

A partial internal model for longevity risk

Søren Fiig Jarner* and Thomas Møller†

August 12, 2013

Abstract. This paper proposes a simple partial internal model for longevity risk within the Solvency 2 framework. The model is closely linked to the mechanisms associated with the so-called Danish longevity benchmark, where the underlying mortality intensity and the trend is estimated yearly based on mortality experience from the Danish life and pension insurance sector and on current data from the entire Danish population. Within this model, we derive an estimate for the 99.5% percentile for longevity risk, which differs from the longevity stress of 20% from the standard model. The new stress explicitly reflects the risk associated with unexpected changes in the underlying population mortality intensity on a one-year horizon and with a 99.5% confidence level. In addition, the model contains a component, which quantifies the unsystematic longevity risk associated with a given insurance portfolio. This last component depends on the size of the specific portfolio.

Key words and phrases. Solvency 2, mortality, longevity stress, Danish longevity benchmark, systematic and unsystematic risk.

*The Danish Labour Market Supplementary Pension Fund, Kongens Vænge 8, DK-3400 Hillerød, Denmark; University of Copenhagen, Universitetsparken 5, DK-2100 København Ø, Denmark; Email: sj@atp.dk

†PFA Pension, Sundkrogsvej 4, DK-2100 København Ø, Denmark; University of Copenhagen, Universitetsparken 5, DK-2100 København Ø, Denmark; Email: thm@pfa.dk

1 Introduction

Within the forthcoming Solvency 2 regime, the solvency capital requirement for life and pension insurance companies in the EU is calculated by taking into account the explicit risks associated with the assets and liabilities. The various risks are divided into a number of modules, e.g. market risk, counterparty risk, life underwriting risk, health underwriting risk and operational risk, where each risk module contains a set of calibrated stress scenarios for the main underlying risks. In a number of so-called quantitative impact studies proposed by the European Commission, the life and pension insurance companies have assessed the economic consequences of the proposed regulation. According to the report EIOPA (2011) on the fifth quantitative impact study, the main risks for the life and pension insurance companies in the life underwriting risk module were lapse risk and longevity risk.

In the Solvency 2 directive, Article 105, 3(b), longevity risk is described as *"the risk of loss, or of adverse change in the value of insurance liabilities, resulting from changes in the level, trend, or volatility of mortality rates, where a decrease in the mortality rate leads to an increase in the value of insurance liabilities (longevity risk)"*, see European Commission (2009). This formulation has also been applied in the subsequent specifications for the sixth quantitative impact study, the so-called long term guarantee assessment, see EIOPA (2012) and EIOPA (2013). In the standard model, longevity risk is quantified by reducing the mortality intensity by 20% regardless of age, underlying assumptions regarding trend, or size of the portfolio. This follows e.g. from European Commission (2011) (not published): *"The capital requirement for longevity risk referred to in point (b) of Article 105(3) of Directive 2009/138/EC shall be equal to the loss in basic own funds of insurance and reinsurance undertakings that would result from an instantaneous permanent decrease of 20% in the mortality rates used for the calculation of technical provisions."* Hence, the risk is defined as the economic loss stemming from an instantaneous, but permanent, decrease in the mortality intensity used for calculating the technical provisions. Following the general framework under Solvency 2, the risk is calibrated using a 99.5% confidence level (Value-at-Risk measure) on a one-year time horizon.

There has been a long-standing discussion about the calibration of the various stresses in the Solvency 2 standard model. The longevity stress has been particularly challenging, since longevity risk is by nature a long term risk, whereas Solvency 2 is based on a one-year time horizon risk assessment. In 2010 the longevity stress was reduced from 25% to the current level of 20% in connection with the launch of the fifth quantitative impact study (QIS 5). For a discussion on calibration of the longevity risk stress, see CEIOPS (2010) and Nielsen (2010).

An alternative to the standard model approach is to use a full or partial internal model, see Chapter VI of European Commission (2011), which describes the requirements for the application of internal models, where the companies may develop alternative models for all or parts of their risks. This is the route we take in this paper.

1.1 The Danish FSA longevity benchmark

In December 2010, the Danish financial supervisory authority (The Danish FSA) introduced a longevity benchmark, which should be used as best estimate by all Danish life and pensions insurance companies for market based valuation of life insurance liabili-

ties. This index is composed of a current mortality intensity and a set of assumptions concerning the expected future mortality improvements. The current mortality intensity is estimated from 5 years of data for a portfolio of insured lives from several large Danish life insurance companies, and hence this mortality intensity will in general differ from the mortality intensity estimated from the total Danish population. We shall also refer to these mortality intensities as the *sector mortality intensity* for the population of insured lives, and the *population mortality intensity* for the average mortality intensity for the total Danish population.

The life and pension insurance companies are obliged to perform a yearly estimation procedure prescribed by the Danish FSA in order to determine whether the companies' mortality experiences deviate significantly from the underlying sector mortality. This estimation is performed within a Poisson regression model, where the underlying sector mortality is fixed. More precisely, the model specifies a parametrization for the company's mortality intensity in terms of the underlying sector mortality via a set of fixed regressors. If the company's current mortality intensity deviates significantly from the benchmark, the company must use an adjusted intensity for the current mortality, where the adjustment is estimated in the regression model. If the difference is not significant, the benchmark intensity is used as best estimate for the current mortality. With this Danish benchmark model, the life and pension companies' current mortality assumptions are adapted to the observed sector mortality intensity via the yearly estimation and test procedures. In particular, this ensures that the mortality assumptions for the current mortality are based on the most recent data.

