Deterministic Shock vs. Stochastic Value-at-Risk – An Analysis of the Solvency II Standard Model Approach to Longevity Risk*

Matthias Börger

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

Helmholtzstraße 22, 89081 Ulm, Germany

Phone: +49 731 50 31257. Fax: +49 731 50 31239

Email: m.boerger@ifa-ulm.de

September 2010

Abstract

In general, the capital requirement under Solvency II is determined as the 99.5% Value-at-Risk of the Available Capital. In the standard model's longevity risk module, this Value-at-Risk is approximated by the change in Net Asset Value due to a pre-specified longevity shock which assumes a 25% reduction of mortality rates for all ages.

We analyze the adequacy of this shock by comparing the resulting capital requirement to the Value-at-Risk based on a stochastic mortality model. This comparison reveals structural shortcomings of the 25% shock and therefore, we propose a modified longevity shock for the Solvency II standard model.

We also discuss the properties of different Risk Margin approximations and find that they can yield significantly different values. Moreover, we explain how the Risk Margin may relate to market prices for longevity risk and, based on this relation, we comment on the calibration of the cost of capital rate and make inferences on prices for longevity derivatives.

*The author is very grateful to Andreas Reuß, Jochen Ruß, Hans-Joachim Zwiesler, and Richard Plat for their valuable comments and support, and to the Continuous Mortality Investigation for the provision of data.

1 Introduction

As part of the Solvency II project, the capital requirements for European insurance companies will be revised in the near future. The main goal of the new solvency regime is a more realistic modeling and assessment of all types of risk insurance companies are exposed to. In principle, the Solvency Capital Requirement (SCR) will then be determined as the 99.5% Value-at-Risk (VaR) of the Available Capital over a 1-year time horizon, i.e., the capital required today to cover all losses which may occur over the next year with at least 99.5% probability. For a detailed overview and a discussion of the Solvency II proposals, we refer to Eling et al. (2007), Steffen (2008), and Doff (2008). A comparison with other solvency regimes can be found in Holzmüller (2009).

Insurance companies are encouraged to implement (stochastic) internal models to assess their risks as accurately as possible. However, the implementation of such internal models is rather costly and sophisticated. Therefore, the European Commission with support of the Committee of Insurance and Occupational Pension Supervisors (CEIOPS) has established a scenario based standard model which all insurance companies will be allowed to use in order to approximate their capital requirements. In this model, the overall risk is split into several modules, e.g. for market risk, operational risk, or life underwriting risk, and submodules for which separate SCRs are computed. These SCRs are then aggregated under the assumption of a multivariate normal distribution with pre-specified correlation matrices to allow for diversification effects. The setup and calibration of the standard model is currently being established by a series of Quantitative Impact Studies (QIS) in which the effects of the new capital requirements are analyzed. Even though this standard model certainly has some shortcomings (for critical discussions see, e.g., Doff (2008) or Devineau and Loisel (2009)), most small and medium-size insurance companies are expected to rely on this model. But also larger companies are likely to adopt at least a few modules for their (partial) internal models. Hence, a reasonable setup and calibration of the standard model is crucial in order to ensure the financial stability of the European insurance markets.

One prominent risk annuity providers and pension funds are particularly exposed to is longevity risk, i.e. the risk that insured on average survive longer than expected. The importance of this risk is very likely to increase even further in the future: A general decrease in benefits from public pay as you go pension schemes in most industrialized countries, often in combination with tax incentives for private annuitization, will almost certainly lead to a further increasing demand for annuity products. Moreover, longevity risk constitutes a systematic risk as it cannot be diversified away in a large insurance portfolio. It is also non-hedgeable since, currently, no liquid and deep market exists for the securitization of this risk.

In the Solvency II standard model, longevity risk is explicitly accounted for as part of the life underwriting risk module. The SCR is, in principle, computed as the change in liabilities due to a longevity shock that assumes a permanent reduction of mortality rates. Up to QIS4, this reduction was set to 25% and was mainly based on what insurance companies in the United Kingdom (UK) in 2004 regarded as consistent with the general 99.5% VaR concept of Solvency II (cf. CEIOPS (2007)). However, the shock's magnitude has been widely discussed over the last years. For instance, some participants of QIS4 regard a shock of 25% as very high (cf. CEIOPS (2008c)) whereas CEIOPS (2010) is still convinced that a reduction of mortality rates by 25% is adequate. In CEIOPS favor, one could argue that the shock should be rather large as the standard model is meant to be conservative in order to give incentives for the implementation of internal models. Nevertheless, only very recently, the European Commission (2010) has decided to reduce the longevity stress to 20% for QIS5 without giving detailed reasoning for this recalibration. The computations in this paper are still based on the value of 25% but, as we mainly focus on the structure of the shock as opposed to its magnitude, we are convinced that all major findings remain valid also for a longevity stress of only 20%.

Besides the magnitude, also the structure of the longevity stress has been the topic of ongoing discussions. A reason given by QIS4 participants for not applying the standard model is that "the form of the longevity stress within the SCR standard formula does not appropriately reflect the actual longevity risk, specifically it does not appropriately allow for the risk of increases in future mortality improvements" (CEIOPS (2008c)). CEIOPS (2009a, 2010) acknowledges the feedback that a gradual change in mortality rates may be more appropriate than a one-off shock and also analyzes a possible longevity stress dependency on age and duration but finally decides to stick to the equal one-off shock for all ages and durations. A consequence of this decision could be an unnecessarily high capitalization of insurance companies in case longevity risk is overestimated by the current longevity shock. In the converse case, a company's default risk could be significantly higher than the accepted level of 0.5%.

Therefore, a thorough analysis is required whether or not the change in liabilities due to a 25% longevity shock constitutes a reasonable approximation of the 99.5% VaR of the Available Capital. This analysis is carried out in this paper. We examine the adequacy of the shock's structure, i.e. an equal shock for all ages and maturities, and its calibration. Moreover, we explain how it could possibly be improved. In the second part of the paper, we discuss the Risk Margin under Solvency II for the case of longevity risk. Due to its complexity, different approximations have been proposed whose properties and performances are yet rather unclear. We examine these approximations, in particular the assumption of a constant ratio of SCRs and liabilities over time which has turned out to be very popular in practice. Moreover, we contribute to the ongoing debate on the calibration of the cost of capital rate by relating the Risk Margin to (hypothetical) market prices for longevity risk. This approach offers a new perspective compared to the shareholder return models which have been applied in the calibration so far (cf. CEIOPS (2009b)). Using the same relation and assuming the cost of capital rate to be fixed, we can finally make inferences on the pricing of longevity derivates.

The impact and the significance of longevity risk on annuity or pension portfolios have already been analyzed by several authors. However, not all of their work directly relates to capital requirements under a particular solvency regime. For instance, Plat (2009) focuses on the additional longevity risk pension funds can be exposed to due to fund specific mortality compared to general population mortality. Stevens et al. (2010) compare the risks inherent in pension funds with different product designs, and Dowd et al. (2006) measure the remaining longevity risk in an annuity portfolio in case of imperfect hedges by various longevity bonds. Others, like Hari et al. (2008a), Olivieri and Pitacco (2008a,b), and Olivieri (2009), do analyze capital requirements for certain portfolios but consider approaches which differ from the 1-year 99.5% VaR concept of Solvency II, e.g. they assume longer time horizons and different default probabilities. Therefore, no final conclusion regarding the standard model approach is possible in their setting. Nevertheless, Olivieri and Pitacco (2008a,b) come to the conclusion that the "shock scenario referred to by the standard formula can be far away from the actual experience of the insurer, and thus may lead to a biased allocation of capital". They find that the standard formula contains some strong simplifications and argue that internal models should be adopted instead. At least, the magnitude of mortality reductions should be recalibrated to capital requirements derived from an exemplary internal model.

The remainder of this paper is organized as follows: In Section 2, the relevant quantities, i.e. the Available Capital, the Risk Margin, and the Solvency Capital Requirement are defined. For the latter, different definitions based on the 25% longevity shock (the shock approach) and based on the VaR of the Available Capital (the VaR approach) are given. Moreover, the practical computation of these SCRs is described. Since the computation in the VaR approach requires stochastic modeling of mortality, in Section 3, we discuss the suitability of different mortality models and explain why we decided to use the forward model of Bauer et al. (2008, 2010). Subsequently, we introduce the forward mortality modeling framework, in which this model is specified. The full specification of the model and an improved calibration algorithm are presented in the appendix. In Section 4, we establish the setting in which the SCRs from both longevity stress approaches are to be compared. In particular, assumptions and simplifications are introduced which are necessary to exclude risks other than longevity risk and to ensure that a direct comparison of SCRs based on a one-off shock approach and a gradual change in mortality as in the model of Bauer et al. (2008, 2010) is possible. The comparison is then performed in Section 5. In the process, various term structures, ages, levels of mortality rates, and portfolios of contracts are considered. Based on these comparisons, a modified longevity stress for the standard model is proposed in Section 6, which is still scenario based but allows for the shock magnitude to depend on age and maturity. In Section 7, the Risk Margins based on the 25% shock and the modified longevity stress are compared and the properties of different Risk Margin approximations are analyzed. Moreover, we discuss the adequacy of the current cost of capital rate and argue how the pricing of longevity derivatives may relate to the Solvency II capital requirements and the Risk Margin in particular. Finally, Section 8 concludes.

2 The Solvency Capital Requirement under Solvency II

2.1 General definitions

Intuitively, the Solvency Capital Requirement (SCR) under Solvency II is defined as the amount of capital necessary at time $t = 0$ to cover all losses which may occur until $t = 1$ with a probability of at least 99.5%. However, in order to give a precise definition of the SCR, we first need to introduce the notion of the Available Capital.

By definition, the Available Capital at time t, AC_t, is the difference between market value of assets and market value of liabilities at time t. Thus, it is a measure of the amount of capital which is available to cover future losses. In general, the market value of a company's assets can be derived rather easily: Either asset prices are directly obtainable from the financial market (mark-to-market) or the assets can be valued by well established standard methods (mark-to-model). The market value of liabilities, however, is difficult to determine. There is no liquid market for such liabilities and due to options and guarantees embedded in insurance contracts, the structure of the liabilities is typically rather complex such that standard models for asset valuation cannot be applied directly. Therefore, under Solvency II a market value of liabilities is approximated by the so-called Technical Provisions which consist of the Best Estimate Liabilities (BEL) and a Risk Margin (RM).

The Risk Margin can be interpreted as a loading for non-hedgeable risk and has to "ensure that the value of technical provisions is equivalent to the amount that (re)insurance undertakings would be expected to require to take over and meet the (re)insurance obligations" (CEIOPS (2008a)). Thus, in case of a company's insolvency, the Risk Margin should be large enough for another company to guarantee the proper run-off of the portfolio of contracts. It is computed via a cost of capital approach (cf. CEIOPS (2008a)) and reflects the required return in excess of the risk-free return on assets backing future SCRs. Hence, the Risk Margin RM can be defined as (cf. CEIOPS (2009b))

$$RM := \sum_{t \geq 0} \frac{CoC \cdot SCR_t}{(1 + i_{t+1})^{t+1}}, \tag{1}$$

where SCR_t is the SCR at time t, i_t is the annual risk-free interest rate at time zero for maturity t, and CoC is the cost of capital rate, i.e. the required return in excess of the risk-free return.

The SCR is defined as the 99.5% VaR of the Available Capital over 1 year, i.e. the smallest amount x for which (cf. Bauer et al. (2010))

$$P(AC_1 > 0 | AC_0 = x) \geq 99.5\%. \tag{2}$$

However, since this implicit definition is rather unpractical in (numerical) computations, one often uses the following approximately equal definition (cf. Bauer et al. (2010))

$$SCR^{VaR} := argmin_x \left\{ P\left(AC_0 - \frac{AC_1}{1 + i_1} > x \right) \leq 0.005 \right\}. \tag{3}$$

From Equations (1) and (3), a mutual dependence between Available Capital and SCR becomes obvious: The SCR is computed as the VaR of the Available Capital, and the Available Capital depends on the SCR via the Risk Margin. In order to solve this circular relation, CEIOPS (2009a) states that – whenever a life underwriting risk stress is based on the change in value of assets minus liabilities – the liabilities should not include a Risk Margin when computing the SCR. Thus, it is assumed that the Risk Margin does not change in stress scenarios and that the change in Available Capital can be approximated by the change in Net Asset Value

$$NAV_t := A_t - BEL_t.$$

Here, A_t denotes the market value of assets and BEL_t the Best Estimate Liabilities at time t. For simplicity, we will refer to the latter as the liabilities only in what follows.

2.2 The Solvency Capital Requirement for longevity risk

As mentioned in the introduction, the Solvency II standard model follows a modular approach where the modules' and submodules' SCRs are computed separately and then aggregated according to pre-specified correlation matrices. Thus, for the submodule of longevity risk the SCR should, in principle, be computed as (cf. Equation (3))

$$SCR_{long}^{VaR} := argmin_x \left\{ P\left(NAV_0 - \frac{NAV_1}{1+i_1} > x \right) \leq 0.005 \right\}, \tag{4}$$

where BEL_t and A_t in the definition of NAV_t correspond to the liabilities of all contracts which are exposed to longevity risk and the associated assets, respectively.

