It Takes Two:
Why Mortality Trend Modeling is more than modeling one Mortality Trend

Matthias Börger

Institute for Finance and Actuarial Sciences (ifa), Ulm

Lise-Meitner-Straße 14, 89081 Ulm, Germany

Phone: +49 731 20644257

Email: m.boerger@ifa-ulm.de

Jochen Russ

Institute for Finance and Actuarial Sciences (ifa), Ulm & Institute of Insurance, Ulm University

Lise-Meitner-Straße 14, 89081 Ulm, Germany

Phone: +49 731 20644233

Email: j.russ@ifa-ulm.de

Johannes Schupp (corresponding author)

Institute for Finance and Actuarial Sciences (ifa), Ulm & Institute of Insurance, Ulm University

Lise-Meitner-Straße 14, 89081 Ulm, Germany

Phone: +49 731 20644241

Email: j.schupp@ifa-ulm.de

January 29, 2019

Abstract

Increasing life expectancy and thus decreasing mortality rates constitute a global trend that can be observed in almost all countries worldwide. Estimating the current rate at which mortality rates decrease and modeling the future rate of decrease is important for e.g. demographers and actuaries. This task is commonly referred to as mortality trend modeling.

In many applications however one needs to carefully distinguish between two different mortality trends: The actual (but unobservable) mortality trend (AMT) prevailing at a certain point in time and the estimated mortality trend (EMT) that an observer would estimate given the (observable) realized mortality up to that point in time. Since the AMT is not observable, an actuary or demographer might misestimate the AMT at any point in time. In particular, he would typically not be able to distinguish between a recent chance in the actual trend and a "normal" random

fluctuation around the previous long term trend. Depending on the question at hand, the AMT or the EMT or both need to be considered and modeled in analyses.

The paper provides a clear definition of and distinction between the actual mortality trend and the estimated mortality trend, discusses their connection, and explains which of the two is relevant for which kind of question. Moreover, a numerically efficient combined model for both trends is specified and calibrated to mortality data. The model component for the actual mortality trend builds on recent findings that mortality appears to evolve log-linear over time with random changes in slope. The model component for the estimated mortality trend is specified such that, given the assumed dynamics for the actual mortality trend, the estimated mortality trend matches the actual trend as close as possible. This provides valuable information on how best estimate mortality assumptions should be derived from the available data in general.

Finally, we apply the combined model in practical examples and illustrate the importance of distinguishing between AMT and EMT. We show that, if the AMT is wrongfully assumed observable, the hedge effectiveness of a longevity hedge or the SCR for longevity risk are typically misestimated significantly.

1. Introduction and Motivation

Increasing life expectancy and thus decreasing mortality rates constitute a global trend that can be observed in almost all countries worldwide. Estimating the current rate at which mortality rates decrease and modeling the future rate of decrease is commonly referred to as mortality trend modeling. This task is important for e.g. demographers and actuaries and e.g. constitutes an important input for the pricing, reserving, and risk management of annuity and pension products.

In this paper, we deal with an aspect that is often ignored in the existing literature, i.e. the fact that one needs to distinguish between two different mortality trends: The (unobservable) actual mortality trend (AMT) prevailing at a certain point in time and the estimated mortality trend (EMT) that an observer would estimate given the realized mortality up to that point in time. These trends differ, since obviously, an observer will not be able to perfectly estimate the (unobservable) actual mortality trend. Moreover, an observer might not always be able to distinguish between a recent chance in the actual trend and a simple random fluctuation around the previous long term trend. This is particularly relevant immediately after a trend change or immediately after a rather strong random fluctuation. We will show that not distinguishing between these two trends might have significant undesired consequences.

Even very simple examples can motivate why one needs to distinguish between the two trends (we will discuss some examples in much more detail in Section 4 where we will also quantify the error that would result from not distinguishing between the two trends). As a first example, assume that one is interested in a confidence band for the cash-flow of a portfolio of annuity contracts over the next, say, 10 years or its present value. In such a run-off simulation, only actual future mortality and hence the survival rates derived from a model for the AMT need to be used.[1]

As a second example, assume that one is interested in the question how reserves for the same book of annuity contracts can change from one year to the next. Now, one needs to model actual mortality over one year as well as the evolution of the EMT over this year, since the EMT (i.e. an observes estimate for the then current trend) would be the basis for the calculation of policy reserves. Hence, a model with two components is required: the development of actual mortality (based on the AMT) and the development of the EMT which at any point in time is some function of the realized mortality up to that point in time.

These simple examples already show that for some applications a model for the future development of the AMT is needed, for other applications, a model for the future development of the EMT is needed, and frequently AMT and EMT need to be modelled simultaneously. A

[1] Nevertheless, even here, it is important to distinguish between the AMT and the EMT since one has to be aware, that today's AMT is not known and therefore the mortality trend at the start of the simulation is uncertain. We will get back to this issue below.

model that only captures one of these two trends might therefore not be suitable for certain analyses.

Therefore, in this paper, we specify a mortality model that simultaneously and consistently projects AMT and EMT. We choose a model structure that makes the combined AMT/EMT model highly efficient also in Monte Carlo simulations where the EMT needs to be estimated in each simulation path: For the AMT component, we use the model proposed by Börger and Schupp (2018) and further refined in Schupp (2018). This model builds on the model structure of Cairns et al. (2006), and the trend processes for the time dependent parameters take into account recent findings that mortality appears to evolve log-linear over time with random changes in slope. We show how the model parameters can be estimated from historical data and discuss how parameter uncertainty, in particular in the unobservable current AMT, can be accounted for. Finally, we illustrate the adequacy of the AMT model component and its parameter estimation by a concrete example calibration for English and Welsh males.

For the EMT component, we discuss and compare different model specifications. The EMT should match the AMT as closely as possible, and therefore, the assumed dynamics for the AMT provide valuable information on how the EMT should be specified. Since we assume the AMT to evolve piecewise linearly, it appears reasonable to also model a linear EMT. We therefore consider a linear regression approach and analyze the effect of different weightings within the regression. Obviously, more recent data points are more informative with respect to the current AMT than data points further in the past. The optimized weightings which are derived in this paper also provide a general indication how mortality trends should be estimated from available data.

Finally, we apply the combined model in practical examples and illustrate the importance of distinguishing between AMT and EMT. If the AMT is wrongfully assumed observable, the hedge effectiveness of a longevity hedge or the SCR for longevity risk are typically misestimated significantly.

The majority of the mortality models proposed in the literature have been developed in order to project actual mortality. Often, these models are also used to derive the EMT, typically in form of the (deterministic) central projection. For instance, Richards et al. (2014) use the Lee-Carter model (see Lee and Carter (1992)) both as an AMT and EMT model in order to compute the Solvency Capital Requirement (SCR) for longevity risk under Solvency II. They project actual mortality for one year, then recalibrate the model based on the extended data set, and finally interpret the central projection of the recalibrated model as the EMT after one year. Cairns (2013), Cairns et al. (2014), and also Cairns and El Boukfaoui (2018) essentially apply the same concept in the field of longevity hedging, but explicitly distinguish between a simulation model for the AMT and a valuation model for the EMT. These models can be chosen independently of each other. However, the repeated full (re)calibration of a valuation model as part of a Monte Carlo simulation framework is typically very expensive from a computational point of view. For this reason, Börger et al. (2014) propose a model in which only the mortality trend (as opposed to the full model) is reestimated based on the available additional data. Thus, they combine the projection of actual and estimated mortality in a single model, but create some inconsistency by interpreting the EMT also as the AMT at any point in time.

