account for higher capital charges for longevity risk. We plan to address this issue and the potential implications for longevity risk management in future works.

References

Bauer, D., Börger, M., and Ruß, J. (2010). On the pricing of longevity-linked securities. *Insurance: Mathematics and Economics*, **46**(1):139–149.

Blake, D., Cairns, A. J. G., Dowd, K., and Kessler, A. R. (2019). Still living with mortality: The longevity risk transfer market after one decade. *British Actuarial Journal*, **24**, e1:1–80.

Börger, M. (2010). Deterministic shock vs. stochastic value-at-risk — an analysis of the Solvency II standard model approach to longevity risk. *Blätter der DGVFM*, **31**(2):225–259.

Börger, M., Schönfeld, J., and Schupp, J. (2019a). *Calibrating mortality processes with trend changes to multi-population data*. Working Paper, Ulm University.

Börger, M. and Schupp, J. (2018). Modeling trend processes in parametric mortality models. *Insurance: Mathematics and Economics*, **78**:369–380.

Börger, M., Schupp, J., and Ruß, J. (2019b). *It Takes Two: Why Mortality Trend Modeling is more than modeling one Mortality Trend*. Working Paper, Ulm University.

Boyer, M. M. and Stentoft, L. (2013). If we can simulate it, we can insure it: An application to longevity risk management. *Insurance: Mathematics and Economics*, **52**(1):35–45.

Cairns, A. J. G., Blake, D., and Dowd, K. (2006). A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *The Journal of Risk and Insurance*, **73**(4):687–718.

Chen, H. and Cox, S. H. (2009). Modeling mortality with jumps: Applications to mortality securitization. *The Journal of Risk and Insurance*, **76**(3):727–751.

Hunt, A. and Blake, D. (2014). A general procedure for constructing mortality models. *North American Actuarial Journal*, **18**(1):116–138.

Lee, R. D. and Carter, L. R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, **87**(419):659–671.

Leung, M., Fung, M. C., and O'Hare, C. (2018). A comparative study of pricing approaches for longevity instruments. *Insurance: Mathematics and Economics*, **82**:95–116.

Levantesi, S. and Menzietti, M. (2017). Maximum market price of longevity risk under solvency regimes: The case of Solvency II. *Risks*, **5**(2):1–21.

Li, J., Li, J. S.-H., Tan, C. I., and Tickle, L. (2019). Assessing basis risk in index-based longevity swap transactions. *Annals of Actuarial Science*, **13**(1):166–197.

Li, J. S.-H., Chan, W.-S., and Cheung, S.-H. (2011). Structural changes in the Lee-Carter mortality indexes: detection and implications. *North American Actuarial Journal*, **15**(1):13–31.

Li, J. S.-H., Chan, W.-S., and Zhou, R. (2017). Semicoherent multipopulation mortality modeling: The impact on longevity risk securitization. *The Journal of Risk and Insurance*, **84**(3):1025–1065.

Lin, Y. and Cox, S. H. (2005). Securitization of mortality risks in life annuities. *The Journal of Risk and Insurance*, **72**(2):227–252.

Loeys, J., Panigirtzoglou, N., and Ribeiro, R. M. (2007). *Longevity: A market in the making*. JPMorgan Global Market Strategy.

Schupp, J. (2019). *On the modeling of variable mortality trend processes*. Working Paper, Ulm University.

Tan, C. I., Li, J., Li, J. S.-H., and Balasooriya, U. (2014). Parametric mortality indexes: From index construction to hedging strategies. *Insurance: Mathematics and Economics*, **59**:285–299.

Villegas, A. M. and Haberman, S. (2014). On the modeling and forecasting of socioeconomic mortality differentials: An application to deprivation and mortality in England. *North American Actuarial Journal*, **18**(1):168–193.

Villegas, A. M., Haberman, S., Kaishev, V. K., and Millossovich, P. (2017). A comparative study of two-population models for the assessment of basis risk in longevity hedges. *ASTIN Bulletin*, **47**(3):631–679.

Villegas, A. M., Millossovich, P., and Vladimir, K. (2018). StMoMo: Stochastic mortality modeling in R. *Journal of Statistical Software*, **84**(3):1–38.

Wang, S. (2007). Normalized exponential tilting. *North American Actuarial Journal*, **11**(3):89–99.

Wang, S. S. (2002). A universal framework for pricing financial and insurance risks. *ASTIN Bulletin*, **32**(2):213–234.

Wüthrich, M. V. (2016). *Market-consistent actuarial valuation*. New York: Springer.

Wüthrich, M. V. and Merz, M. (2013). *Financial modeling, actuarial valuation and solvency in insurance*. New York: Springer.

Xu, Y., Sherris, M., and Ziveyi, J. (2019). Market price of longevity risk for a multi-cohort mortality model with application to longevity bond option pricing. *The Journal of Risk and Insurance*.