In the current specification of the Solvency II standard model, however, the SCR for longevity risk – as an approximation of SCR_{long}^{VaR} – is determined as the change in Net Asset Value due to a longevity shock at time $t = 0$ (cf. CEIOPS (2008a)), i.e.

$$SCR_{long}^{shock} := NAV_0 - (NAV_0 | longevity\, shock). \tag{5}$$

The longevity shock is a permanent reduction of the mortality rates for each age by 25%. This shock is only meant to account for systematic changes in mortality and does not account for small sample risk. Therefore, we also disregard any small sample risk in the VaR approach as an inclusion would blur the results of our comparison in Section 5. In what follows, we only consider the SCR for longevity risk and thus omit the index $\cdot_{long}$ for simplicity.

3 The mortality model

3.1 Model requirements

The computation of the SCR for longevity risk via the VaR approach obviously requires stochastic modeling of mortality. In the literature, a considerable number of stochastic mortality models has been proposed and for an overview we refer to Cairns et al. (2008). However, only very few of these models are suitable for the computation of a VaR over a 1-year time horizon.

From an annuity provider's perspective, longevity risk in the 1-year setting of Solvency II consists of two components: First, there is the risk that next year's realized mortality will be (significantly) below its expectation, e.g., due to a mild winter with less people than usual dying from flu. The second component is the risk of a decrease in expected mortality beyond next year, for which a cure for cancer is the classical example. A newly discovered medication against cancer would have to be tested thoroughly first and it would take some time until it would be available to a group of people large enough such that mortality on a population scale would be affected. Thus, a noticeable effect on next year's realized mortality is rather unlikely. However, (long-term) mortality assumptions would certainly have to be revised. Both components of the longevity risk lead to higher than expected liabilities at $t = 1$: in the first case, because more insured would still be alive than assumed and in the second case, because for those who are still alive liabilities at $t = 1$ on a best estimate basis would be larger than anticipated at $t = 0$. Hence, in order to properly assess longevity risk over a 1-year time horizon, a stochastic mortality model must account for both components of the risk.

The most common mortality models belong to the class of spot models where only realized (period) mortality is modeled. To account for anticipated changes in mortality over time (typically by assuming a decrease in mortality), spot models contain a mortality trend assumption. However, in most spot models this trend is fixed as part of the calibration process and scenarios of realized mortality are derived as random deviations from this trend. This means that the liabilities at $t = 1$ are always computed based on the same trend assumption as at $t = 0$ and thus, these models do not account for the second component of the longevity risk.

Some authors have overcome the issue of a fixed mortality trend though and explicitly allow for changes in this trend. For instance, Cox et al. (2009) propose a very comprehensive model which includes components for mortality jumps and for mortality trend reductions. However, besides the issue that the calibration of their model significantly relies on expert judgment, they only allow for temporary trend changes. The long-term mortality trend assumption is fixed and thus, only parts of the second component of the longevity risk are covered by their model. Very similar arguments hold for the models of Biffis (2005) and Hari et al. (2008b) where the process which models the mortality trend is mean reverting with the mean fixed at $t = 0$. Milidonis et al. (2010) take a different approach and use a Markov regime switching model with different trend and volatility assumptions in the two mortality regimes under consideration. However, in their model, the anticipated long-term mortality trend is rather fixed by the stationary distribution of the Markov chain and hence, the model does also not sufficiently account for possible changes in the anticipated long-term mortality trend. A way to overcome this issue is presented by Sweeting (2009) who incorporates a trend change component into the model of Cairns et al. (2006a). The trend parameters in the two stochastic factors can change each year and, in contrast to the previous models, the current trend parameters always determine the best estimate mortality trend until infinity. However, due to only three possible scenarios for each trend parameter in a 1-year simulation (upward/downward movement by a fixed amount or remaining unchanged), the range of possible overall trend changes over one year is very limited. This makes a reasonable computation of a 99.5% quantile impossible as, for such a computation to be sound, the distribution of the magnitude of trend changes would have to be (at least approximately) continuous. Therefore, we conclude that also very recently developed spot models which allow for trend changes are not directly applicable in the Solvency II framework.

A class of mortality models which overcomes the outlined drawbacks of spot models are so-called forward mortality models. Such models can be seen as an extension of spot models in the time/maturity dimension as they require the expected future mortality as input and model changes in this quantity over time. Thus, they simultaneously allow for random evolutions of realized mortality and for changes in the long-term mortality expectation. Even changes in expected mortality for only certain future time periods can be modeled in contrast to trend-varying spot models where changes in the trend parameters influence expected mortality at all future points in time. This makes forward models usually more complex than spot models but for a reasonable computation of a 1-year VaR this additional complexity seems inevitable. Moreover, even if an adequate spot model existed, a forward model would still offer the advantage that no nested simulations would be required for the computation of the liabilities at $t = 1$. In a forward modeling framework, these liabilities can be computed directly based on the current mortality surface, whereas, when using a spot model, for each 1-year simulation path another set of paths is needed to compute the liabilities by Monte

Carlo simulation.

For our analyses in the following sections, we use a slightly modified version of the forward mortality model introduced by Bauer et al. (2008, 2010) which we refer to as the BBRZ model. Forward mortality models have been proposed by other authors as well, e.g. Dahl (2004), Miltersen and Persson (2005), or Cairns et al. (2006b). However, to our knowledge, they do not provide concrete specifications of their models. Thus, the BBRZ model is the only forward model which is readily available for practical applications. In the following subsection, we introduce the forward mortality framework in which the BBRZ model is specified. The full specification of the model is given in the appendix where also an improved calibration algorithm is presented.

3.2 The forward mortality framework

Following Bauer et al. (2008, 2010), for the remainder of this paper we fix a time horizon T^* and a filtered probability space $(\Omega, \mathcal{F}, \mathbf{F}, P)$, where $\mathbf{F} = (\mathcal{F})_{0\le t\le T^*}$ satisfies the usual conditions. We assume a large but fixed underlying population of individuals, which is reasonable as we disregard any small sample risk in our setting. Each age cohort in this population is denoted by its age x_0 at time $t_0 = 0$. We also assume that the best estimate forward force of mortality with maturity T as from time t,

$$\mu_t(T, x_0) := -\frac{\partial}{\partial T} \log\left\{ E_P\left[{}_Tp_{x_0}^{(T)} \middle| \mathcal{F}_t \right] \right\} \overset{T\ge t}{=} -\frac{\partial}{\partial T} \log\left\{ E_P\left[{}_{T-t}p_{x_0+t}^{(T)} \middle| \mathcal{F}_t \right] \right\}, \tag{6}$$

is well defined. Here, ${}_{T-t}p_{x_0+t}^{(T)}$ denotes the proportion of $(x_0 + t)$-year olds at time $t \le T$ who are still alive at time T, i.e. it is the survival rate or the "realized survival probability". Moreover, we assume that $(\mu_t(T, x_0))$ satisfies the stochastic differential equations

$$d\mu_t(T, x_0) = \alpha(t, T, x_0)\, dt + \sigma(t, T, x_0)\, dW_t,\ \mu_0(T, x_0) > 0,\ x_0, T \ge 0, \tag{7}$$

where $(W_t)_{t\ge0}$ is a d-dimensional standard Brownian motion independent of the financial market, and $\alpha(t, T, x_0)$ as well as $\sigma(t, T, x_0)$ are continuous in t. Furthermore, the drift term $\alpha(t, T, x_0)$ has to satisfy the drift condition (cf. Bauer et al. (2008))

$$\alpha(t, T, x_0) = \sigma(t, T, x_0) \times \int_t^T \sigma(t, s, x_0)'\, ds. \tag{8}$$

This implies that a forward mortality model is fully specified by the volatility $\sigma(t, T, x_0)$ and the initial curve $\mu_0(T, x_0)$.

From definition (6), we can deduce the best estimate survival probability for an x_0-year old to survive from time zero to time T as seen at time t:

$$E_P\left[{}_Tp_{x_0}^{(T)} \middle| \mathcal{F}_t \right] = E_P\left[e^{-\int_0^T \mu_u(u,x_0)\, du} \middle| \mathcal{F}_t \right] = e^{-\int_0^T \mu_t(u,x_0)\, du}. \tag{9}$$

For $t = 1$, these are the survival probabilities we need in order to compute the liabilities and the Net Asset Value after one year. Note that $\mu_t(u, x_0) = \mu_u(u, x_0)$ for $u \le t$ since the volatility $\sigma(t, u, x_0)$ must obviously be equal to zero for $u \leq t$. Inserting the dynamics (7) into (9) yields

$$E_P\left[{}_Tp_{x_0}^{(T)}\middle|\,\mathcal{F}_t\right] = E_P\left[{}_Tp_{x_0}^{(T)}\right] \times e^{-\int_0^T\left(\int_0^t \alpha(s,u,x_0)\,ds+\int_0^t \sigma(s,u,x_0)\,dW_s\right)du}, \tag{10}$$

which means that we actually do not have to specify the forward force of mortality. For our purposes, it is sufficient to provide – besides the volatility $\sigma(s,u,x_0)$ – best estimate T-year survival probabilities as at time $t = 0$ which can, from a technical point of view, be obtained from any generational mortality table. Given these quantities, we are able to compute the SCR for longevity risk via the VaR approach empirically by means of Monte Carlo simulation.

Since we use a deterministic volatility (cf. appendix) the forward force of mortality is Gaussian. Hence, it could become negative for extreme scenarios with survival probabilities becoming larger than 1. However, for the 50,000 paths we have randomly chosen for the empirical computation of the VaR in the subsequent analysis such a scenario does not occur and therefore, we regard this as a negligible shortcoming for practical considerations, even in tail scenarios. Moreover, survival probabilities larger than 1 can also be regarded as unproblematic with respect to longevity risk as they merely make a Gaussian forward mortality model more conservative.

An appealing feature of the Gaussian setting is that – given a deterministic "market price of risk" process $(\lambda(t))_{t\geq 0}$ – the volatilities under the real world measure P and under the equivalent martingale measure Q generated by $(\lambda(t))_{t\geq 0}$ via its Radon-Nikodym density (see, e.g., Harrison and Kreps (1979) or Duffie and Skiadas (1994))

$$\left.\frac{\partial Q}{\partial P}\right|_{\mathcal{F}_t} = \exp\left\{\int_0^t \lambda(s)'\,dW_s - \frac{1}{2}\int_0^t \|\lambda(s)\|^2\,ds\right\},$$

coincide (cf. Bauer et al. (2008)). Hence, risk-adjusted survival probabilities, i.e. survival probabilities under the measure Q, can easily be derived from their best estimate counterparts via

$$\begin{aligned}
& E_Q\left[{}_Tp_{x_0}^{(T)}\middle|\,\mathcal{F}_t\right] \\
= \; & E_Q\left[e^{-\int_0^T \mu_u(u,x_0)\,du}\middle|\,\mathcal{F}_t\right] = E_Q\left[e^{-\int_0^T\left(\mu_t(u,x_0)+\int_t^u \alpha(s,u,x_0)\,ds+\int_t^u \sigma(s,u,x_0)\,dW_s\right)du}\middle|\,\mathcal{F}_t\right] \\
= \; & e^{-\int_0^T\int_t^u \sigma(s,u,x_0)\,\lambda(s)\,ds\,du} \times E_Q\left[e^{-\int_0^T\left(\mu_t(u,x_0)+\int_t^u \alpha(s,u,x_0)\,ds+\int_t^u \sigma(s,u,x_0)\,(dW_s-\lambda(s)\,ds)\right)du}\middle|\,\mathcal{F}_t\right] \\
= \; & e^{-\int_0^T\int_t^u \sigma(s,u,x_0)\,\lambda(s)\,ds\,du} \times E_P\left[{}_Tp_{x_0}^{(T)}\middle|\,\mathcal{F}_t\right],
\end{aligned} \tag{11}$$

since $\left(\tilde{W}_t\right)_{t\geq 0}$ with $\tilde{W}_t = W_t - \int_0^t \lambda(u)\,du$ is a Q-Brownian motion by Girsanov's Theorem (see e.g. Theorem 3.5.1 in Karatzas and Shreve (1991)). We will make use of this relation when we analyze the Risk Margin in Section 7.

4 The model setup

In this section, we describe the setting for our analysis of the longevity stress in the Solvency II standard model which includes assumptions on the interest rate evolution, the reference company's asset strategy, the contracts under consideration, as well as the best estimate mortality.

Our reference company is situated in the UK and is closed to new business. We set $t = 0$ in 2007 and as risk-free interest rates we deploy the risk-free term structure for year-end 2007 provided by CEIOPS (2008b) as part of QIS4. For the approximative definition of the SCR (cf. Equation (3)) to coincide with the exact definition (cf. Equation (2)), we assume that the company's assets and Technical Provisions coincide at $t = 0$. Since we want our company to be solely exposed to longevity risk, we also assume that it only invests in risk-free assets and that it is completely hedged against changes in interest rates. As a consequence of the latter, a deterministic approach for the interest rate evolution is sufficient and the company specific risk-free term structure at time t can be deduced from the 2007 term structure. We denote by $i(t, T)$ the annual interest rate for maturity T at time t, $t \leq T$. The company's risk-free assets are traded only when premium payments are received or survival benefits are paid and in each case the asset cash flow coincides with the liability cash flow. Hence, differing values of assets and Technical Provisions at $t > 0$ are always and only due to changes in (expected) mortality. Finally, we disregard any operational risk.