All the aforecited papers have in common that first actual mortality is simulated and then the EMT is derived from an extended data series. Plat (2011) turns this intuitive approach around when he proposes a stochastic EMT model in which, in a second step only, actual mortality rates are derived such that they are consistent with the simulated EMT change; for instance, if expected mortality increases from one year to the next, the actual mortality rates must have been larger than anticipated. Since the simulation of actual mortality is rather cumbersome without an explicit AMT, the model is particularly designed for applications in which the EMT is required in each year, e.g. for SCR computations. A similar argument holds for the class of forward mortality models. These models primarily project changes in expected mortality rates, but also provide actual mortality rates when expected mortality turns into realized mortality. Forward mortality models have been proposed by Bauer et al. (2008), Bauer et al. (2010), Zhu and Bauer (2011), or Hunt and Blake (2015, 2016). However, the specification and calibration of these models is fairly complex.

The remainder of this paper is organized as follows: In Section 2, we introduce the AMT component of our combined mortality model. We discuss how its parameters can be estimated and how parameter uncertainty can be accounted for. Moreover, we derive a full calibration for the case of English and Welsh males. Section 3 discusses how the EMT should be derived from observed mortality. In particular, we compare different approaches for weighting the available data such that the EMT adequately reacts to trend changes in the AMT. In Section 4, we provide three practical examples in which both, the AMT and the EMT are required, and we show how important it is to properly distinguish between these two trends. Finally, Section 5 concludes.

2. AMT Model Component

As outlined in the Introduction, we will now derive a full specification of a mortality model that consists of two components: the model for the unobservable actual mortality trend (AMT) and the model for the estimated mortality trend (EMT) an actuary or demographer would derive from the observable data available at any point in time. In this section, we introduce the AMT model component. The EMT model component will be discussed in Section 3.

2.1 AMT Model Specification

For the AMT, we use the model proposed by Börger and Schupp (2018) and further refined by Schupp (2018). This model builds on the well-known CBD model of Cairns et al. (2006), i.e. annual probabilities of death are modeled as

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

where $\bar{x}$ is the average age of the age range under consideration. The time dependent parameter $\kappa_t^{(1)}$ determines the general level of mortality, whereas the slope parameter $\kappa_t^{(2)}$ describes the increase of mortality with age.

In many cases, these time dependent parameters are projected by a two-dimensional random walk with drift as originally proposed by Cairns et al. (2006). However, Figure 1 indicates that assuming a constant and fixed drift might not be a reasonable assumption particularly for long-term mortality projections. The figure shows the logarithm of probabilities of death for 65-year old males in different countries all over the world; the data has been obtained from the HMD (2018). We observe for all populations that mortality trends appear linear, but change their slope once in a while. Thus, in order to project mortality consistently with historical observations, it appears reasonable to allow for trend changes also in the future. Börger and Schupp (2018) show that a random walk with constant drift might significantly underestimate the long-term uncertainty in future mortality.

Therefore, Börger and Schupp (2018) and Schupp (2018) propose a trend process which projects piecewise linear trends with random changes in slope.[2] For any future year t, the "observable" processes $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ are modeled as the sum of the actual but unobservable true mortality processes and some random noise:

$$\kappa_t^{(i)} = \hat{\kappa}_t^{(i)} + \varepsilon_t^{(i)},\ i = 1,2.$$

The noise term $\varepsilon_t^{(i)}$ accounts for annual fluctuations which are, e.g., due to flu waves, very hot summers, or catastrophes. The vector $\varepsilon_t = \left(\varepsilon_t^{(1)}, \varepsilon_t^{(2)}\right)'$ is assumed to follow a two-dimensional Normal distribution with mean zero and covariance matrix Σ.[3]

The actual mortality processes $\hat{\kappa}_t^{(i)},\ i = 1,2$ are projected linearly with the current unobservable AMTs $\hat{d}_t^{(i)},\ i = 1,2$ as slopes:

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

The AMTs remain unchanged until the next trend change occurs:

$$\hat{d}_t^{(i)} = \begin{cases} \hat{d}_{t-1}^{(i)} & ,\ if\ no\ trend\ change\ occurred\ in\ t-1 \\ \hat{d}_{t-1}^{(i)} + \lambda_{t-1}^{(i)} & ,\ if\ a\ trend\ change\ by\ \lambda_{t-1}^{(i)}\ occurred\ in\ t-1 \end{cases}.$$

[2] For a literature overview on alternative models with time dependent drift see e.g. Börger and Schupp (2018).

[3] Note that the covariance matrix is assumed constant over time even though heteroscedasticity can usually be observed in the historical data (see Figure 1). However, since the noise does not impact long-term mortality evolutions (in contrast to the innovations in the random walk), this simplification appears appropriate.

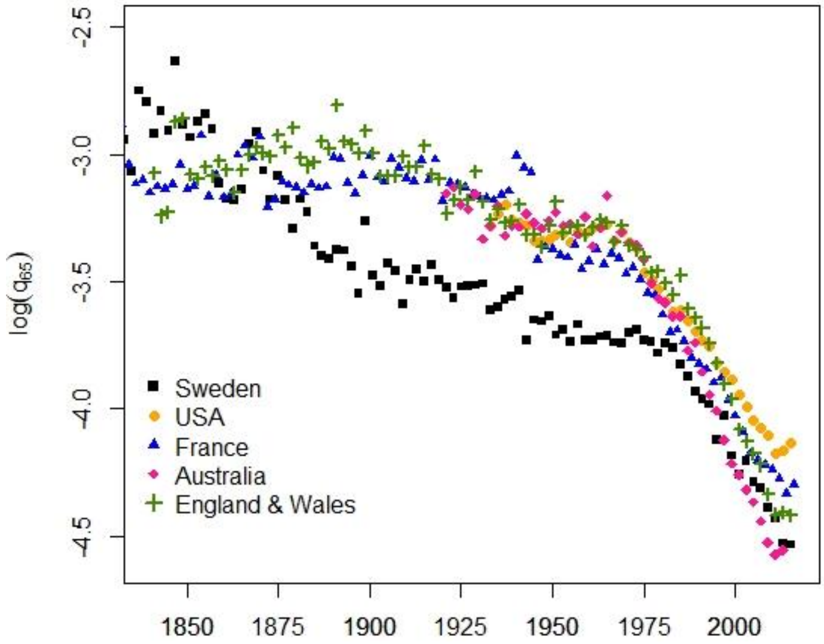

Figure 1: Logarithm of probabilities of deaths for 65-year old males in selected countries

The trend change intensities $\lambda_t^{(i)}$, $i = 1{,}2$, are derived as the product of their absolute magnitudes $M_t^{(i)}$ and their signs $S_t^{(i)}$:

$$\lambda_t^{(i)} = S_t^{(i)} \cdot M_t^{(i)}, i = 1{,}2.$$

Based on analyses of historical trend changes, Börger and Schupp (2018) propose modeling the magnitudes $M_t^{(i)}$ by Lognormal distributions $LN(\mu^{(i)},\ \sigma^{(i)})$. For the signs $S_t^{(i)}$, they use a Bernoulli distribution with attainable values -1 and 1 and probability 0.5 each. The probability of observing a trend change in $\hat{d}_t^{(i)}$ in any particular year is denoted by $p^{(i)}, i = 1, 2$. Moreover, trend changes in $\hat{d}_t^{(1)}$ and $\hat{d}_t^{(2)}$ are assumed to occur independently as indicated by the occurrences of trend changes in the historical data for a large set of populations.