Zeddouk, F. and Devolder, P. (2019). Pricing of longevity derivatives and cost of capital. *Risks*, **7**(2):1–29.

Zhou, R., Li, J. S.-H., and Tan, K. S. (2015). Economic pricing of mortality-linked securities: A tâtonnement approach. *The Journal of Risk and Insurance*, **82**(1):65–96.

Appendix

A Model calibration

This appendix provides the model calibration to the historical mortality experience of English and Welsh males.

Data

For the reference population, we use data of the male population of England and Wales for the years 1841 to 2016 over the age range of 60 to 109, which is available in the Human Mortality Database (data downloaded on 01 July 2018 from: http://www.mortality.org). Regarding the subpopulations, we rely on mortality data of English males sorted by quintiles of the Index of Multiple Deprivation (IMD) for the years 2001 to 2016 over the age range of 60 to 89, which we obtained from the Office for National Statistics. Note that the relative modeling setting can deal with a shorter data history and a different age range for the subpopulations than for the reference population. The IMD is the official measure of relative deprivation for small areas (so-called Lower Super Output Areas) in England. By combining information from seven domains, including income, employment rate, education, and health, the index measures deprivation in 32,844 small areas with an average of 1,500 residents each. Based on their IMD-score, these areas are then grouped into quintiles (ordered from the most to the least deprived areas) to obtain five homogeneous subpopulations of the male English population. We calibrate the underlying multi-population CBD model structure via a standard two-stage maximum likelihood estimation (MLE) approach assuming binomially distributed deaths, see Villegas et al. (2018).

Extrapolation for high ages

When calibrating the CBD model to different socioeconomic subpopulations, one needs to be aware that the fitted mortality lines in the logit plot typically cross over somewhere in the highest age range. The reason for this phenomenon lies in the structure of socioeconomic mortality differentials, which are usually more distinct for younger ages and diminish for high ages, see e.g. Villegas et al. (2017). To avoid this crossing over, we adjust the extrapolation by setting the mortality rates equal to those of the reference population as soon as its mortality curve would be crossed, i.e., we set

$$q_{x,t}^{[p]} := \begin{cases} \max\left\{q_{x,t}^{[R]}; q_{x,t}^{[p]}\right\}, & \text{for } p = 1, 2 \\ \min\left\{q_{x,t}^{[R]}; q_{x,t}^{[p]}\right\}, & \text{for } p = 4, 5 \end{cases}$$

and leave the middle subpopulation, which might lie (depending on the mortality scenario) above or below the reference population, unadjusted.

Stochastic trend process

For the calibration of the stochastic AMT process, we follow Schupp (2019) and apply an iterative pseudo MLE approach. The technical details of this calibration algorithm can be found in Schupp (2019). Since the calibration of stochastic mortality trend processes typically involves a considerable amount of uncertainty, we explicitly account for parameter uncertainty

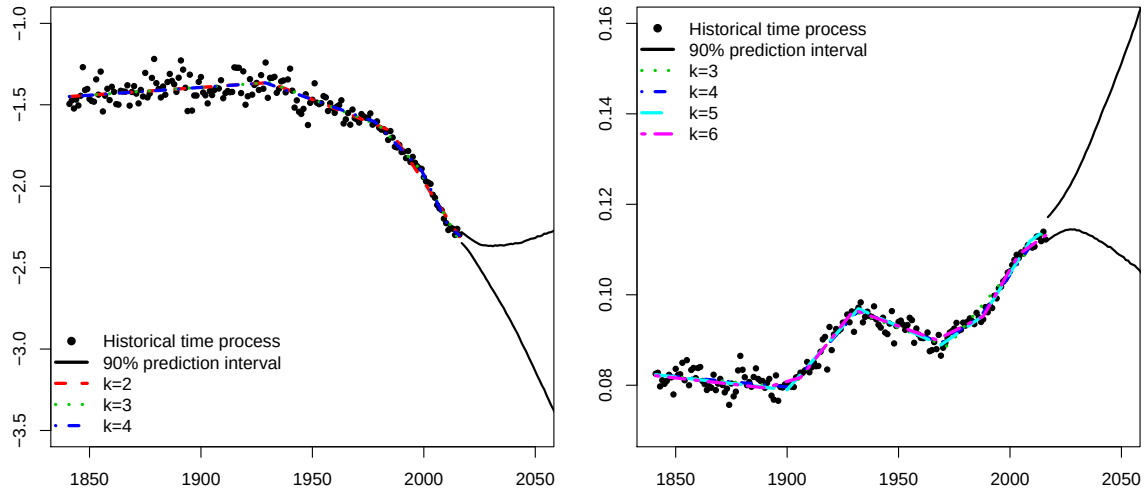

(a) $\kappa_t^{(1)[R]}$ (b) $\kappa_t^{(2)[R]}$

Figure 4: Historical trend processes for English and Welsh males (dotted), best possible realizations for the actual trend processes given different numbers of trend changes k (colored dashed lines), and 90% prediction intervals (black solid lines)