As standard contracts we consider immediate and deferred life annuities which pay a fixed annuity amount yearly in arrears if the insured is still alive. We assume that these contracts do not contain any options or guarantees for death benefits. Moreover, there is no surplus participation and any charges are disregarded. Thus, in case an x_0-year old at $t_0 = 0$ is still alive at time t, the liabilities for an exemplary immediate annuity contract paying £1 are

$$BEL_t = \sum_{T>t} \frac{1}{(1 + i(t,T))^{T-t}} \cdot E_P \left[{}_T p_{x_0}^{(T)} \middle| \mathcal{F}_t \right]. \qquad (12)$$

The best estimate survival probabilities at $t = 0$ are as for UK Life Office Pensioners in 2007. More specifically, we use the basis table PNMA00 which contains amounts based mortality rates for normal entries in 2000, and a projection according to the average of projections used by 5 large UK insurance companies in 2006. For details on this average projection, we refer to Bauer et al. (2010) and Grimshaw (2007).

As a consequence of the assumptions on payment dates and the asset evolution, we have

$$A_1 = A_0 \left(1 + i(0, 1)\right) + CF_1, \qquad (13)$$

where CF_1 denotes the company's (stochastic) cash flow at $t = 1$. In case of immediate annuities, this cash flow is always negative, for deferred annuities it is positive as long as premiums are paid. For a mixed portfolio of running and deferred annuities, it may be positive or negative. Equation (13) implies that

$$NAV_0 - \frac{NAV_1}{1 + i(0,1)} = \frac{BEL_1 - CF_1}{1 + i(0,1)} - BEL_0,$$

and hence, the SCR formula in the VaR approach, Equation (4), simplifies to

$$SCR^{VaR} = argmin_x \left\{ P\left(\frac{BEL_1 - CF_1}{1 + i(0,1)} - BEL_0 > x \right) \leq 0.005 \right\}. \qquad (14)$$

Thus, we can disregard the evolution of the assets in our computations. In our subsequent analysis, this

SCR is computed empirically based on the forward mortality model introduced in the previous section. For $t = 1$, Equation (10) describes how best estimate survival probabilities at the end of the year can be simulated. For each path, the Best Estimate Liabilities and the cash flow at $t = 1$ are then computed based on these probabilities and the SCR finally corresponds to the empirical 99.5% quantile of the loss variable

$$\frac{BEL_1 - CF_1}{1 + i(0,1)} - BEL_0.$$

Analogously to the SCR formula in the VaR approach, also the SCR formula in the shock approach, Equation (5), can be reformulated in terms of the liabilities:

$$SCR^{shock} = (BEL_0 | longevity\, shock) - BEL_0. \tag{15}$$

Before we can start the comparison of the different SCR formulas, however, we need to ensure that the comparison does not get blurred simply because of the differing structures of the SCR formulas in combination with peculiarities of our modeling setup. To this end, an important observation can be made from the simplified SCR formulas: The SCR in the VaR approach only depends on the expected mortality at the end of the year, i.e. at $t = 1$, because cash flows occur only yearly in our setting. Hence, it does not matter whether changes in mortality emerge gradually over the year or in a one-off shock at $t = 0$. Moreover, every gradual mortality evolution can be expressed in a one-off shock at $t = 0$, namely in a shock which transforms the expected survival probabilities at $t = 0$ into the ones obtained at $t = 1$ via the gradual evolution. Therefore, if we express a longevity scenario which yields the SCR in the VaR approach by a one-off shock denoted by $shock^{VaR}$, we obtain

$$\begin{aligned}
SCR^{VaR} &= argmin_x \left\{ P\left(\frac{BEL_1 - CF_1}{1 + i(0,1)} - BEL_0 > x \right) \leq 0.005 \right\} \\
&= \left(\frac{BEL_1 - CF_1}{1 + i(0,1)} \middle| \, shock^{VaR} \right) - BEL_0 \\
&= \left(BEL_0 | \, shock^{VaR} \right) - BEL_0 \\
&= SCR^{shock}.
\end{aligned}$$

The question for the subsequent analysis is now whether $shock^{VaR}$ can be the 25% shock. If the SCRs computed according to Equations (14) and (15) coincide a 25% reduction in mortality rates corresponds to a 200-year scenario for the mortality evolution over one year. Or in other words, if the SCRs differ significantly, the 25% shock approach is not a reasonable approximation of the VaR concept.

5 Comparison of the SCR formulas

We start our analysis with a life annuity of £1000 for a 65-year old. The liabilities and the SCRs for the shock approach and the VaR approach are given in Table 1 and we observe that the shock approach requires about 26% more capital than the VaR approach. Even though the standard model longevity stress is meant to be conservative, this deviation seems rather large. Also in relation to the liabilities at $t = 0$, the deviation

	BEL_0	$BEL_1 - CF_1$	SCR	SCR/BEL_0
Shock approach	12,619.28	14,238.81	869.87	6.9%
VaR approach	12,619.28	14,050.62	691.59	5.5%

Table 1: SCRs for a 65-year old

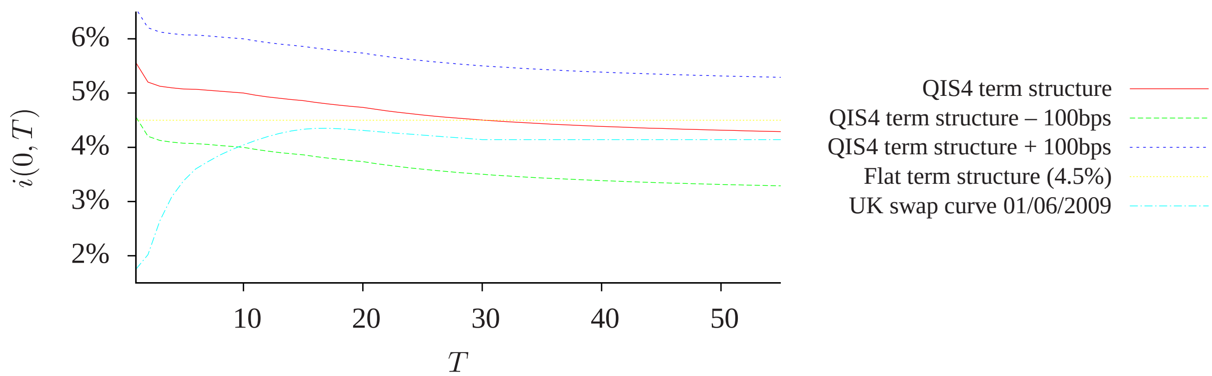

Figure 1: Different term structures for SCR computation

of about 1.4% is significant.

A natural question now is what happens if we vary certain parameters in our computations, e.g. the age, the best estimate survival probabilities, or the payment dates of the survival benefits. However, in order to ensure that the deviation in SCRs is not primarily due to the term structure under consideration, let us first recompute the SCRs in Table 1 for different interest rates. In addition to the QIS4 risk-free term structure, we consider term structures shifted upwards and downwards by 100bps, a flat term structure at 4.5%, and the UK swap curve as at 01/06/2009 (obtained from Bloomberg on 12/06/2009). For the latter, we interpolate the quotes for the first 30 years with cubic splines and assume a flat yield curve thereafter. The differences between these term structures are illustrated in Figure 1 and Table 2 provides the resulting liabilities and SCRs.

As expected, the liabilities and also the SCRs increase with decreasing interest rates. Moreover, from the values of SCR^{shock}/BEL_0 and SCR^{VaR}/BEL_0, we can observe that the SCR in the shock approach is slightly more volatile than its counterpart in the VaR approach and thus more sensitive to changing interest rates. Nevertheless, the relative deviation between the SCRs

$$\frac{\Delta SCR}{SCR^{VaR}} = \frac{SCR^{shock} - SCR^{VaR}}{SCR^{VaR}} \tag{16}$$

remains rather constant as can be seen in the last column. Hence, the interest rates obviously do not have a significant impact on our analysis.

In order to analyze the SCRs for different ages, we consider ages between 55 and 105 as given in the first column of Table 3. In the third and fourth column we see that the SCR in the shock approach first increases with age and then decreases again and that the SCR relative to the liabilities BEL_0 becomes rather large for

Term structure	BEL_0	SCR^{shock}	$\frac{SCR^{shock}}{BEL_0}$	SCR^{VaR}	$\frac{SCR^{VaR}}{BEL_0}$	$\frac{\Delta SCR}{SCR^{VaR}}$
QIS4	12,619.28	869.87	6.9%	691.59	5.5%	25.8%
QIS4 - 100bps	14,002.91	1,090.17	7.8%	848.91	6.1%	28.4%
QIS4 + 100bps	11,448.34	701.71	6.1%	569.47	5.0%	23.2%
Flat (4.5%)	13,040.42	887.65	6.8%	710.19	5.5%	25.0%
Swap 01/06/2009	13,646.54	950.05	7.0%	753.73	5.5%	26.0%

Table 2: SCRs for different term structures

Age	BEL_0	SCR^{shock}	$\frac{SCR^{shock}}{BEL_0}$	SCR^{VaR}	$\frac{SCR^{VaR}}{BEL_0}$	$\frac{\Delta SCR}{SCR^{VaR}}$	$\frac{\Delta SCR}{BEL_0}$
55	15,671.10	657.23	4.2%	729.88	4.7%	-10.0%	-0.5%
65	12,619.28	869.87	6.9%	691.59	5.5%	25.8%	1.4%
75	8,941.83	1,009.81	11.3%	513.27	5.7%	96.7%	5.6%
85	4,940.13	1,003.43	20.3%	304.89	6.2%	229.1%	14.1%
95	2,549.75	818.58	32.1%	214.38	8.4%	281.8%	23.7%
105	1,413.19	646.23	45.7%	180.79	12.8%	257.4%	32.9%

Table 3: SCRs for different ages

very old ages. These observations are due to the structure of the longevity shock since the effect of the 25% reduction increases as the mortality rates increase. In the VaR approach, the SCR decreases with age and decreasing liabilities, and only the ratio of SCR and liabilities increases which seems more intuitive. In the last two columns, we can observe that for ages above 85, the SCR in the shock approach more than triples the SCR in the VaR approach and also in relation to the liabilities the deviation is huge. These observations clearly question the adequacy of the longevity stress in the standard model.

On the other hand, for age 55 the SCR for the current shock calibration is already smaller than its counterpart in the VaR approach. Therefore, a simple adjustment of the shock's magnitude such that the SCRs for old ages approximately coincide may lead to a significant underestimation of the longevity risk for younger ages. Hence, the shortcoming of the standard model longevity stress seems to be more a structural one and, in principle, an age-dependent stress with smaller relative reductions for old ages might be more appropriate. This coincides with epidemiological findings: The number of relevant causes of death is larger for old ages than for young ages (cf. Tabeau et al. (2001)) which means that even if some cause of death is contained, older people are more likely to die from another cause of death. Therefore, significant reductions in mortality rates seem to be more difficult to achieve for old ages compared to young ages.

Table 4 shows SCRs for various mortality levels for a 65-year old. The different levels have been established by shifting the mortality rates in the basis table as stated in the first column of the table. Here, two opposing trends can be observed. With increasing mortality, the SCR in the shock approach increases whereas the SCR in the VaR approach decreases – in absolute value as well as relative to the liabilities. This leads to an almost identical SCR for a 20% mortality downward shift and a further increasing deviation in SCRs for mortality upward shifts. The reason for the increase in SCR^{shock} is again the structure of the shock: the higher the mortality rates, the larger the shocks in absolute terms. For small mortality rates, an increased

Mortality shift	BEL_0	SCR^{shock}	$\frac{SCR^{shock}}{BEL_0}$	SCR^{VaR}	$\frac{SCR^{VaR}}{BEL_0}$	$\frac{\Delta SCR}{SCR^{VaR}}$	$\frac{\Delta SCR}{BEL_0}$
-20%	13,296.28	849.03	6.4%	844.65	6.4%	0.5%	0.03%
-10%	12,940.87	860.24	6.7%	758.61	5.9%	13.4%	0.79%
original	12,619.28	869.87	6.9%	691.59	5.5%	25.8%	1.41%
+10%	12,325.42	878.43	7.1%	637.37	5.2%	37.8%	1.96%
+20%	12,054.68	886.19	7.4%	593.84	4.9%	49.2%	2.43%

Table 4: SCRs for different mortality levels

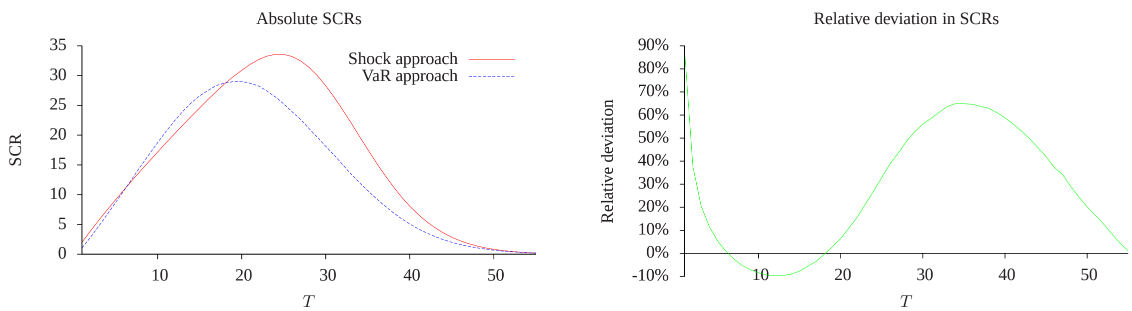

Figure 2: SCRs for 1-year endowment contracts with maturity T

long-term risk due to longer survival does not compensate for the smaller shocks. In contrast to this kind of "relative volatility" assumption in the shock approach, the volatility in the BBRZ model does not depend on the mortality level. This results in the decreasing SCR with increasing mortality as earlier deaths reduce the long-term risk. The automatic adjustment to the mortality level under consideration is a convenient property of the standard model approach and the results emphasize that the BBRZ model should be calibrated to data at the mortality level in view for the application. In case such data is not sufficiently available, the deficit may be reduced by calibrating the model to a more extensive set of data and then adjusting the volatility to the mortality level of the population under consideration, e.g. by some kind of scaling.