This decomposition of the trend change intensities has several convenient implications. First, the distributions of future trend changes are symmetric, i.e. the slope in the AMT increases and decreases with equal probability and magnitude. Thus, the prevailing AMT is always the best estimate for the trend at any future point in time. Furthermore, the distribution of $\lambda_t^{(i)}$ has no probability mass at zero and only very little mass around zero. Thus, simulated trend changes can be considered as rather “significant”. At the same time, the heavy tail of the Lognormal distribution implies that strong trend changes can occur which is in line with some of the trend changes we observe in Figure 1 (e.g. in Sweden around 1980).

2.2 Parameter Estimation and Uncertainty

The estimation of the model parameters consists of two main steps. First, the CBD model is fitted to historical data. Here, the historical time series should be as long as possible such that it contains as many trend changes as possible. In the second step, the parameters of the trend processes $\kappa_t^{(i)}$ are estimated from the historical realizations $\kappa_t^{(i)}, t \leq t_0$, where t_0 denotes the final year of the historical data, i.e. the starting point of a simulation. The parameters to be estimated are:

- the starting values for the actual trend processes $\hat{\kappa}_{t_o}^{(i)}$,
- the prevailing AMTs $\hat{d}_{t_0}^{(i)}$,
- the probabilities of observing a trend change in a certain year, $p^{(i)}$
- the parameters of the Lognormal distributions for the trend change magnitudes, $\mu^{(i)}$ and $\sigma^{(i)}$,
- and the covariance matrix Σ of the two-dimensional noise vector ε_t.

The parameter estimation is carried out separately for each $\kappa_t^{(i)}$ process and the correlation between $\varepsilon_t^{(1)}$ and $\varepsilon_t^{(2)}$ is estimated once all other parameters have been determined. As Schupp (2018) explains, parameter estimation is complex due to the dependence of realized $\kappa_{t,t\leq t_0}^{(i)}$ on potential but unknown trend changes in previous years. In particular, a full maximum likelihood estimation of all model parameters seems impossible from a practical point of view. Therefore, we estimate parameters from specific realizations for the actual trend process $\hat{\kappa}_t^{(i)}, t \leq t_0$.

For a trend process realization with k past trend changes the obvious estimate for the future trend change probability is $p^{(i)} = k/N$, where N is the number of data points in the time series $\kappa_t^{(i)}, t \leq t_0$. For each of the (rather low number of) feasible values for $p^{(i)}$ or k, respectively, the other parameters are then optimized using a pseudo maximum likelihood algorithm. For details on this rather complex algorithm, we refer to Schupp (2018).

We then have a full set of parameter estimates for each number of possible past trend changes k. From these parameter sets, some "optimal" parameter estimates could be chosen given some optimality criterion. However, since the parameter estimation involves a certain degree of uncertainty, we rather interpret these parameter sets as possible outcomes of an empirical distribution from which parameter values can be drawn randomly for each simulation path. The probabilities for the different parameter sets S_k are derived as Bayesian weights (see Burnham and Anderson (2002)):

$$P(S_k) = \frac{exp\left(-\frac{1}{2}\left(BIC(S_k) - \min_i\{BIC(S_i)\}\right)\right)}{\sum_{j=0}^{N-1} exp\left(-\frac{1}{2}\left(BIC(S_j) - \min_i\{BIC(S_i)\}\right)\right)},$$

where $BIC(S_k)$ denotes the Bayesian Information Criterion for the parameter set S_k. It can be determined from the likelihood of a specific trend process realization which is a byproduct of the pseudo maximum likelihood algorithm (see Schupp (2018) for details); the number of parameters in the penalty term of the BIC corresponds to the degrees of freedom of a trend process realization with k trend changes, i.e. $3 \cdot k + 3$.

2.3 Example Calibration

In this subsection, we present a model calibration for English and Welsh males. We use the entire data set which is available in the HMD (2018) for ages 60 to 109, i.e. from 1841 to 2016. In the subsequent sections, we use this example in order to illustrate different topics with respect to modeling the EMT and to present some concrete examples for the application of our model which underline the importance of differentiating between AMT and EMT.

Figure 2 shows the historic trend processes $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ and the best possible realizations for the actual trend processes $\hat{\kappa}_t^{(1)}$ and $\hat{\kappa}_t^{(2)}$ given different numbers of past trend changes k; the depicted k are those with Bayesian weights significantly different from zero. We observe that there might have been three, four, or five trend changes for $\kappa_t^{(1)}$ and five, six, or eight trend changes for $\kappa_t^{(2)}$.

Table 1 and Table 2 provide the estimates for $p^{(i)}$, $\mu^{(i)}$, $\sigma^{(i)}$, $\hat{\kappa}_{t_o}^{(i)}$, $\hat{d}_{t_0}^{(i)}$, and $\Sigma_{(i,i)}$ as well as the probabilities for all parameter sets. The trend change probability $p^{(i)}$ obviously increases with k. The mean of the trend change magnitude $\mu^{(i)}$, on the other hand, is rather constant. Combining these observations implies that the different parameter sets yield long-term mortality projections which differ significantly in terms of uncertainty. For $\kappa_t^{(2)}$, however, almost all probability is concentrated at $k = 5$, i.e. parameter uncertainty is rather small. For $\kappa_t^{(1)}$ most probability is assigned to the case $k = 4$, but there is also a significant probability for the case $k = 3$. Interestingly, the start values for the actual trend process, $\hat{d}_{t_0}^{(1)}$, differ significantly. While a trend change in 2011 is very likely, there is not yet enough data to be fairly sure (see also the difference between the orange and blue curves in the left panel of Figure 2). Thus, parameter uncertainty is significant here.[4]

Finally, Table 3 contains the covariance estimates between $\varepsilon_t^{(1)}$ and $\varepsilon_t^{(2)}$ for all combinations of parameter sets for $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$. They are derived from the correlations in the noise terms of the best possible trend process realizations and the variance estimates $\Sigma_{(1,1)}$ and $\Sigma_{(2,2)}$.

[4] We also tested the inclusion of additional trend changes after the most recent detected one for each k as done in Börger and Schupp (2018). However, the likelihood for such additional trend changes was negligible in any case.