k	$\hat{\kappa}_{t=0}^{(1)[R]}$	$\hat{d}_{t=0}^{(1)[R]}$	$\mathbb{P}$
2	-2.3099	$-2.10 \cdot 10^{-4}$	6.24%
3	-2.3374	$-2.43 \cdot 10^{-4}$	0.54%
4	-2.3020	$-1.15 \cdot 10^{-4}$	93.22%

k	$\hat{\kappa}_{t=0}^{(2)[R]}$	$\hat{d}_{t=0}^{(2)[R]}$	$\mathbb{P}$
3	0.1144	$5.85 \cdot 10^{-4}$	43.14%
4	0.1156	$6.96 \cdot 10^{-4}$	3.89%
5	0.1143	$3.55 \cdot 10^{-4}$	34.67%
6	0.1131	$3.49 \cdot 10^{-4}$	18.30%

Table 3: Empirical distributions for the AMT starting values

in the starting values of the simulation as well as in the trend change parameters (cf. Börger et al. (2019a)).

Figure 4 shows the historical trend processes $\kappa_t^{(i)[R]}, i = 1, 2$ and the best possible realizations for the underlying trend processes $\hat{\kappa}_t^{(i)[R]}, i = 1, 2$ for all relevant potential numbers (k) of historical trend changes, which are identified by the calibration algorithm. The parameter estimates and assigned weights are summarized in Table 3. We account for parameter uncertainty by sampling the starting values from the respective empirical distribution.

Moreover, we obtain the following trend change parameters:

$$\left(\hat{p}^{(1)[R]}, \hat{\mu}_M^{(1)[R]}, \hat{\sigma}_M^{(1)[R]}\right) = (0.0223, -4.5453, 0.4105)\,,$$
$$\left(\hat{p}^{(2)[R]}, \hat{\mu}_M^{(2)[R]}, \hat{\sigma}_M^{(2)[R]}\right) = (0.0246, -7.4134, 0.2027)$$

with corresponding covariance matrices of standard errors of

$$SE^{(1)} = \begin{pmatrix} 1.353 \cdot 10^{-4} & 3.535 \cdot 10^{-5} & -1.754 \cdot 10^{-5} \\ 3.535 \cdot 10^{-5} & 4.616 \cdot 10^{-2} & 3.331 \cdot 10^{-4} \\ -1.754 \cdot 10^{-5} & 3.331 \cdot 10^{-4} & 2.322 \cdot 10^{-2} \end{pmatrix},$$

$$SE^{(2)} = \begin{pmatrix} 1.860 \cdot 10^{-4} & -1.211 \cdot 10^{-3} & 8.060 \cdot 10^{-4} \\ -1.211 \cdot 10^{-3} & 4.285 \cdot 10^{-2} & -2.153 \cdot 10^{-2} \\ 8.060 \cdot 10^{-4} & -2.153 \cdot 10^{-2} & 2.124 \cdot 10^{-2} \end{pmatrix}.$$

We follow Börger et al. (2019a) and account for parameter uncertainty in these parameters as follows: For each simulation path, we generate a multivariate normal random vector with mean equal to the estimated trend change parameters and covariance matrix $SE^{(i)}$. Subsequently, the first component of the vector is transformed to a Beta distribution with same mean and variance to obtain trend change probabilities between zero and one. Similarly, the third component is transformed to a Gamma distribution with unchanged mean and variance to ensure positivity.

Finally, we estimate the following covariance structure for the random noise around the AMT:

$$\Sigma^{[R]} = \begin{pmatrix} 3.865 \cdot 10^{-4} & 1.720 \cdot 10^{-5} \\ 1.720 \cdot 10^{-5} & 2.036 \cdot 10^{-6} \end{pmatrix}.$$

For simplicity, we neglect parameter uncertainty in the covariance matrix since it is typically not material.

EMT weights

As suggested by Börger et al. (2019b), we determine the optimal EMT weighting by minimizing the mean squared errors between the unobservable AMT and the derived EMT in a Monte Carlo simulation, which gives

$$\left(\psi^{(1)}, \psi^{(2)}\right) = (2.225, 2.752).$$

Time processes for socioeconomic mortality differentials

The RWD for socioeconomic mortality differentials is calibrated via a standard MLE approach. For simplicity, we do not account for parameter uncertainty for the subpopulations.

Contact details:

Arne Freimann
Address: Lise-Meitner-Straße 14, 89081 Ulm, Germany
Email: a.freimann@ifa-ulm.de
Phone: +49 (731) 20 644-253