Finally, we want to analyze the impact of varying payment dates on the deviation in SCRs. To this end, we do not consider an annuity contract but simple endowment contracts for a 65-year old with only one payment of £1000 conditional on survival up to time T. Thus, a combination of these contracts for maturities $T = 1, .., 55$ equals the payoff structure of the life annuity. In the left panel of Figure 2, the SCRs are plotted for maturities up to $T = 55$. We observe that until $T = 20$, the SCRs in both approaches are rather close in absolute value. Thereafter, the shock approach demands significantly more SCR which explains the larger SCR for the life annuity. Finally, the SCRs in both approaches converge to zero due to an extremely low probability of survival. Note though that in the VaR approach, the sum of SCRs in Figure 2 is about 5% larger than the SCR for the corresponding life annuity because diversification between different maturities is disregarded when computing the SCR for each endowment separately. In the shock approach, the sum of the SCRs for the endowments coincides with the SCR for the life annuity since the stress in the standard model does not allow for any diversification. In the right panel of Figure 2, the relative deviation in SCRs (cf. Equation (16)) is displayed. It varies considerably over time ranging from about -10% to more than

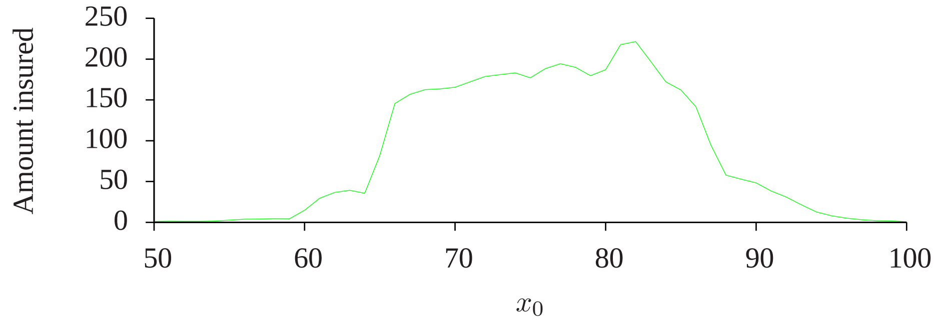

Figure 3: Age composition of the portfolio of immediate annuity contracts

	BEL_0	$BEL_1 - CF_1$	SCR	SCR/BEL_0
Shock approach	36,394.73	42,939.27	4,283.83	11.8%
VaR approach	36,394.73	40,587.52	2,055.90	5.7%

Table 5: SCRs for a portfolio of immediate life annuities

80%. Hence, we can conclude that a longevity stress independent of the maturity under consideration is not adequate.

So far, we have only compared SCRs for single contracts. Now, we want to investigate if the depicted shortcomings of the standard model longevity stress also lead to deviating SCRs for a realistic portfolio of immediate life annuities. In order to ensure a reasonable age composition, we choose a portfolio according to Continuous Mortality Investigation (CMI) data on the amounts based exposures of UK Life Office Pensioners in 2006. We only consider males and normal entries and rescale the data by dividing by 100,000 to obtain results in a handy range. The age composition of the portfolio is illustrated in Figure 3.

From Table 5, we observe that, for this portfolio of contracts, the SCR in the shock approach more than doubles the SCR in the VaR approach. Given the age composition of the portfolio, this is consistent with the observed SCRs for different ages (see Table 3). Thus, the aforecited structural shortcomings of the standard model longevity stress also affect the SCR for a realistic portfolio of contracts. According to Table 5, an insurance company which applies the standard model would have to raise additional solvency capital of about 6% of its liabilities compared to a company which implements the BBRZ model for VaR computation in an internal model. This is a huge amount given that, according to a stylized balance sheet for all undertakings in QIS4, Net Asset Value accounts for only about 18% of the total liabilities (cf. CEIOPS (2008c)).

When analyzing the SCRs for different ages, we found that the relation between the SCRs seems to turn around for young ages. To further investigate this, we consider deferred life annuities for different ages which again pay £1000 in arrears starting at age 65. Table 6 contains the corresponding SCRs.

The SCRs increase in absolute value with age and for age 60, the SCRs in both approaches almost coincide. However, for younger ages the SCR in the VaR approach is (significantly) larger than its counterpart in the shock approach with the relative deviation growing, the younger the age under consideration. Here, the

Age	BEL_0	SCR^{shock}	$\frac{SCR^{shock}}{BEL_0}$	SCR^{VaR}	$\frac{SCR^{VaR}}{BEL_0}$	$\frac{\Delta SCR}{SCR^{VaR}}$	$\frac{\Delta SCR}{BEL_0}$
30	3,205.97	217.90	6.8%	382.66	11.9%	-43.1%	-5.1%
35	3,851.54	268.30	7.0%	428.53	11.1%	-37.4%	-4.2%
40	4,623.92	329.89	7.1%	489.79	10.6%	-32.7%	-3.5%
45	5,549.28	404.85	7.3%	561.71	10.1%	-27.9%	-2.8%
50	6,676.64	495.51	7.4%	631.44	9.5%	-21.5%	-2.0%
55	8,100.16	604.29	7.5%	688.98	8.5%	-12.3%	-1.1%
60	9,978.02	733.27	7.4%	724.06	7.3%	1.3%	0.1%

Table 6: SCRs for deferred annuities at different ages

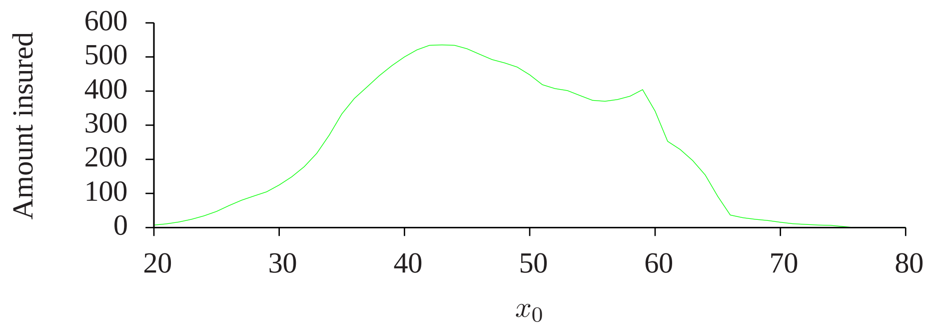

Figure 4: Age composition of the portfolio of deferred annuity contracts

longevity stress in the standard model seems to (significantly) underestimate the longevity risk. Thus, we again find that a longevity shock independent of age does not seem appropriate.

As for the immediate annuities, we also want to analyze the consequences of the highlighted shortcomings of the shock approach for a realistic portfolio of deferred annuity contracts. This time, we derive the age composition of our portfolio from CMI data on lives based exposures of male UK Personal Pensioners in 2006. We confine the data to pension annuities in deferment and omit data on ages below 20 as the BBRZ model does not cover those ages. The latter restriction is unproblematic for our purposes as exposures for those ages are very small in the CMI data. Moreover, we assume yearly benefits in arrears of £0.01 for all contracts to keep all values in a handy range. The resulting overall amount insured for each age is displayed in Figure 4. The benefit payments are assumed to commence at age 65 or, in case age 65 has already been reached, after one more year of deferment.

The results in Table 7 confirm the observation for single ages that the shock approach demands significantly less SCR for young ages. The significance of this difference is underlined by the fact that the shock SCR of 6629.31 corresponds to a 1.5% default probability in the VaR approach. Thus, under the assumption that the BBRZ model assesses longevity risk correctly, the true default probability of our reference company would be three times as large as accepted if the company used the Solvency II standard model. This also holds for the case when premiums are paid during the deferment period. In a further analysis, we found that, in such a setting, the liabilities were obviously much lower but that the SCRs changed only slightly.

For a combined portfolio of immediate and deferred annuities the deviations in SCRs from the two ap-

	BEL_0	$BEL_1 - CF_1$	SCR	SCR/BEL_0
Shock approach	88,165.37	100,062.89	6,629.31	7.5%
VaR approach	88,165.37	101,485.14	7,976.68	9.1%

Table 7: SCRs for a portfolio of deferred life annuities

proaches would obviously cancel each other out to some extent. Thus, despite its depicted structural shortcomings the standard model longevity stress may yield a reasonable overall value for the longevity SCR in some cases. However, there are insurance companies in the market whose portfolios have a strong focus on certain age groups. Therefore, an equal reduction of mortality rates for all ages and maturities does not seem appropriate for the Solvency II standard model. To further illustrate this, we can cite the magnitudes of shocks which would yield SCRs for the two considered portfolios equal to those in the VaR approach. For the immediate annuities, a shock of about 13.0% would have been sufficient, whereas for the deferred annuities, a shock of about 29.4% would have been required.

Therefore, in the following section we propose a modification of the longevity stress in the standard model which overcomes the shortcomings of the current approach but does still not require stochastic simulation of mortality.

6 A modification of the standard model longevity stress

In this section, we explain how the longevity stress in the Solvency II standard model could be modified in order to overcome the shortcomings highlighted in the previous section which result from an equal shock for all ages and maturities.

We keep the structure of a one-off shock which means the integration of the longevity stress into the standard model does not change. From our point of view, such a one-off shock is an acceptable approximation of gradual changes in expected mortality over only one year as the distinction between a one-off and a gradual mortality evolution only affects cash flows (premium payments and/or survival benefits) during the first year. Cash flows thereafter only depend on the expected mortality at $t = 1$ and not on how this expectation emerged (cf. Section 4). Regarding the cash flows in the first year, we see in the left panel of Figure 2 that the corresponding SCR is typically very small compared to those for later years. Furthermore, the implementation of gradual mortality changes would make the standard model more complex as, currently, all risks other than longevity risk are implicitly excluded by computing the shocked liabilities at $t = 0$.

We modify the longevity shock according to the volatility in the BBRZ model which introduces a dependency structure of the shock magnitude on age and maturity. Instead of the 1-year mortality rates, we now shock the (expected) T-year survival probabilities. These are the quantities which are typically required for computing the liabilities of longevity prone contracts like annuities (cf. Equation (12)). Thus, it seems plausible to specify a longevity shock in terms of extreme changes in these quantities over one year. Moreover, the increasing uncertainty with time in the expected future mortality evolution is accounted for explicitly and thus, in a more plausible way. When shocking the annual mortality rates, this increasing uncertainty is only considered implicitly by the accumulation of an increasing number of shocked annual probabilities. If one wanted to stick to a shock of the annual mortality rates the shocks for different maturities T should thus be specified such that the accumulated shock is reasonable.

For the modified shock, we set each T-year survival probability to its individual 99.5% quantile in a 1-year simulation of the BBRZ model. From Equation (10), factors can easily be derived for each age and maturity which, when multiplied with the best estimate T-year survival probabilities, yield the desired percentiles. Thus, once the BBRZ model is calibrated, supervisory authorities would only have to provide insurance companies with a matrix of these factors and the computational efforts for calculating the SCR based on this modified shock would basically remain as in the 25% shock approach.

In principle, the modified shock yields an SCR which is larger than the one in the VaR approach since any diversification effects between different ages and maturities are disregarded when computing the 99.5% quantiles of each survival probability individually. In case of the portfolios considered in the previous section, this gives a markup of about 5.2% of the SCR for the immediate annuities and of about 9.4% for the deferred annuities. However, we regard this inaccuracy as acceptable given the considerable structural improvement of the modified shock compared to the original 25% shock. Moreover, such a markup might actually be desirable since the longevity stress in the standard model is supposed to be conservative in order to give incentives for the implementation of internal models. From this perspective, the modified shock also offers the additional advantage that a markup is distributed over all ages and maturities according to the actual risk, whereas, in the 25% shock, this distribution is not directly risk-related.

As indicated in the previous section, the BBRZ model implies the application of an absolute volatility and obviously, the same holds for the modified shock. The shock factors are only adequate as long as the mortality level in the insurance portfolio in view (approximately) coincides with the level of mortality experience to which the BBRZ model has been calibrated. However, this drawback could be dealt with by calibrating the BBRZ model to the mortality experience of different populations (with different mortality levels) such that the resulting shock factors allow for possibly different risk profiles – in magnitude as well as in structure.