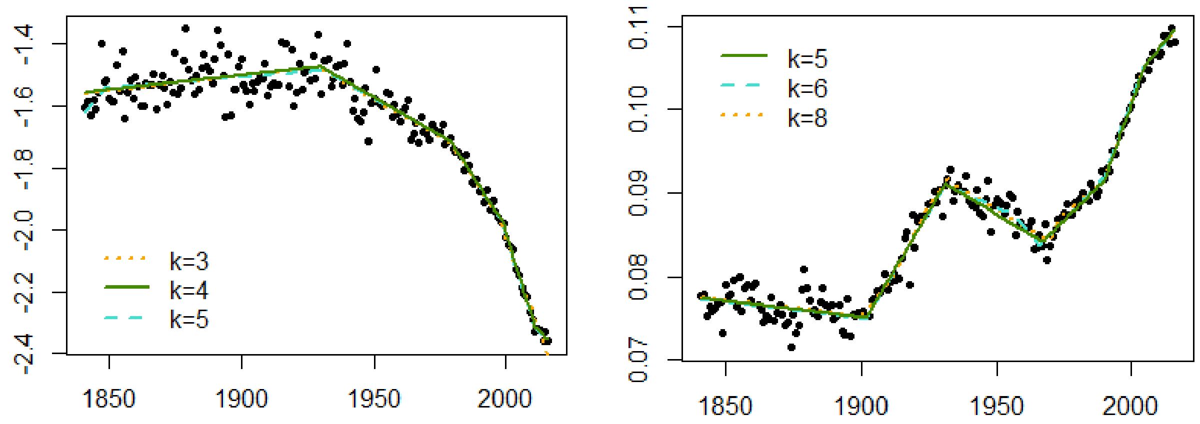

Figure 2: Historical trend processes $\kappa_t^{(1)}$ (left) and $\kappa_t^{(2)}$ (right) for English and Welsh males (dotted) and best possible realizations for the actual trend processes $\hat{\kappa}_t^{(1)}$ and $\hat{\kappa}_t^{(2)}$ given different numbers of trend changes k (solid lines)

k	$p^{(1)}$	$\mu^{(1)}$	$\sigma^{(1)}$	$\hat{\kappa}_{t_o}^{(1)}$	$\hat{d}_{t_0}^{(1)}$	$\Sigma_{(1,1)}$	$P(S_k)$
3	0.0170	−4.836	0.2398	−2.403	−0.0235	$6.62 \cdot 10^{-4}$	5.63%
4	0.0227	−4.520	0.4698	−2.356	−0.0084	$2.26 \cdot 10^{-4}$	94.00%
5	0.0284	−4.543	0.4387	−2.360	−0.0086	$2.63 \cdot 10^{-4}$	0.37%

Table 1: Empirical distribution for the AMT model parameters in $\hat{\kappa}_t^{(1)}$

k	$p^{(2)}$	$\mu^{(2)}$	$\sigma^{(2)}$	$\hat{\kappa}_{t_o}^{(2)}$	$\hat{d}_{t_0}^{(2)}$	$\Sigma_{(2,2)}$	$P(S_k)$
5	0.0284	−7.431	0.1355	0.1096	$3.95 \cdot 10^{-4}$	$4.64 \cdot 10^{-7}$	98.90%
6	0.0341	−7.526	0.3503	0.1096	$3.91 \cdot 10^{-4}$	$4.96 \cdot 10^{-7}$	1.05%
7	0.0398	−7.553	0.5053	0.1087	$3.85 \cdot 10^{-4}$	$4.46 \cdot 10^{-7}$	0.00%
8	0.0455	−7.659	0.4697	0.1095	$3.92 \cdot 10^{-4}$	$3.94 \cdot 10^{-7}$	0.05%

Table 2: Empirical distribution for the AMT model parameters in $\hat{\kappa}_t^{(2)}$

		$\kappa_t^{(2)}$		
	$\Sigma_{(1,2)}$	$k=5$	$k=6$	$k=8$
$\kappa_t^{(1)}$	$k=3$	$9.52 \cdot 10^{-6}$	$9.43 \cdot 10^{-6}$	$8.80 \cdot 10^{-6}$
	$k=4$	$7.06 \cdot 10^{-6}$	$8.00 \cdot 10^{-6}$	$7.16 \cdot 10^{-6}$
	$k=5$	$6.86 \cdot 10^{-6}$	$7.76 \cdot 10^{-6}$	$7.09 \cdot 10^{-6}$

Table 3: Covariance estimates between $\varepsilon_t^{(1)}$ and $\varepsilon_t^{(2)}$ for all combinations of parameter sets for $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$

This fully calibrated AMT model can be used to answer questions for which only the realized future mortality evolution is required, e.g. to simulate a portfolio run-off. In order to illustrate the adequacy of the mortality scenarios which the model generates, Figure 3 shows the median and mean projections as well as the 95% prediction intervals for the remaining period life expectancy of 65-year old English and Welsh males. We see that, in terms of AMT, the mean projection lies somewhere between the historical trends before and after 2011. This is in line with the empirical distribution for $\hat{d}_{t_0}^{(1)}$ which assigns substantial weight to both values for the AMT. The median projection yields slightly smaller life expectancies than the mean projection which illustrates the "upward" skewness implied by the empirical distribution for $\hat{d}_{t_0}^{(1)}$. The prediction intervals look highly plausible in the sense that the model output permits both, an ongoing steep increase in life expectancy as well as a leveling off in the near future. Comparisons of the presented trend process with alternative processes within the CBD model can be found in Börger and Schupp (2018).

3. EMT Model Component

In this section, we introduce the EMT component of our mortality model. As explained in the Introduction, the EMT is required in addition to the AMT in many practical applications, e.g. for the computation of risk capital over limited time horizons or for computing terminal payoffs of longevity hedging instruments with a fixed term. Therefore, we analyze and compare possible approaches for deriving the EMT, outline shortcomings of commonly used EMT approaches, and discuss the availability of information in order to derive the EMT at any point in time.

3.1 EMT Modeling Setup

Since the AMT is unobservable, any prediction of future mortality must build on some estimate for the AMT. For time t_0, one example for a (probabilistic) estimate has already been presented in the subsection on parameter uncertainty in the AMT model component. In many practical applications, however, an estimate for the AMT is required also at later points in time. Here, typically a point estimate is sufficient and we denote such deterministic EMTs at some point in time T by $\tilde{d}_T^{(i)}$.[5] An obvious choice for a deterministic EMT at time t_0 would be the mean of the empirical distribution for $\hat{d}_{t_0}^{(i)}$. For English and Welsh males at $t_0 = 2016$, this would be $\tilde{d}_{2016}^{(1)} = -0.00925$ and $\tilde{d}_{2016}^{(2)} = 0.000395$. In principle, it would be possible to determine the EMT in this way also within a simulation model at any future point in time. However, this would be fairly complex and time consuming since empirical distributions would have to be derived in every simulation path for every relevant point in time. This might not be feasible in a Monte Carlo simulation framework with thousands of simulation paths. Hence, more efficient alternatives need to be considered in practical applications.

[5] EMT-related quantities are denoted by $\tilde{}$ and AMT-related quantities by $\hat{}$ throughout the paper.