In case the provision of a matrix of shock factors is regarded unpractical or too complex for the standard model, the surface formed by the factors could be approximated by a function of age and maturity. Then only this function would have to be provided by the supervisory authorities. The fitting of such a function would also allow for an extrapolation of the shock factors to ages below 20 for which the volatility in the BBRZ model as specified in Bauer et al. (2008) is not defined. Figure 5 shows the log log of the shock factors in the current calibration of the BBRZ model for all feasible combinations of initial ages x_0 and maturities T. At first sight, the surface looks like a fairly even plane for all maturities $T \geq 10$. Therefore, a function of the form $\exp\{\exp\{f(x_0, T)\}\}$ might be applicable as an approximation of the shock factors, where $f(x_0, T)$ is a plane with a correction for short maturities. However, the surface of the shock factors may look different for a different calibration of the BBRZ model and hence, also this approximating function might have to be derived individually for each population in view. We leave this question for future work.

Figure 6 shows the ratio of shock factors for T-year survival probabilities based on the modified shock and the original 25% shock of the mortality rates. Unfortunately, a general comparison is not possible as

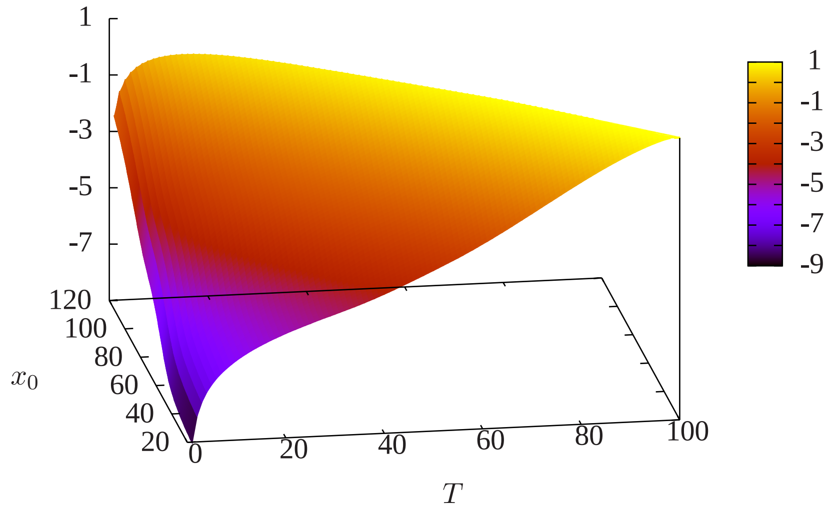

Figure 5: Log log of factors for modified shock approach

the magnitude of the original shock depends on the level of the mortality rates. Here, we apply the 25% reduction to the aforementioned Life Office Pensioners mortality table. We observe that for combinations of ages up to about 60 and maturities up to about 50 the shocks are very similar. However, for higher initial ages the ratios decrease significantly which means the modified shock demands significantly less SCR for those ages. On the other hand, for rather young initial ages and long maturities, the ratios increase rapidly as can be seen for maturities T up to 70 in the figure. Therefore, the 25% shock may significantly underestimate the risk of changes in long-term mortality trends. All observations made from Figure 6 are in line with those made in Section 5.

7 The Risk Margin for longevity risk

7.1 Risk Margin approximations

In Section 5, we found structural shortcomings of the longevity stress currently implemented in the Solvency II standard model. These shortcomings may also affect the Risk Margin which is computed as the cost of capital for future SCRs in the run-off of an insurance portfolio (cf. Equation (1)). We will now analyze this point by comparing the Risk Margin in the 25% shock approach with its counterpart in the modified shock approach (as a proxy of the VaR approach) introduced in the previous section.

However, an exact computation of the Risk Margin would require the determination of each year's SCR, SCR_t, conditional on the mortality evolution up to time t. Since this is practically impossible, the Risk Margin is usually approximated. In the literature and in CEIOPS (2008a) in particular, several such approximations have been proposed and we will consider the following four:

(I) Approximation of the SCR_t by assuming a mortality evolution up to time t according to its best

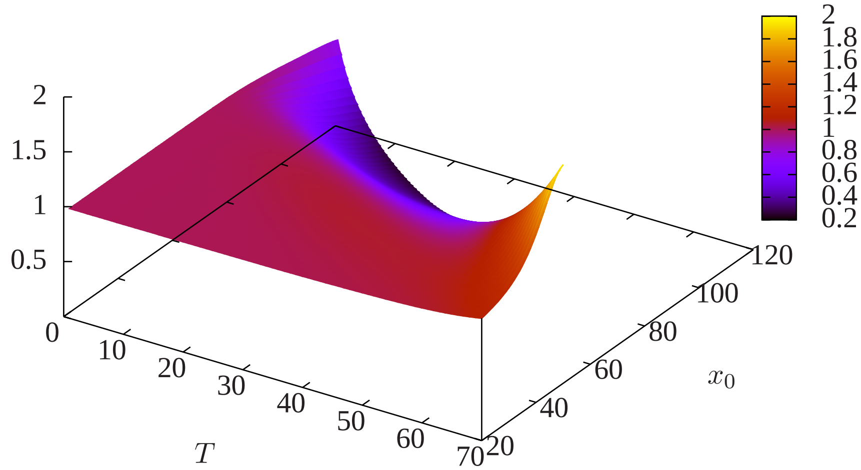

Figure 6: Ratio of shock factors from modified and 25% longevity shock

estimate. Thus, we have

$$RM^{(I)} = \sum_{t\geq 0} \frac{CoC}{(1+i(0,t+1))^{t+1}} \cdot SCR_t^{BE}.$$

In practice, this calculation method is generally seen as yielding the "exact" Risk Margin and therefore, we are going to use it as a benchmark for the other approximation methods.

(II) Approximation of the SCR_t using the proxy formula in CEIOPS (2008a, TS.XI.C.9), i.e.

$$RM^{(II)} = \sum_{t\geq 0} \frac{CoC}{(1+i(0,t+1))^{t+1}} \cdot 25\% \cdot q_t^{av} \cdot 1.1^{(dur_t-1)/2} \cdot dur_t \cdot BEL_t,$$

where q_t^{av} is the expected average 1-year death rate at time t weighted by sum assured and dur_t is the modified duration of the liabilities at time t.

(III) Approximation of the SCR_t as the fraction SCR_0/BEL_0 of the liabilities BEL_t, i.e. assuming a constant ratio of SCRs and liabilities over time (cf. CEIOPS (2008a, TS.II.C.28)):

$$RM^{(III)} = \sum_{t\geq 0} \frac{CoC}{(1+i(0,t+1))^{t+1}} \cdot \frac{SCR_0}{BEL_0} \cdot BEL_t.$$

(IV) Direct approximation of the Risk Margin via the modified duration of the liabilities (cf. CEIOPS (2008a, TS.II.C.26)):

$$RM^{(IV)} = CoC \cdot dur_0 \cdot SCR_0.$$

Table 8 contains the Risk Margins for the 25% longevity shock and the modified longevity shock, the two portfolios of annuity contracts and all four approximation methods as far as they are well defined. The cost of capital rate CoC is set to its current calibration in the Solvency II standard model, i.e. 6%. For the

Portfolio	Method	BEL_0	25% longevity shock			Modified longevity shock		
			RM	Rel. dev.	$\frac{RM}{BEL_0}$	RM	Rel. dev.	$\frac{RM}{BEL_0}$
Immediate annuities	(I)	36,394.73	2,383.87		6.6%	1,143.08		3.1%
	(II)	36,394.73	2,751.71	15.4%	7.6%	n/a	n/a	n/a
	(III)	36,394.73	1,957.24	-17.9%	5.4%	988.11	-13.6%	2.7%
	(IV)	36,394.73	2,051.70	-13.9%	5.6%	1,035.80	-9.4%	2.9%
Deferred annuities	(I)	88,165.37	11,240.74		12.8%	12,159.44		13.8%
	(II)	88,165.37	10,126.91	-9.9%	11.5%	n/a	n/a	n/a
	(III)	88,165.37	10,034.70	-10.7%	11.4%	13,206.71	8.6%	15.0%
	(IV)	88,165.37	10,488.60	-6.7%	11.9%	13,804.08	13.5%	15.7%

Table 8: Risk Margins for different portfolios, approximation methods, and longevity shocks

modified longevity shock, approximation method (II) is obviously not applicable since it is based on the 25% shock of the mortality rates. Also note that this method is actually not admissible for the portfolio of deferred annuities because the average age in this portfolio does not reach the required 60 years (cf. CEIOPS (2008a)). Thus, the value for the deferred annuities in case of a 25% shock is only given for comparison.

Regarding the two portfolios, we find that the Risk Margin in relation to the liabilities is generally larger for the portfolio of deferred annuities which seems reasonable given the more long-term risk resulting in larger uncertainty. This finding is supported by the graphs in the left panels of Figure 7 where the evolutions of the SCRs according to approximation method (I) are displayed for both portfolios and shock approaches. For the immediate annuities, the SCRs decrease toward zero quite directly whereas for the deferred annuities, the SCRs first increase before fading out as less and less insured survive.

Furthermore, for the “exact” calculation method (I) and the immediate annuities portfolio, we observe that the Risk Margin in the 25% shock approach more than doubles its counterpart in the modified shock approach. This is not surprising as we have observed the same relation for the SCRs at $t = 0$ (see Table 5). Moreover, in the top left panel of Figure 7, we can see that this relation holds thereafter as well. For the portfolio of deferred annuities, however, the Risk Margins are similar even though the SCRs at $t = 0$ differ significantly (see Table 7). Here, a compensation over time occurs which becomes obvious in the bottom left panel of Figure 7: The higher SCRs in the modified shock approach for smaller t to some extent compensate for the smaller SCRs later on, i.e. when the portfolio becomes similar to the immediate annuities portfolio. Nevertheless, depending on the portfolio under consideration, the structural shortcomings of the 25% shock also adversly affect the Risk Margin.

With respect to the Risk Margin approximations, we find that they seem rather crude in general with the closest approximation still being about 6.7% away from the “exact” value. However, the deviations appear less significant in light of the large uncertainty regarding the correct cost of capital rate (cf. CEIOPS (2009b)) and the value from calculation method (I) itself only being an approximation of unknown accuracy. Nevertheless, we think that the wide range of values for the approximated Risk Margin for each portfolio and longevity shock combination is problematic. The Solvency II regime is supposed to ease and improve comparisons of the solvency situations of European insurance companies, but such comparisons might get

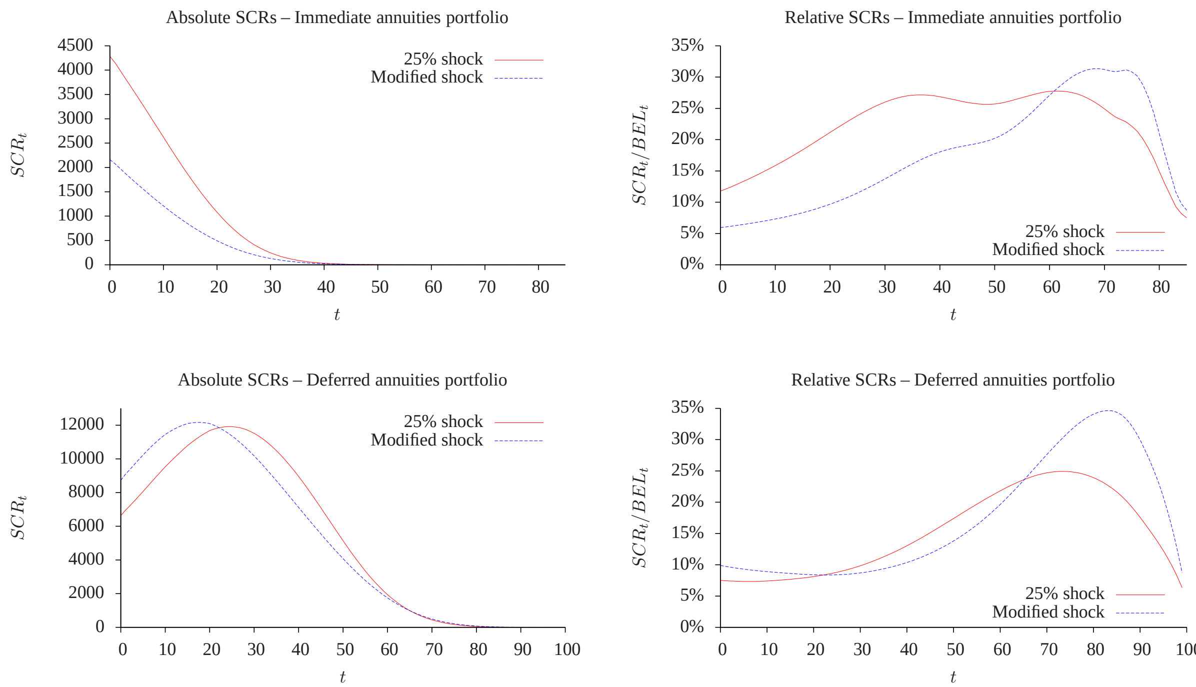

Figure 7: Evolution of the SCRs in portfolio run-offs

blurred simply by the use of different Risk Margin approximations. Moreover, instead of assessing their risk as accurately as possible, insurance companies might be tempted to apply the Risk Margin approximation which yields the smallest value.

Finally, we have a closer look at the Risk Margin for approximation method (III) which is very popular in practice. From Table 8, we observe that the assumption, that future SCRs are proportional to future liabilities, results in the largest relative deviations in three of the four combinations of portfolios and longevity shocks. In each of these cases, the “exact” Risk Margin is underestimated. This is due to the ratio SCR_0/BEL_0 being a very crude if not inadequate proxy for ratios of future SCRs and liabilities as can be seen in the right panels of Figure 7. For both portfolios and both longevity shocks, the ratio is not even approximately constant but, in general, increases with time. This observation is in line with the findings of Haslip (2008) for non-life insurance. Thus, approximation method (III) does not seem appropriate in its current form. However, an adjustment of the proxy, e.g., in form of a dependence on the average age in the portfolio may improve results considerably.

7.2 The Cost of Capital rate

As already indicated, there is a large uncertainty and hence, an ongoing discussion regarding the correct cost of capital rate. It is currently set to 6%, but in CEIOPS (2009b) and references therein, values between 2% and 8% have been derived from different shareholder return models. An alternative idea for the calibration of the cost of capital rate has been brought forward by Olivieri and Pitacco (2008c) who rely on reinsurance premiums. However, since data on such premiums is not sufficiently available their idea is currently more

Portfolio of contracts	25% longevity shock		Modified longevity shock	
	RM	λ	RM	λ
Immediate annuities	2,383.87	18.6%	1,143.08	13.2%
Deferred annuities	11,240.74	8.7%	12,159.44	8.9%

Table 9: Risk Margins and corresponding Sharpe ratios λ

of theoretical interest.

In what follows, we take a different approach to finding an adequate cost of capital rate. We compare the Risk Margin based on the current rate to hypothetical market prices of longevity dependent liabilities. The Risk Margin should coincide with the markup of the risk-adjusted liabilities over their best estimate counterparts in such a market as it is supposed to provide a risk adjustment of the Best Estimate Liabilities. This idea is also in line with the Risk Margin's more specific interpretation in the cost of capital setting, i.e. the provision of sufficient capital to guarantee a proper run-off of a portfolio. If a market for longevity risk existed an insolvent insurer could guarantee the portfolio run-off by transferring the risk to the market at the cost of the Best Estimate Liabilities and the Risk Margin.

We assume that, in the hypothetical longevity market, risk-adjusted survival probabilities are derived from their best estimate counterparts as explained in Section 3. For the market price of longevity risk process $\lambda(t)$, we assume a time-constant Sharpe ratio λ, i.e. the simplest possible process, since currently, there is no information available on the structure of a market price of longevity risk process (cf. Bauer et al. (2010)). The comparison is then performed by finding the Sharpe ratio which yields a markup in the liabilities equal to the cost of capital Risk Margin.

Table 9 shows the Risk Margins and the corresponding Sharpe ratios for all combinations of portfolios and longevity shocks. For each portfolio, we observe that the Sharpe ratios increase with the Risk Margin which is what one would expect. Moreover, we see that the Sharpe ratios for the deferred annuities are smaller than those for the immediate annuities. This means that the risk adjustment of the survival probabilities perceives a larger risk in the longer maturities of the deferred annuities than the Risk Margin does as the Sharpe ratios are chosen such that the markups in the hypothetical longevity market coincide with the Risk Margins. Nevertheless, we observe that the Sharpe ratios are all of reasonable magnitude. For the modified longevity shock, where the computation of the Risk Margin and the risk-adjustment of the survival probabilities are performed based on (essentially) the same model/volatility, the Sharpe ratios do not seem too large in particular. For comparison, Bauer et al. (2010) find that Sharpe ratios for longevity risk might lie somewhere between 5% and 17% and Loeys et al. (2007) regard a value of 25% as reasonable. The moderate values of λ we have found suggest that, at least in the case of longevity risk, the current cost of capital rate of 6% is not overly conservative. This is particularly the case if we take into account that, in a shock scenario which leads to a company's insolvency and for which the Risk Margin is to be provided, investors typically require higher risk compensation than usual.

So far, we have assumed the existence of a market for longevity risk and based on that we have made inferences on the Risk Margin and the cost of capital rate in particular. However, it is very likely that in practice, the Solvency II regime and its Risk Margin requirements will come into effect before a deep and liquid market for longevity risk exists. Hence, it makes sense to look at the relation between the Risk Margin and the market-consistent valuation of longevity risk also from the opposite perspective. Thus, we now want to analyze what conclusions for the pricing of longevity-linked securities can possibly be drawn from Risk Margin requirements once the latter have been fixed.

Given an appropriate mortality model, the pricing of longevity derivatives narrows down to the specification of a reasonable market price of longevity risk process (or parameter) – $\lambda(t)$ in our setting. Hence, we need to answer the question whether such a process can be reasonably derived from Risk Margin requirements. In the literature, several ideas have already been proposed on how such processes could be obtained, e.g. from market annuity quotes (cf. Lin and Cox (2005)) or other asset classes like stocks (cf. Loeys et al. (2007)), but none of these approaches is really satisfactory. For a thorough discussion of this issue, we refer to Bauer et al. (2010).

A company which is completely hedged against longevity risk does not have to provide any solvency capital for longevity risk anymore. In our setting, where the reference company is solely exposed to longevity risk this obviously reduces the future cost of solvency capital to zero. Thus, neglecting any possible credit risk arising from a longevity transaction, the company might be interested in securitizing its longevity risk as long as the transaction price does not exceed the present value of the future cost of capital in case it keeps the risk. By definition, the Risk Margin is equal to this present value for a typical insurance company and therefore, it corresponds to the maximum price such a company would be willing to pay for longevity risk securitization. Consequently, the Sharpe ratios in Table 9 reflect the level of risk compensation a typical market supplier of longevity risk might accept and hence, what a reasonable market price of longevity risk could be.

Obviously, this approach to finding a market price of longevity risk also has its shortcomings, in particular because a company's decision of securitizing its longevity risk is influenced by several other effects. For instance, the maximum price a company is willing to accept may be lower because it may expect its own cost of capital to be lower than the Risk Margin, e.g. because of diversification effects with other risks than longevity. On the other hand, a company might accept a higher Sharpe ratio for strategic reasons, e.g. the abandonment of a line of business, or due to difficulties in raising capital and the risk of increasing cost of capital in the future. Obviously, the relevance of these effects will vary between companies and as a consequence, different companies will accept different market prices of longevity risk. The market's appetite for longevity risk will then decide on which company can securitize its risk at an acceptable price and what the market price for longevity risk will finally be. Nevertheless, we believe that solvency requirements and the Solvency II Risk Margin in particular, can provide valuable insights into and a reasonable starting point for the pricing of longevity derivatives.

8 Conclusion

In the Solvency II standard model, the SCR for longevity risk is computed as the change in Net Asset Value due to a permanent 25% reduction in mortality rates. This scenario based approach is to approximate the

99.5% VaR of the Available Capital over one year which can only be determined exactly via stochastic simulation of mortality. However, this standard model longevity stress has come under some criticism for its possibly unrealistically simple structure. In particular, the reduction in mortality rates does not depend on age and maturity which may lead to an incorrect assessment of the true longevity risk. To assess this issue, we compare the standard model longevity stress to the 99.5% VaR for longevity risk which we compute by the forward mortality model of Bauer et al. (2008, BBRZ model).

We find that, in our setting, the SCRs from both approaches differ considerably in most cases. In general, the VaR approach yields a larger SCR for young ages whereas the shock approach demands more SCR for old ages. Moreover, we observe varying SCRs for different maturities. Thus, depending on the composition of an insurance portfolio, the longevity stress in the standard model may significantly overestimate or underestimate the true longevity risk. In the former case, companies would be forced to hold an unnecessarily large amount of capital whereas in the latter case, a company's default risk would be significantly higher than the accepted level of 0.5%. For instance, for an exemplary but realistic portfolio of deferred annuities we observe that the standard model longevity stress demands an SCR which corresponds to a default probability of 1.5% in the VaR approach.

Hence, the current longevity stress in the Solvency II standard model seems to have some crucial structural shortcomings. In particular, the shock magnitude's independence of age and maturity does not seem appropriate. Even though we have performed our analysis in a rather simple setting, we expect that our observations are valid rather generally. Allowing, e.g., for more complex contracts with surplus participation and various options and guarantees might change the company's exposure to longevity risk. For instance, the risk for certain ages or maturities might be reduced and SCR values may therefore change considerably. However, the inherent longevity risk would still be very similar and we must expect the observed deviations between the two approaches and the structural shortcomings of the standard model longevity stress to remain.

Therefore, we believe that a modification of this longevity stress is necessary so that it more appropriately reflects the risk insurance companies are exposed to. Such a modified longevity stress is proposed in this paper. In order to keep the standard model's longevity risk module as simple as possible, we keep the one-off shock structure. However, we define a different shock – for each age x_0 and maturity T – according to the 99.5% quantile of the expected T-year survival probability for an x_0-year old as in a 1-year simulation of the BBRZ model. This shock can be applied by multiplying the best estimate survival probabilities with corresponding shock factors. Thus, the computational efforts for calculating the SCR would essentially remain similar to the 25% shock but the risk perception would be improved considerably.

We then analyze the cost of capital Risk Margin – in our case for longevity risk only – and find that the proposed approximations in CEIOPS (2008a) lead to significantly different values. Thus, two companies' solvency situations might differ considerably only due to the chosen approximation method for the Risk Margin. Moreover, we observe that the assumption of a constant ratio of SCRs and liabilities over time, which is a popular proxy for future SCRs in practice, often leads to inadequate results. We find mostly increasing ratios for the two portfolios under consideration which means that the approximated Risk Margin might be too small to fulfill its purpose of guaranteeing a proper run-off of a portfolio in case of insolvency.

This finding should be particularly interesting for regulators.
Subsequently, we discuss the adequacy of the current calibration of the cost of capital rate by comparing the Risk Margin to the price of a securitization in a hypothetical market for longevity risk. We observe that the cost of capital rate of 6% corresponds to Sharpe ratios in the market which are of reasonable magnitude. Thus, the rate of 6% does not seem overly conservative for longevity risk. Finally, we explain how market prices for longevity risk may be derived from solvency requirements and the Solvency II Risk Margin in particular. The securitization of longevity risk reduces a company's future cost of solvency capital and the Risk Margin as the present value of these costs therefore provides a maximum price which a typical insurance company might be willing to accept for the securitization of its risk. Even though there are several other effects which may influence the company's decision on securitization, we believe that solvency requirements can provide a reasonable starting point for the pricing of longevity derivatives.

References

Bauer D, Börger M, Ruß J, Zwiesler H-J (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: 139–149.

Bauer D, Bergmann D, Reuß A (2010) On the Calculation of the Solvency Capital Requirement Based on Nested Simulations. Working Paper, Georgia State University and Ulm University.

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

Booth H, Tickle L (2008) Mortality Modelling and Forecasting: A Review of Methods. Annals of Actuarial Science, 3: 3–41.

Booth H, Maindonald J, Smith L (2002) Applying Lee-Carter under conditions of variable mortality decline. Population Studies, 56: 325–336.

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: 687–718.

Cairns A, Blake D, Dowd K (2006) Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk. ASTIN Bulletin, 36: 79–120.

Cairns A, Blake D, Dowd K (2008) Modelling and Management of Mortality Risk: A Review. Scandinavian Actuarial Journal, 2: 79–113.

Cairns A, Blake D, Dowd K, Coughlan G, Epstein D, Ong A, Balevich I (2009) A Quantitative Comparison of Stochastic Mortality Models Using Data from England & Wales and the United States. North American Actuarial Journal, 13: 1–35.

CEIOPS (2007) QIS3 Calibration of the underwriting risk, market risk and MCR. Available at: http://www.ceiops.eu/media/files/consultations/QIS/QIS3/QIS3CalibrationPapers.pdf.

CEIOPS (2008) QIS4 Technical Specifications. Available at: http://www.ceiops.eu/media/docman/Technical%20Specifications%20QIS4.doc.

CEIOPS (2008) QIS4 Term Structures. Available at: http://www.ceiops.eu/media/docman/public_files/consultations/QIS/CE DOC-23-08%20rev%20QIS4%20Term%20Structures%2020080507.xls.

CEIOPS (2008) CEIOPS' Report on its fourth Quantitative Impact Study (QIS4) for Solvency II. Available at: http://www.ceiops.eu/media/files/consultations/QIS/CEIOPS-SEC-82-08%20QIS4%20Report.pdf.

CEIOPS (2009) Consultation Paper No. 49, Draft CEIOPS' Advice for Level 3 Implementing Measures on Solvency II: Standard formula SCR - Article 109 c Life underwriting risk. Available at: http://www.ceiops.eu/media/files/consultations/consultationpapers/CP49/CEIOPS-CP-49-09-L2-Advice-Standard-Formula-Life-Underwriting-risk.pdf.