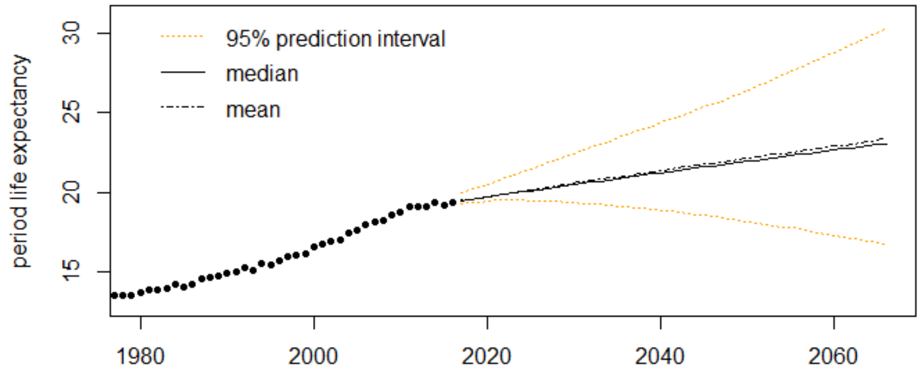

Figure 3: Historical and forecast period life expectancies for 65-year old males in England and Wales

As explained in Subsection 2.1, under our assumptions the prevailing AMT is always the (unobservable) best estimate for the future mortality trend. If we believe that the piecewise linear AMT describes reality in an adequate way, it is obvious to derive the EMT as the linear trend in the most recent available data.[6] We note that in practice, the $\kappa_t^{(i)}$ processes are not directly observable. However, by fitting the CBD model structure to raw mortality rates up to time T, an observer would obtain the same $\kappa_t^{(i)}, t \leq T$ as projected in the AMT simulation. Thus, the $\kappa_t^{(i)}$ can be derived from observable data. Then the EMT can be derived from the $\kappa_t^{(i)}$ which are directly available from the AMT simulation up to T.[7] This makes our modeling framework very efficient. In other mortality models, this might not be the case. For instance, in the model of Lee and Carter (1992), the observer would estimate a process $\kappa_t, t \leq T$ which differs from the simulated one due to interactions between the trend process and the age dependent parameters. Here, additional information in form of new data impacts κ_t estimates for previous years. Thus, the κ_t process could not be assumed observable, and the EMT would need to be determined from the observed raw mortality rates.

The common approach for estimating a mortality trend is to determine the linear trend within a time series of data whose length is usually fixed based on data availability constraints, expert judgement, or irregularities in the data. At the same time, all data points usually have the same impact on the trend estimation. However, this could be problematic if mortality is assumed to follow a piecewise linear trend with changing slope. In this case, the most recent data points are obviously more informative with respect to the prevailing AMT than data points further in the past. In fact, if a trend change occurred in the recent past, all data points before this trend change even blur the EMT derivation instead of supporting it. The EMT would react too slowly to potential trend changes and could thus be far off the AMT. Consequently, the most recent data points should have more weight in the trend estimation. If, however, too much weight was given to just a few data points, the EMT could pick up trend changes in the AMT rather quickly,

[6] Note that, by this approach, the (unknown) level of the mortality process, $\hat{\kappa}_T^{(i)}$, is estimated as a byproduct.

[7] When deriving the EMT in practice, the $\kappa_t^{(i)}$ would typically be available only up to a few years before T as data collection and preparation takes its time. However, this does not affect the general concept presented in this paper, and for simplicity we therefore assume that mortality data is available directly.

but this would increase the risk of wrongly identifying noise as a trend change. We will discuss the question of a suitable weighting in the next subsection. Here, the assumed dynamics for the AMT will serve as a basis to derive a reasonable trade-off, i.e. some kind of optimal weighting, based on the available information on the volatility of the noise and how likely trend changes of which intensity are. In order to be consistent with the AMT model component, the weighting should be fixed such that the EMT is as close as possible to the AMT "on average".

3.2 Comparison of EMT Weight Specifications

In order to analyze the impact of weighting, we compare different weight specifications for our example population of English and Welsh males. The base case is "constant weights" for an estimation period whose length is to be optimized. When deriving an EMT at time T, the weight for the data point in year $t \leq T$ is

$$w_{const}(t,T) = \begin{cases} 1 & , if\ T - h_{const} < t \leq T \\ 0 & , if\ t \leq T - h_{const} \end{cases},$$

where h_{const} is the "weighting parameter", i.e. the length of the estimation period in this case. As alternatives we consider weights which decay linearly or exponentially going backward in time:

$$w_{lin}(t,T) = \max\left(0; 1 - \frac{1}{h_{lin}}(T-t)\right), t \leq T,$$

$$w_{exp}(t,T) = \frac{1}{\left(1 + {}^{1}/_{h_{exp}}\right)^{T-t}}, t \leq T.$$

Here, the weighting parameters h_{lin} and h_{exp} determine the speed of the linear or exponential decay, respectively.[8] Note that the weights $w_{exp}(t,T)$ have already been proposed and used in a similar context by (Börger, et al., 2014).

In order to optimize the weighting parameters, we consider two different criteria and compare the results. In the first approach, we determine the weighting parameters such that the EMTs $\tilde{d}_T^{(i)}$, estimated from $\kappa_{1841}^{(i)}, \dots, \kappa_T^{(i)}$, are as close as possible in terms of mean squared error to the AMTs $\hat{d}_T^{(i)}$, $i = 1, 2$. In the second approach, we minimize the mean squared error between the remaining cohort life expectancy of a 65-year old based on the EMT at time T, $\tilde{e}_{65,T}$, and the actual cohort life expectancy. We denote the latter by $\hat{e}_{65,t_\omega}$ and derive it from an AMT simulation from T to t_ω, where t_ω is the point in time when the cohort has died out. Here the weighting parameters for $\tilde{d}_T^{(1)}$ and $\tilde{d}_T^{(2)}$ are determined simultaneously and also the estimates for $\hat{\kappa}_T^{(i)}$ impact the weighting parameter calibrations. We consider the cohort life expectancy here since it is highly dependent on the future mortality trend – in contrast to, e.g., the period life expectancy – and is similar in structure to present values of annuities which are of high

[8] In the case of exponential weighting, the weights are different from zero for all t. Nevertheless, a limited estimation period of sufficient length can be considered for practical purposes as the impact of data points far in the past is negligible due to the fast decay.

relevance to actuaries. Thus, in total we compare six different sets of weighting parameter estimates (two optimization criteria x three weight specifications).

In order to ensure that our parameter estimates are not biased by the fixed historical $\kappa_t^{(i)}, t \leq t_0$, we first simulate the AMT model 40 years into the future and optimize the weighting parameters based on the EMTs at $T = t_0 + 40$. The mean squared errors (mse's) for each set of weighting parameter estimates are computed from 100,000 simulation paths.