CEIOPS (2009) Consultation Paper No. 42, Draft CEIOPS' Advice for Level 2 Implementing Measures on Solvency II: Article 85(d) - Calculation of the Risk Margin. Available at: http://www.ceiops.eu/media/files/consultations/consultationpapers/CP42/CEIOPS-CP-42-09-L2-Advice-TP-Risk-Margin.pdf.

CEIOPS (2010) Solvency II Calibration Paper. Available at: http://www.ceiops.eu/media/files/publications/submissionstotheec/CEIOPS-Calibration-paper-Solvency-II.pdf.

Chan W, Li S, Cheung S (2008) Testing Deterministic versus Stochastic Trends in the Lee-Carter Mortality Indexes and its Implication for Projecting Mortality Improvements at Advanced Ages. Available at:, http://www.soa.org/library/monographs/retirement-systems/living-to-100-and-beyond/2008/january/mono-li08-6a-chan.pdf.

Continuous Mortality Investigation (CMI) (2009) Working Paper 37 – Version 1.1 of the CMI Library of Mortality Projections. Available at: www.actuaries.org.uk.

Cox S, Lin Y, Pedersen H (2009) Mortality Risk Modeling: Applications to Insurance Securitization. Insurance: Mathematics and Economics, 46: 242–253.

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

Devineau L, Loisel S (2009) Risk Aggregation in Solvency II: How to Converge the Approaches of the Internal Models and Those of the Standard Formula? Bulletin Français d'Actuariat, 18: 107–145.

Doff R (2008) A Critical Analysis of the Solvency II Proposals. The Geneva Papers on Risk and Insurance Issues and Practice, 33: 193–206.

Dowd K, Cairns A, Blake D (2006) Mortality-Dependent Financial Risk Measures. Insurance: Mathematics and Economics, 38: 427–440.

Duffie D, Skiadas C (1994) Continuous-time security pricing: A utility gradient approach. Journal of Mathematical Economics, 23: 107–131.

Eling M, Schmeiser H, Schmit J (2007) The Solvency II Process: Overview and Critical Analysis. Risk Management and Insurance Review, 10: 69–85.

European Commission (2010) QIS5 Technical Specifications. Available at: http://www.ceiops.eu/media/files/consultations/QIS/QIS5/QIS5-technical_specifications_20100706.pdf.

Grimshaw D (2007) Mortality Projections. Presentation at the Current Issues in Life Assurance seminar, 2007.

Hanewald K (2009) Mortality Modeling: Lee-Carter and the Macroeconomy. Working Paper, Humboldt University Berlin.

Hari N, De Waegenaere A, Melenberg B, Nijman T (2008) Longevity Risk in Portfolios of Pension Annuities. Insurance: Mathematics and Economics, 42: 505–519.

Hari N, De Waegenaere A, Melenberg B, Nijman T (2008) Estimating the term structure of mortality. Insurance: Mathematics and Economics, 42: 492–504.

Harrison M, Kreps D (1979) Martingales and arbitrage in multiperiod security markets. Journal of Economic Theory, 20: 381–408.

Haslip G (2008) Risk Assessment. The Actuary, Dec 2008.

Holzmüller I (2009) The United States RBC Standards, Solvency II and the Swiss Solvency Test: A Comparative Assessment. The Geneva Papers on Risk and Insurance Issues and Practice, 34: 56–77.

Human Mortality Database (2009) University of California, Berkeley, USA, and Max Planck Institute for Demographic Research, Germany. Available at: www.mortality.org (Data downloaded on 12/04/2009).

Karatzas I, Shreve S (1991) Brownian Motion and Stochastic Calculus. Graduate Texts in Mathematics, 113. Springer, New York, NY (USA).

Lee R, Carter L (1992) Modeling and Forecasting US mortality. Journal of the American Statistical Association, 87: 659–671.

Lee R, Miller T (2001) Evaluating the Performance of the Lee-Carter Method for Forecasting Mortality. Demography, 38: 537–549.

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

Loeys J, Panigirtzoglou N, Ribeiro R (2007) Longevity: A Market in the Making. JPMorgan Global Market Strategy.

Milidonis A, Lin Y, Cox S (2010) Mortality Regimes and Pricing. To appear in: North American Actuarial Journal.

Miltersen K, Persson S (2005) Is mortality dead? Stochastic force of mortality determined by no arbitrage. Working Paper, Norwegian School of Economics and Business Administration, Bergen and Copenhagen Business School.

Olivieri A (2009) Stochastic Mortality: Experience-Based Modeling and Application Issues Consistent with Solvency 2. Working Paper, University of Parma.

Olivieri A, Pitacco E (2008) Solvency Requirements for Life Annuities: Some Comparisons. Working Paper, University of Parma and University of Trieste.

Olivieri A, Pitacco E (2008) Stochastic Mortality: the Impact on Target Capital. CAREFIN Working Paper, University Bocconi.

Olivieri A, Pitacco E (2008) Assessing the Cost of Capital for Longevity Risk. Insurance: Mathematics and Economics, 42: 1013–1021.

Plat R (2009) Stochastic Portfolio Specific Mortality and the Quantification of Mortality Basis Risk. Insurance: Mathematics and Economics, 45: 123–132.

Steffen T (2008) Solvency II and the Work of CEIOPS. The Geneva Papers on Risk and Insurance Issues and Practice, 33: 60–65.

Stevens R, De Waegenaere A, Melenberg B (2010) Longevity risk in pension annuities with exchange options: The effect of product design. Insurance: Mathematics and Economics, 46: 222–234.

Sweeting P (2009) A Trend-Change Extension of the Cairns-Blake-Dowd Model. Working Paper, The Pensions Institute.

Tabeau E, van den Berg Jeths A, Heathcote C (2001) Towards an Integration of the Statistical, Demographic and Epidemiological Perspectives in Forecasting Mortality. In Tabeau, E., van den Berg Jeths, A., Heathcote, C., Eds.: Forecasting Mortality in Developed Countries. Kluwer Academic Publishers, Dordrecht.

Thatcher A, Kannisto V, Vaupel J (1998) The Force of Mortality at Ages 80 to 120. In Odense Monographs on Population Aging 5. Odense University Press, Odense, Denmark.

Vaupel J (1986) How Change in Age-Specific Mortality Affects Life Expectancy. Population Studies, 40: 147–157.

Wilmoth J (1993) Computational Methods for Fitting and Extrapolating the Lee-Carter Model of Mortality Change. Technical Report, University of California, Berkeley.

Appendix

A Specification and calibration of the BBRZ-model

A.1 Model specification

In general, a forward mortality model in the framework of Section 3 is fully specified by the volatility $\sigma(t,T,x_0)$ and the initial curve $\mu_0(T,x_0)$. For our purposes, however, the volatility and T-year survival probabilities are sufficient, where for the latter we refer to Section 4. Regarding the volatility, Bauer et al. (2008) propose a 6-factor specification of which we use a slightly modified version in this paper. We keep the general functional structure of the volatility vector $\sigma(t,T,x_0) = (\sigma_1(t,T,x_0),\ldots,\sigma_6(t,T,x_0))$ and only deploy a different functional form for what they refer to as the "correction term". Instead of a standard Gompertz form, we use a logistic Gompertz form $\mu_x = \frac{\exp(ax+b)}{1+\exp(ax+b)} + c$ as introduced by Thatcher et al. (1998). In contrast to the standard Gompertz law, this prevents the volatility from becoming unrealistically large for very old ages. Hence, for $T \geq t$ the volatility vector $\sigma(t,T,x_0)$ consists of the following components:

$$\begin{aligned}
&\text{general:} && \sigma_1(t,T,x_0) = c_1 \times \left(\tfrac{\exp(a(x_0+T)+b)}{1+\exp(a(x_0+T)+b)} + c\right);\\
&\text{short-term:} && \sigma_2(t,T,x_0) = c_2 \times \left(\tfrac{\exp(a(x_0+T)+b)}{1+\exp(a(x_0+T)+b)} + c\right) \times \exp(\log(0.1)(T-t));\\
&\text{young age:} && \sigma_3(t,T,x_0) = c_3 \times \left(\tfrac{\exp(a(x_0+T)+b)}{1+\exp(a(x_0+T)+b)} + c\right)\\
& && \quad \times \exp\left(\tfrac{\log(0.5)}{20^2}(T-t-20)^2 + \tfrac{\log(0.5)}{17.5^2}(x_0+T-37.5)^2\right);\\
&\text{middle age:} && \sigma_4(t,T,x_0) = c_4 \times \left(\tfrac{\exp(a(x_0+T)+b)}{1+\exp(a(x_0+T)+b)} + c\right)\\
& && \quad \times \exp\left(\tfrac{\log(0.5)}{20^2}(T-t-20)^2 + \tfrac{\log(0.5)}{12.5^2}(x_0+T-67.5)^2\right);\\
&\text{old age:} && \sigma_5(t,T,x_0) = c_5 \times \left(\tfrac{\exp(a(x_0+T)+b)}{1+\exp(a(x_0+T)+b)} + c\right)\\
& && \quad \times \exp\left(\tfrac{\log(0.5)}{20^2}(T-t-20)^2 + \tfrac{\log(0.5)}{30^2}(x_0+T-110)^2\right);\\
&\text{long-term:} && \sigma_6(t,T,x_0) = c_6 \times \left(\tfrac{\exp(a(x_0+T)+b)}{1+\exp(a(x_0+T)+b)} + c\right)\\
& && \quad \times \exp\left(\tfrac{\log(0.5)}{80^2}(T-t-120)^2\right).
\end{aligned}$$

For $T < t$, the volatility has to be zero obviously, since the forward force of mortality for maturity T is already known at time t and hence deterministic.

A.2 Generational mortality tables for model calibration

Since forward mortality models reflect changes in expected future mortality they need to be calibrated to quantities which contain information on the evolution of expected mortality. Such quantities could, e.g., be market prices of longevity derivatives, market annuity quotes, or generational mortality tables. However, since data on longevity derivatives' prices and annuity quotes is very sparse and/or blurred due to charges, profit margins etc., currently, generational mortality tables seem to be the most appropriate starting point for the calibration of forward mortality models. Hence, Bauer et al. (2008, 2010) calibrate their model to series of generational mortality tables for US and UK pensioners. However, they have only available 4 or 7 tables, respectively, which is a rather thin data base for model calibration.

Moreover, it is not clear how good the volatility derived from such historical mortality tables fits the future volatility since the methodology of constructing mortality projections has changed significantly over the past decades, i.e. from the technique of age shifting to stochastic (spot) mortality models. Therefore, the transition in projection methods often lead to rather sudden and strong "jumps" in predicted mortality rates from one table to the next which might result in an unreasonably high volatility in the forward model.

Therefore, we believe that it is more appropriate to calibrate forward mortality models to a series of historical generational mortality tables which have all been constructed using the same and up-to-date projection methodology. In order to build such tables, we use the model of Lee and Carter (1992). With this model, a generational mortality table is constructed for each historical year based on the mortality data which was available in that year. We opt for the Lee-Carter model because it has become a standard model in the literature on mortality forecasting (cf. Booth and Tickle (2008)) and it has also been used by the CMI in the construction of some of their recent projections for UK annuitants and pensioners (cf. CMI (2009) and references therein). However, we think that the model choice is not crucial in our setting for two reasons: First, for the calibration of the forward model, we are not interested in absolute values of projected mortality rates but only in their changes over time. Secondly, the conclusions drawn in this paper with respect to the structure of the longevity stress in the Solvency II standard model should still be valid even if the true volatility in the forward model was generally overestimated or underestimated.

In the Lee-Carter model, the (central) mortality rate $m(x,t)$ for age x in year t is parameterized as

$$\log\{m(x,t)\} = \alpha_x + \beta_x \cdot \kappa_t + \epsilon_{x,t},$$

where the α_x describe the average mortality rate for each age x over time, β_x specifies the magnitude of changes in mortality for age x relative to other ages, $\epsilon_{x,t}$ is an age and time dependent normally distributed error term, and κ_t is the time trend (often also referred to as mortality index). The latter is usually assumed to follow an ARIMA(0,1,0) process, i.e.

$$\kappa_t = \kappa_{t-1} + \gamma + e_t,$$

where γ is a constant drift and e_t is a normally distributed error term with mean zero.

For fitting this model to historical mortality data, we use the weighted least squares algorithm introduced by Wilmoth (1993) and adjust the κ_t by fitting a Poisson regression model to the annual number of deaths at each age (see Booth et al. (2002) for details). Since mortality data usually becomes very sparse for old ages, we only fit the model for ages 20 to 95 and use the following extrapolation up to age 120: We extract the α_x for $x \geq 96$ from a logistic Gompertz form which we fit in least squares to the α_x for $x \leq 95$. For the β_x, we assume the average value of β_{91} to β_{95} to hold for all $x \geq 91$.