Table 4 shows the optimal weighting parameters and the mse's for both EMTs $\tilde{d}_T^{(1)}$ and $\tilde{d}_T^{(2)}$ based on the first optimization criterion (EMT close to AMT). Figure 4 illustrates the resulting weights where the solid lines correspond to $\tilde{d}_T^{(1)}$ and the dashed lines to $\tilde{d}_T^{(2)}$. We observe that the weighting parameters for $\tilde{d}_T^{(2)}$ are generally smaller than for $\tilde{d}_T^{(1)}$. The main reason is that the trend change probability for $\hat{\kappa}_T^{(2)}$ is on average larger than for $\hat{\kappa}_T^{(1)}$, and therefore, $\tilde{d}_T^{(2)}$ needs to react to trend changes more often. Furthermore, we see that the linear and exponential weights outperform the constant weights in terms of mse. This is as expected since in the case of constant weights (i.e. the case of no weighting within an estimation period of optimal length) older data points which possibly lie before the most recent trend change materially bias the EMT derivation. This observation further supports the idea that weighting should be considered in general when deriving mortality trends from historical data. The mse differences between linear and exponential weights are rather small in this example of English and Welsh males, with linear weights having the slight edge. However, we have found that for other populations the exponential weights perform better. This particularly depends on the likelihood of trend changes, their magnitudes, and the volatilities of the noise terms in the $\kappa_t^{(i)}$ processes. More precisely, the smaller the noise's volatility and the larger the trend change probability and magnitude, the better the exponential weights perform with their more pronounced focus on the most recent data points.

	weights	constant	linear	exponential
$\tilde{d}_T^{(1)}$	$h^{(1)}$	9.0	10.5	2.1
	mse	$1.45 \cdot 10^{-5}$	$1.31 \cdot 10^{-5}$	$1.35 \cdot 10^{-5}$
$\tilde{d}_T^{(2)}$	$h^{(2)}$	8.0	9.5	1.8
	mse	$3.19 \cdot 10^{-8}$	$2.88 \cdot 10^{-8}$	$2.95 \cdot 10^{-8}$

Table 4: Optimal weighting parameters and mse's for both EMTs $\tilde{d}_T^{(1)}$ and $\tilde{d}_T^{(2)}$ based on the optimization criterion "EMT close to AMT"

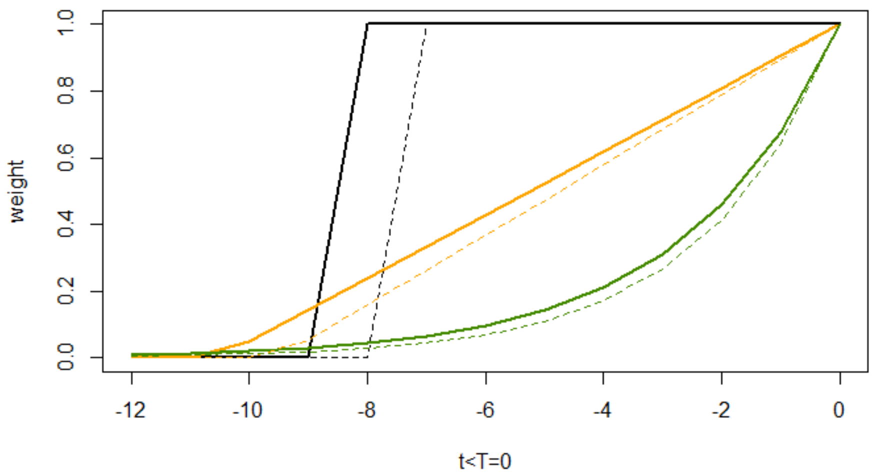

Figure 4: Optimal weights for both EMTs $\tilde{d}_T^{(1)}$ (solid lines) and $\tilde{d}_T^{(2)}$ (dashed lines) based on the optimization criterion "EMT close to AMT"

Table 5 contains the optimal weighting parameters with their corresponding mse's in case of the second optimization criterion ($\tilde{e}_{65,T}$ based on EMT close to actual life expectancy $\hat{e}_{65,t_\omega}$). It also gives probabilities that the life expectancy estimates $\tilde{e}_{65,T}$ are smaller (larger) than 95% (105%) of the actual life expectancies $\hat{e}_{65,t_\omega}$. Moreover, for comparison we have added results for a fixed estimation period of 30 years (i.e. constant weights within a predetermined 30-year estimation period) as an example for an estimation period which is typically used in practice.

First of all we observe that the optimal weighting parameters in Table 5 are very similar to those in Table 4. Thus, the concrete optimization criterion has only a small impact on the weights. Furthermore, the linear weights again outperform the alternatives in terms of mse. However, as also indicated by the misestimation probabilities, the differences between the three weighting approaches are fairly small. Results for the fixed estimation period of 30 years, on the other hand, are considerably worse. Compared to the case of an estimation period with optimal length (constant weights approach), the mse is about 91% and the combined misestimation probability 10.1 percentage points (i.e. 26.1% – 16.0%) larger. Clearly, available information on the dynamics of actual mortality should be taken into account when deriving best estimate mortality projections. The estimation period should be optimized based on this information, independent of whether weighting is applied or not.

weights	constant	linear	exponential	fixed period
$h_{\cdot}^{(1)}$	10.0	11.0	2.2	30
$h_{\cdot}^{(2)}$	8.0	10.5	2.1	30
mse	1.59	1.56	1.58	3.04
$P(\tilde{e}_{65,T} < 95\% \cdot \hat{e}_{65,t_\omega})$	8.7%	8.5%	8.7%	15.0%
$P(\tilde{e}_{65,T} > 105\% \cdot \hat{e}_{65,t_\omega})$	7.3%	7.1%	7.0%	11.1%

Table 5: Optimal weighting parameters and mse's for both EMTs $\tilde{d}_T^{(1)}$ and $\tilde{d}_T^{(2)}$ based on the cohort life expectancy optimization criterion

We would like to stress that the optimal weighting parameters were derived assuming a concrete specification for the dynamics of the true (unknown) AMT. If the real dynamics of future mortality deviate from this assumption, other parameters might be optimal. We consider this as an aspect of model risk and will not further deal with this issue in this paper.

Coming back to our example of English and Welsh males aged 65 and older, the EMTs $\tilde{d}_T^{(1)}$ and $\tilde{d}_T^{(2)}$ at future points in time T should be derived using linear weights; the weighting parameters should be in the ranges as given in Table 4 and Table 5.

4. Distinction between AMT and EMT

In this section, we use several practical examples to illustrate how important it is to differentiate between AMT and EMT. At the same time, we use numerical examples to explain how the combined AMT/EMT model from the previous sections can be applied in practice and how risk is misestimated if the difference between AMT and EMT is ignored. For the EMTs, we always use linear weights with weighting parameters as given in Table 5. Without loss of generality, we assume that annuitants'/pensioners' mortality rates are exactly as for males in England and Wales and that the number of annuitants/pensioners is large enough such that there is no unsystematic mortality risk. Furthermore, in oder to focus on effects resulting from stochastic mortality, we assume throughout this section that interest rates are deterministic and constant at 2%.

4.1 Example 1: Hedge Effectiveness of a Value Hedge

In this first example, we consider a pension fund whose members are all aged 45 at t_0. Some hedge provider has offered a value hedge agreement in order to mitigate the risk of the pension fund being underfunded at the members' retirement age of 65. Thus, the hedge agreement has a term of 20 years and pays the pension fund the best estimate liabilities according to up-to-date mortality assumptions at $T = t_0 + 20$. We assume that the hedge provider and the pension fund's trustees have agreed on deriving mortality assumptions at time T based on the EMT with optimal weighting (see Subsection 3.2). The premium for this hedge, payable at expiry, is the pension fund's expected liabilities as estimated at t_0 plus some hedging fee. Thus effectively, the hedge provider will fill up the pension fund's liabilities at expiry of the hedge if required and will receive money if the pension fund's liabilities are smaller than expected. While the hedge agreement transfers the longevity risk up to time T (including the reserving risk at time T), the pension fund will still be exposed to the risk that mortality after time T might not evolve as expected. Thus the remaining risk corresponds to the difference between the value of the actual pension payouts and its estimate at time T, i.e. the payoff of the value hedge. This is the risk we consider in more detail. In fact, it consists of two components: the risk that the mortality assumptions at time T are inadequate, i.e. the risk that the EMT differs from the AMT, and the risk that the AMT might change during the pension payout phase.