The (deterministic) projection of future mortality rates is then conducted by setting the error terms $\epsilon_{x,t}$ and e_t equal to their mean zero for all future years. Due to the lognormal distribution of the $m(x,t)$, the thus derived mortality rates are actually not the expectations of the future $m(x,t)$, which are typically listed in a generational mortality table, but their medians. However, even though not fully correct, this approach to projecting mortality within the Lee-Carter model is widely accepted (see, e.g., Wilmoth (1993), Lee and

Miller (2001), or Booth and Tickle (2008)) and should be unproblematic in particular in our setting, where only changes in projected mortality over time are considered. Finally, we derive approximate 1-year initial mortality rates $q(x,t)$ from the central mortality rates as (cf. Cairns et al. (2009))

$$q(x,t) \approx 1 - \exp\{-m(x,t)\}.$$

The generational mortality tables are constructed for the male general population of England and Wales. For our purposes, it would obviously be preferable to build tables from annuitant or pensioner mortality data. However, a sufficient amount of such data is only available for a limited number of years and ages. Therefore, we make the assumption that the volatility derived from the general population tables is also valid for insured's mortality. As discussed in Sections 5 and 6, an adjustment of the volatility to the mortality level of the population in view may be appropriate to account for this issue but for simplicity we do not consider such an adjustment here.

Mortality data, i.e. deaths and exposures, for years 1947 to 2006 has been obtained from the Human Mortality Database (2009). A generational mortality table is then constructed based on a Lee-Carter fit to each set of 30 consecutive years of data which yields a series of 31 tables. We decided to use only 30 years of data for each Lee-Carter fit as we think it is not reasonable to calibrate a mortality model with constant parameters b_x and γ to a significantly longer time series of historical data. Changes in age dependent mortality reduction rates have been observed for most countries in the past (see, e.g., Vaupel (1986) and Booth and Tickle (2008)), including England and Wales, and using a larger set of data would thus imply the risk of extrapolating outdated mortality trends from the far past into the future. Moreover, Chan et al. (2008), Hanewald (2009), and Booth et al. (2002), amongst others, find structural breaks in the time trend κ_t for England and Wales as well as for other industrialized countries. A significantly shorter time series of data, on the other hand, might lead to the extrapolation of noise and merely temporary mortality trends. Nevertheless, the choice of 30 years is still rather arbitrary. Additionally, we assume a gap year for data collection between the fitting period for the Lee-Carter model and the starting year of the corresponding mortality table, e.g., from the data for years 1947 to 1976 a table with starting year 1978 is derived. Thus, in total, we obtain a series of 31 mortality tables.

A.3 Calibration algorithm

Based on the series of generational mortality tables derived in the previous subsection, we are now able to calibrate the parameters in the forward model. In analogy to Bauer et al. (2010), we fit the correction term to the forward force of mortality for a 20-year old in the most recent mortality table using least squares and obtain parameter values of $a = 0.1069$, $b = 12.57$, and $c = 0.0007896$.

For the parameters c_i, $i = 1, .., d$ with $d = 6$ in our case, Bauer et al. (2008) present a calibration algorithm based on Maximum Likelihood estimation. However, due to numerical issues they can only use a small part of the available data, i.e. six 1-year survival probabilities for different ages and maturities from each generational mortality table, the choice of which obviously implies some arbitrariness. Here, we propose a new 2-step calibration algorithm which makes use of all available data and the fact that the deterministic volatility specified in Appendix A.1 can be re-written in the form $\sigma_i(s,u,x_0) = c_i \cdot r_i(s,u,x_0)$, $i = 1,..,d$, with the c_i the only free parameters.

Let $t_0 = 0$ be the year of the first generational mortality table, i.e. 1978 in our case, and let $t_1,..,t_{N-1}$ denote the (hypothetical) compilation years of all later tables. We define by

$$p_{x_0+T}^{(t_n):T\to T+1} := \frac{E\left[\left. {}_{T+1}p_{x_0}^{(T+1)} \right| \mathcal{F}_{t_n}\right]}{E\left[\left. {}_{T}p_{x_0}^{(T)} \right| \mathcal{F}_{t_n}\right]} = \exp\left\{-\int_T^{T+1} \mu_{t_n}(u,x_0)\, du\right\}$$

the 1-year forward survival probability from T to $T+1$ for an x_0-year old at time zero as seen at time t_n, which can be obtained from the mortality tables. However, given a set of forward survival probabilities at time t_n, $n \in \{0,1,..,N-2\}$, the fixed volatility specification only allows for certain changes in the forward survival probabilities up to time t_{n+1}. We denote by $\bar{p}_{x_0+T}^{(t_{n+1}):T\to T+1}$ the "attainable" forward survival probabilities at time t_{n+1} with respect to the forward survival probabilities at time t_n and the given volatility specification, and these probabilities satisfy the relation

$$\begin{aligned}\bar{p}_{x_0+T}^{(t_{n+1}):T\to T+1} &= p_{x_0+T}^{(t_n):T\to T+1} \\ &\quad \times \exp\left\{-\int_T^{T+1}\int_{t_n}^{t_{n+1}} \alpha(s,u,x_0)\, ds\, du - \int_T^{T+1}\int_{t_n}^{t_{n+1}} \sigma(s,u,x_0)\, dW(s)\, du\right\}.\end{aligned}$$

Inserting the drift condition (8), some further computations yield

$$\begin{aligned}\log\left\{\bar{p}_{x_0+T}^{(t_{n+1}):T\to T+1}\right\} &= \log\left\{p_{x_0+T}^{(t_n):T\to T+1}\right\} \\ &\quad -\sum_{i=1}^{d} c_i^2 \int_T^{T+1}\int_{t_n}^{t_{n+1}} r_i(s,u,x_0)\int_s^u r_i(s,v,x_0)\, dv\, ds\, du \\ &\quad -\sum_{i=1}^{d}\sum_{m=1}^{M} |c_i|\, N_{i,m}^{(t_{n+1})} \left(\int_{t_n+\frac{m-1}{M}}^{t_n+\frac{m}{M}} \left[\int_T^{T+1} r_i(s,u,x_0)\, du\right]^2 ds\right)^{\frac{1}{2}} \\ &= \log\left\{p_{x_0+T}^{(t_n):T\to T+1}\right\} \\ &\quad -\sum_{i=1}^{d} a_i^{(t_{n+1})} \int_T^{T+1}\int_{t_n}^{t_{n+1}} r_i(s,u,x_0)\int_s^u r_i(s,v,x_0)\, dv\, ds\, du \\ &\quad -\sum_{i=1}^{d}\sum_{m=1}^{M} b_{i,m}^{(t_{n+1})} \left(\int_{t_n+\frac{m-1}{M}}^{t_n+\frac{m}{M}} \left[\int_T^{T+1} r_i(s,u,x_0)\, du\right]^2 ds\right)^{\frac{1}{2}},\end{aligned}$$

where the $N_{i,m}^{(t_{n+1})}$, $i = 1,..,d$, $m = 1,..,M$ are standard normally distributed random variables, $a_i^{(t_{n+1})} \geq 0$, and $b_{i,m}^{(t_{n+1})} \in \mathbb{R}$, $i = 1,..,d$. Obviously, the above equations hold for any $M \in \mathbb{N}$, but in case forward survival probabilities for different ages x_0 and maturities T are considered the choice of M has a significant influence on the correlation between the evolutions of these probabilities. For $M = 1$, they would be fully correlated which, in general, is only the case if the volatility vector is 1-dimensional or component-wise constant. These conditions are clearly not fulfilled in our situation and hence, the parameter M should be chosen as large as numerically feasible to meet the actual correlation structure as accurately as possible. We will come back to this issue later.

The idea behind the first calibration step is now to choose the parameters $a_i^{(t_{n+1})}$ and $b_{i,m}^{(t_{n+1})}$ such that the attainable forward probabilities are as close as possible to the actual probabilities listed in the mortality table at t_{n+1}. This is done by minimizing, for each n, the least squares expression

$$LS_{t_{n+1}} := \sum_{\substack{(T,x_0):\, T \geq t_{n+1},\, x_0+t_n \geq 20,\\ x_0+T-t_{n+1} \leq \omega}} \left(\frac{\log\left\{ p_{x_0+T}^{(t_{n+1}):T\to T+1} \right\} - \log\left\{ \bar{p}_{x_0+T}^{(t_{n+1}):T\to T+1} \right\}}{\log\left\{ p_{x_0+T}^{(t_{n+1}):T\to T+1} \right\}} \right)^2 ,$$

where ω is the limiting age. The changes in forward mortality from time t_n to time t_{n+1} can then be described by the parameters $a_i^{(t_{n+1})}$ and $b_{i,m}^{(t_{n+1})}$ in combination with the volatility $\sigma(s,u,x_0)$.

In the second step, we derive values for the parameters c_i, $i = 1,..,d$ from the $a_i^{(t_{n+1})}$ and $b_{i,m}^{(t_{n+1})}$ via Maximum Likelihood estimation. For each $n \in \{0,1,..,N-2\}$, the aggregated change over all ages and maturities in log forward survival probabilities resulting from the i^{th} mortality effect, i.e. the changes driven by the i^{th} component of the vector of Brownian motions W_t (cf. Equation (7)), $i = 1,..,d$, is given by

$$\begin{aligned} ML_i^{(t_{n+1})} \quad &:= \sum_{\substack{(T,x_0):\, T \geq t_{n+1},\, x_0+t_n \geq 20,\\ x_0+T-t_{n+1} \leq \omega}} c_i^2 \int_T^{T+1} \int_{t_n}^{t_{n+1}} r_i(s,u,x_0) \int_s^u r_i(s,v,x_0)\, dv\, ds\, du \\ &+ \sum_{m=1}^{M} |c_i|\, N_{i,m}^{(t_{n+1})} \left(\int_{t_n+\frac{m-1}{M}}^{t_n+\frac{m}{M}} \left[\int_T^{T+1} r_i(s,u,x_0)\, du \right]^2 ds \right)^{\frac{1}{2}} . \end{aligned}$$

This expression is normally distributed with mean

$$m_i^{(t_{n+1})} := E\left[ML_i^{(t_{n+1})} \right] = \sum_{\substack{(T,x_0):\, T \geq t_{n+1},\, x_0+t_n \geq 20,\\ x_0+T-t_{n+1} \leq \omega}} c_i^2 \int_T^{T+1} \int_{t_n}^{t_{n+1}} r_i(s,u,x_0) \int_s^u r_i(s,v,x_0)\, dv\, ds\, du,$$

and variance

$$\begin{aligned} \left(s_i^{(t_{n+1})} \right)^2 \quad &:= \quad Var\left[ML_i^{(t_{n+1})} \right] \\ &= \quad \sum_{m=1}^{M} \left(|c_i| \sum_{\substack{(T,x_0):\, T \geq t_{n+1},\, x_0+t_n \geq 20,\\ x_0+T-t_{n+1} \leq \omega}} \left(\int_{t_n+\frac{m-1}{M}}^{t_n+\frac{m}{M}} \left[\int_T^{T+1} r_i(s,u,x_0)\, du \right]^2 ds \right)^{\frac{1}{2}} \right)^2 , \end{aligned}$$

and according to the independent increments property of a Brownian motion, the $ML_i^{(t_{n+1})}$ are independent for all $i = 1,..,d$ and $n = 0,..,N-2$. Hence, the density of the random vector

$$\left(ML_1^{t_{(1)}}, \ldots, ML_d^{t_{(1)}}, ML_1^{t_{(2)}}, \ldots, ML_d^{t_{(N-2)}}, ML_1^{t_{(N-1)}}, \ldots, ML_d^{t_{(N-1)}} \right)$$

Parameters	c_1	c_2	c_3	c_4	c_5	c_6
	0.07744	0.07456	0.06747	0.25902	0.04215	0.24054

Table 10: Optimal values for the volatility parameters

is the product of the marginal densities $f_{ML_i^{(t_{n+1})}}$ and we need to maximize the likelihood

$$L := \prod_{n=0}^{N-2} \prod_{i=1}^{d} f_{ML_i^{(t_{n+1})}} \left(\hat{M L}_1^{(t_{n+1})}, \ldots, \hat{M L}_d^{(t_{n+1})}; c_1, \ldots, c_d \right)$$

with respect to c_i, $i = 1, .., d$, where the realizations $\hat{M L}_i^{(t_{n+1})}$ are given by substituting c_i^2 by the corresponding $a_i^{(t_{n+1})}$ and $|c_i| \, N_{i,m}^{(t_{n+1})}$ by the corresponding $b_{i,m}^{(t_{n+1})}$. For numerical reasons, we chose to maximize the corresponding log-likelihood

$$\log\{L\} := \sum_{n=0}^{N-2} \sum_{i=1}^{d} - \log \left\{ s_i^{(t_{n+1})} \right\} - \frac{1}{2} \left(\frac{\hat{M L}_i^{(t_{n+1})} - m_i^{(t_{n+1})}}{s_i^{(t_{n+1})}} \right)^2 .$$

Table 10 contains the resulting parameter values for $M = 365$, i.e. a daily approximation of W_t. However, due to different correction terms (cf. Appendix A.1), these parameter values cannot be directly compared to those in Bauer et al. (2010).

As mentioned above, the correlation structure between the changes in mortality for different ages and maturities can only be approximated with the quality of the approximation increasing in M. In order to check the stability of the calibration algorithm and to ensure the reliability of the resulting parameter values, we performed the algorithm for different choices of M and observed a convergence of the optimal parameter values. For instance, from $M = 200$ to $M = 365$ none of the parameter values changed by more than 0.05%. Therefore, we think the calibration algorithm is stable and the results given in Table 10 are reliable.