Assume for a moment that, when assessing this risk at t_0, the pension fund's trustees do not distinguish between AMT and EMT. Instead, they (wrongfully) assume that the AMT which they simulate up to time T is observable and compute the hedge payoff according to this AMT. Then for the trustees, the risk of the pension fund appears to lie in the following difference of present values (PV) at time T:

$$PV_T\left(pension\,payouts \mid AMT_{t_\omega}\right) - PV_T(pension\,payouts \mid AMT_T), \tag{1}$$

where t_ω denotes the year when every pensioner has died and AMT_{t_ω} the realized AMT up to that year. However, in this case the trustees would neglect the first component of the risk the pension fund is exposed to, i.e. the risk of inadequate mortality assumptions at time T based on which the hedge payoff is derived. The actual risk the trustees are facing is rather given by

$$PV_T\left(pension\,payouts \mid AMT_{t_\omega}\right) - PV_T(pension\,payouts \mid EMT_T). \tag{2}$$

Figure 5 shows histograms for both differences based on 100,000 simulation paths and for 1£ of annual pension payout. Obviously, the distribution for the risk described by equation (2) is significantly wider. In fact, its variance is 0.24 compared to only 0.15 in the case where the AMT is assumed to be observable. Thus, given the risk without a hedge, i.e. $Var\left(PV_T\left(pension\,payouts \mid AMT_{t_\omega}\right)\right) = 1.90$, the trustees would assume a hedge effectiveness of 1–0.15/1.90 = 92.1%. The true hedge effectiveness however is only 87.2%.[9] Thus, ignoring the fact that the AMT is not observable makes the hedge appear better than it actually is.

[9] Note that we implicitly assume here that actual mortality evolves according to the dynamics in the AMT model component. Thus, we do not consider any model risk in the AMT model. Taking this model risk into account would reduce the hedge effectiveness in both cases, but the essence of this example would be unaffected.

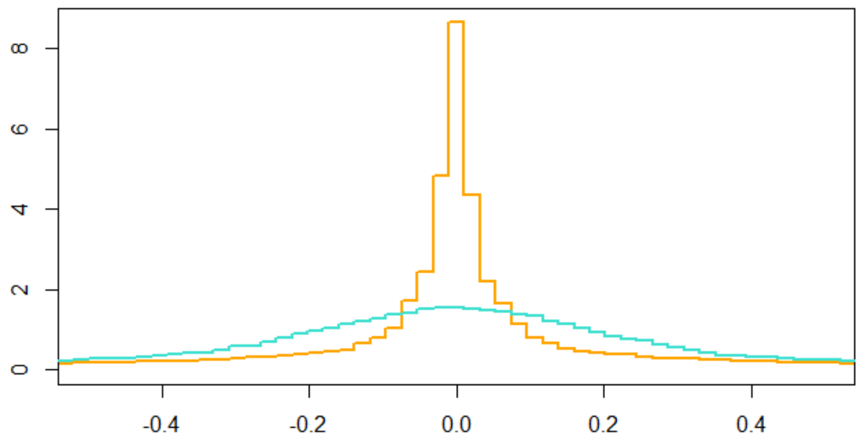

Figure 5: Histograms for differences in Equation (1) (yellow line) and Equation (2) (blue line) for 1£ of annual pension payout

4.2 Example 2: Safety Margins in Annuity Conversion Rates

In a second example, we now consider an insurer with a portfolio of unit-linked deferred annuities. We again assume that all policy holders are aged 45 at t_0. At the end of the accumulation phase (at $T = t_0 + 20$), the fund value is converted into a life-long annuity for which we assume a technical interest rate of 2%. However, the insurer does not want to convert at the actuarially fair rate, but intends to fix the rate such that the probability for losses from increasing longevity during the payout phase amounts to 1%. The surplus which arises in the 99% other cases may be (partially) credited to the policy holders as discretionary benefits. More for the sake of simplicity, we assume that the insurer is trying to achieve this by implementing a safety margin in form of a fixed percentage reduction of the actuarially fair conversion rate.

The risk of the insurer consists of the same two components as in the first example: the EMT might be an imprecise estimate of the AMT and actual longevity might increase after time T. Obviously, both components need to be taken into account when deriving the conversion rate reduction. Thus, the insurer is interested in the 99th percentile of the distribution in Equation (2) (with annuity payouts instead of pension payouts) and printed in blue in Figure 5. He can then compute the percentage reduction of the actuarially fair annuity conversion rate as the ratio of this percentile and the average of the fair annuity conversion rates in the 100,000 simulation paths, i.e. $PV_T(annuity\,payouts \mid EMT_T) = 17.58$. With the 99th percentile being equal to 1.68, the percentage reduction is then 9.6%.

However, if the insurer wrongfully assumes the AMT to be observable, he will determine the 99th percentile of the distribution in Equation (1) (printed in yellow in Figure 5). In this case the percentage reduction of $PV_T(annuity\,payouts \mid AMT_T)$, which the insurer would assume to be the actuarially fair annuity conversion rate, would be only 7.8%. Thus, the insurer would underestimate his risk significantly. In fact, the probability of suffering losses during the annuity payout phase would be 1.7% instead of the intended 1%.

4.3 Example 3: SCR for Longevity Risk

In a final example, we consider an insurer who wants to develop an internal model for Solvency II. Thus the insurer needs to determine the SCR for longevity risk (as any SCR) as the 99.5% Value-at-Risk of the basic own funds (which is essentially the difference between assets and liabilities) over a 1-year time horizon. In this setting, longevity risk consists of two components: the risk of less annuitants than expected dying during the next year and the risk of an unfavorable change in longevity assumptions for the time beyond. Typically, the second component is more relevant.

For simplicity, we assume that only the liabilities are exposed to longevity risk (no hedge instruments at the asset side) and that there is no loss-absorbing capacity of technical provisions. Following Börger (2010), the insurer can then compute the SCR at time t as the 99.5th percentile of the change in liabilities ΔBEL_{t+1} from time t to $t+1$ due to longevity:

$$\Delta BEL_{t+1} = (BEL_{t+1} + CF_{t+1}) \cdot \frac{1}{1+r} - BEL_t,$$

where r is the risk-free interest rate (2% in our case), CF_{t+1} denotes the cashflows (benefits plus expenses minus premiums) of the longevity prone business between t and $t+1$, and BEL_t is the best estimate liabilities at time t. In order to evaluate the distribution of ΔBEL_{t+1}, the insurer needs a combined AMT/EMT model. More precisely, the EMT component is needed to compute the best estimate liabilities and the AMT component to simulate the realized mortality evolution over the 1-year time horizon.

As mentioned above, the change in best estimate liabilities and thus the change in the EMT over one year essentially determines the SCR. Therefore, it is important to ensure that the assumed variability in the EMT over time is realistic. If the impact of the additional data point (for the one year) on the EMT update is too small (e.g. in case of a long estimation period), annual changes of the EMT will likely be underestimated. If the impact of the additional data point is too large, annual EMT changes will be overestimated. Here, the findings in Subsection 3.2 can be very helpful. The optimal weights give a clear recommendation how large the impact of new data points should be and thus by how much the EMT might change over one year.

Now assume that – as part of the internal model validation process – the insurer wants to analyze how large the SCR at some point in time T can be, depending on the (historical) mortality evolution up to that point in time. Besides the new data point in year $T+1$ and the weighting in the EMT derivation, the structure (in particular potential trend changes) in the mortality data up to time T is the main determinant for EMT changes over one year. As in the second example, the insurer sets $T = t_0 + 20$. He simulates 10,000 outer paths up to time T in order to average over virtually all potential mortality evolutions. Then for each of these outer paths, 10,000 inner 1-year paths are simulated in order to empirically derive the 99.5th percentile of the change in liabilities ΔBEL_{T+1}. Exemplarily for a portfolio of annuitants, he considers the cohort of 75-year olds who receive combined annual payments of 1£ at time T. For simplicity, there are no expenses and premiums.

The blue curve in Figure 6 is the density of the distribution for the SCR at time T depending on the mortality evolution up to T. On average, the SCR is 0.40 with a (sample) variance of 0.0064 and (sample) standard deviation of 0.08. The considerable coefficient of variation of about 20% shows that the SCR can vary significantly due to the variability in the EMT changes over one year.

Once again, we want to illustrate how important it is to distinguish between AMT and EMT. If the insurer wrongfully assumes the AMT to be observable and computes the best estimate liabilities based on the AMT, he will obtain the distribution given in yellow in Figure 6. On average, the SCR would be 0.73 and thus significantly larger than in the EMT case. Clearly, longevity risk would be overstated substantially. The reason is that the AMT is very likely to exhibit massive trend changes in the extreme 1-year scenarios which are relevant for the SCR. Annual changes in the EMT, on the other hand, are not that strong as the EMT does not pick up trend changes instantly. Due to the weighted linear regression, trend changes are only taken into account over several years.

While assuming the AMT to be observable implies an overestimation of risk in this example, it was the opposite in the other examples. The fact that misestimations can occur in both directions once again underlines the necessity to clearly distinguish between AMT and EMT.

5. Conclusions

In this paper, we have explained that for virtually all questions that require stochastic mortality modeling, a clear distinction between actual mortality and estimated future mortality is required. In particular, for models with a stochastic mortality trend, at any point in time the actual mortality trend that is being modeled and the estimated mortality trend that an observer would estimate based on the development up to that point in time are different.

We have specified a concrete model that simultaneously and consistently projects AMT and EMT. In line with historic observations, the AMT is modeled as a piecewise linear trend with random changes in slope. Since the AMT at the start of a simulation is unknown, parameter uncertainty needs to be taken into account, and we show how this can be done. Given the assumed dynamics for the AMT, the EMT is best derived by a weighted linear regression on the most recent observable data. In our modeling framework, this can be done on the output of the AMT model component directly, which makes the framework very efficient in Monte Carlo simulations.

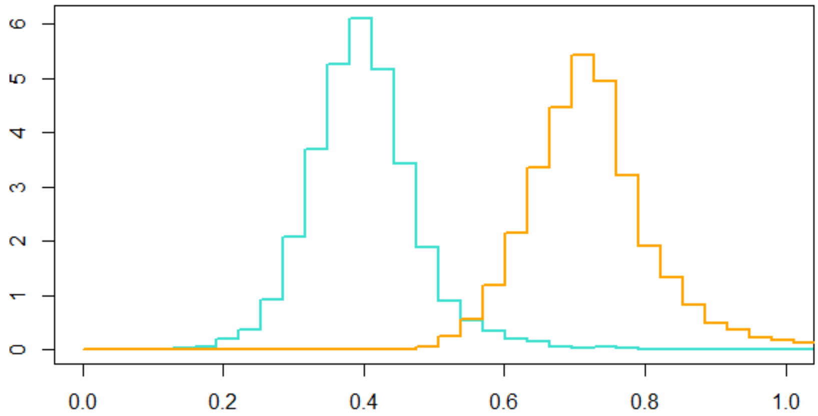

Figure 6: Histograms for the SCR at time T depending on the mortality evolution up to T under the assumption that the AMT is observable (yellow line) or unobservable (blue line)

Finally, we have considered several practical examples to show how our model can be applied and to quantify the error that results if the difference between AMT and EMT is ignored: If the AMT is wrongfully assumed observable, risk is significantly misestimated in all our examples – sometimes underestimated, sometimes overestimated. Therefore, this research should be of interest to anybody who is concerned with mortality projections or longevity risk. This particularly holds for public pension schemes, pension funds, and insurers, but also for regulators and auditors.

References

Bauer, D., Börger, M., Russ, J., & Zwiesler, H. (2008). The Volatility of Mortality. Asia-Pacific Journal of Risk and Insurance 3, 184-211.

Bauer, D., Börger, M., & Russ, J. (2010). On the Pricing of Longevity-Linked Securities. Insurance: Mathematics and Economics 46(1), 139-149.

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., Fleischer, D., & Kuksin, N. (2014). Modeling the Mortality Trend under Modern Solvency Regimes. ASTIN Bulletin 44(1), 1-38.

Börger, M., & Schupp, J: (2018). Modeling Trend Processes in Parametric Mortality Models. Insurance: Mathematics and Economics 78, 369-380.

Burnham, K. P., & Anderson, D. R. (2002). Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach. Springer Science & Business Media.

Cairns, A. (2013). Robust Hedging of Longevity Risk. Journal of Risk and Insurance 80(3), 621-648.

Cairns, A., Blake, D., & Dowd, K. (2006). A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration. Journal of Risk and Insurance 73(4), 687-718.

Cairns, A., Dowd, K., Blake, D., & Coughlan, G. (2014). Longevity Hedge Effectiveness: A Decomposition. Quantitative Finance 14(2), 217-235.

Cairns, A. J., & El Boukfaoui, G. (2018). Basis risk in index based longevity hedges: A guide for longevity hedgers. North American Actuarial Journal (forthcoming).

Human Mortality Database (2018) University of California, Berkeley, USA, and Max Planck Institute for Demographic Research, Germany. Available at: www.mortality.org (Data downloaded on 01/11/2018).

Hunt, A., & Blake, D. (2015). Forward Mortality Rates in Discrete Time I: Calibration and Securities Pricing. Working Paper, Cass Business School.

Hunt, A., & Blake, D. (2016). Forward mortality Rates in Discrete Time II: Longevity Risk and Hedging Strategies. Working Paper, Cass Business School.

Lee, R., & Carter, L. (1992). Modeling and Forecasting U.S. Mortality. Journal of the American Statistical Association, Volume 87(419), 659-671.

Plat, R. (2011). One-year Value-at-Risk for Longevity and Mortality. Insurance: Mathematics and Economics 49(3), 462-470.

Richards, S. J., Currie, I. D., & Ritchie, G. P. (2014). A Value-at-Risk framework for longevity trend risk. British Actuarial Journal 19(1), 116-139.

Schupp, J. (2018). On the Modeling of variable Mortality Trend Processes. Working Paper, Ulm University.

Zhu, N., & Bauer, D. (2011). Coherent Modeling of the Risk in Mortality Projections: A Semi-Parametric Approach. Tech. Rep. 678, Georgia State University, Atlanta.