The assumptions about the expected future mortality improvements in the benchmark are estimated from mortality data for the entire Danish population during the last 30 years. This leads to age- and gender-dependent improvement factors that reflect the average improvements during the last 30 years in the Danish population. The benchmark assumptions for the mortality improvements are updated on a yearly basis. The companies are not required to perform an analysis of the trend in their own population since it is assumed that the portfolio of a single company will in general be too small to establish a significant difference from the benchmark trend. Therefore, the Danish life and pension companies typically use the country specific trend when determining the best estimate for the life insurance liabilities. In particular, this ensures that the companies use assumptions concerning future mortality improvements that are similar to the observed historical improvements.

1.2 The proposed internal model

The Solvency 2 standard model prescribes a pragmatic longevity stress of 20% for all ages and all time points. Since the mortality assumptions for Danish life and pension insurance companies are revised yearly by including recent mortality experience and by following a relatively advanced model, it seems reasonable to propose a more refined longevity stress within this model which is more consistent with a 99.5% confidence level on a one-year time horizon. Therefore, the Danish Society of Actuaries established in 2012 a working group in order to propose a model for the longevity stress, which allows for a measurement of the systematic risk associated with the yearly updating procedure for the longevity benchmark and the unsystematic risk associated with the analysis performed by the individual companies.

In the present paper, we propose a partial internal model for longevity risk within

the Solvency 2 framework. The model is relatively simple, easy to use, and based on the celebrated framework of Lee and Carter (1992) and the extension proposed by Brouhns et al. (2002). Hence, the key components of the model may be easier to comprehend for senior executives than more advanced models.

We derive and calibrate a specific stress for Danish life and pension insurance companies within the Danish regulatory regime consisting of the longevity benchmark. However, the paper can also be viewed as a blueprint for how such stress may be calibrated within other regulatory regimes. Moreover, our approach is not restricted to the Lee-Carter model and could be applied together with more advanced stochastic mortality models.

Our quantification of longevity risk focuses on the impact of new data on the best-estimate (model) mortality used by the company to calculate the value of their liabilities. The procedure for determining the model mortality is *external* to the driving model and the variability stems from using the model mortality algorithm on volatile data. This differs from the approach taken by Börger (2010) and Plat (2011). They consider models where the trend itself is stochastic, and they quantify longevity risk by the change in the expected future trend over one year *within* the model. This approach rests on the use of so-called forward mortality models known from securitization and pricing of mortality derivatives, see e.g. Dahl (2004); Dahl and Møller (2006); Bauer et al. (2010); Cairns (2011). Our approach is more in line with the value-at-risk framework suggested by Richards et al. (2012), see also Olivieri and Pitacco (2009).

1.3 Outline

The paper is organized as follows. Section 2 contains the main results with the calibrated stress and its impact on the remaining life times. In Section 3, we introduce the model framework with joint modeling of the national mortality, the insurance sector mortality and the mortality for an insurance company. In addition, we describe the calibration method adopted in the paper. In Section 4 and 5, the analysis of the systematic and unsystematic risk is carried out, and Section 6 contains a more detailed analysis of the unsystematic mortality risk. Finally, a description of the Danish longevity benchmark and the Poisson variant of the Lee-Carter model can be found in the appendix.

2 Main results

In order to state the results we first need to provide a brief description of the Danish FSA longevity benchmark. The benchmark consists of a current level of mortality for insured lives together with annual rates of improvement of gender- and age-specific mortality intensities. We refer to the former as the (benchmark) *level* and the latter as the (benchmark) *trend*.

The gender-specific, benchmark mortality intensity for a person of age x in year t of gender k (F for females, M for males) takes the form

$$\mu_k^{FSA}(x, t) = \mu_k^{FSA}(x, T) (1 - R_k(x))^{t-T}, \quad (1)$$

where T is the reference year for the benchmark, i.e. the year of the current observed level of mortality. The Danish FSA provides the level, $\mu_k^{FSA}(x, T)$, and the trend, $R_k(x)$, from which the users of the benchmark can construct the full benchmark intensity

surface by use of formula (1). The benchmark is updated yearly to reflect the latest information on Danish mortality.

Given the benchmark each company estimates its own company-specific mortality relative to the benchmark. The company-specific mortality is termed the *model* mortality and has the form

$$\mu_k^{model}(x, t) = \exp\left(\beta_1^k r_1(x) + \beta_2^k r_2(x) + \beta_3^k r_3(x)\right) \mu_k^{FSA}(x, t), \quad (2)$$

where r_1 , r_2 and r_3 are fixed regressors specified by the Danish FSA. The β -parameters are estimated by the company and subject to a significance test in which non-significant parameters are set to zero. The estimation is performed each year and is based on the mortality experience of the company's portfolio over the last 5 years.

Appendix A contains a more comprehensive description of the benchmark including details on the estimation and testing procedure leading to the model mortality.

2.1 Proposed longevity stress

The proposed longevity stress is phrased in terms of the Danish FSA longevity benchmark. The form of the stress is chosen to make it easy to apply in practise for Danish life and pension insurance companies using the benchmark. However, despite its simplicity the proposed stress captures the essential structure of longevity risk as we shall see.

The stress has three components: One related to variability in the benchmark level; one related to variability in the benchmark trend; and one related to estimation uncertainty of the company-specific model mortality. We will talk about the first two components as quantifying the *systematic* risk shared by all companies, and the last component as quantifying the *unsystematic* risk due to the randomness of deaths in the portfolio of a specific company.

The systematic longevity stress is obtained by stressing the benchmark level and trend in the following way

$$\tilde{\mu}_k^{FSA}(x, t) = (1 - S_{level}) \mu_k^{FSA}(x, T) (1 - (1 + S_{trend}) R_k(x))^{t-T}. \quad (3)$$

That is, by reducing the current observed level by the factor S_{level} and increasing the future rates of improvement by the factor S_{trend} . The two factors depend neither on gender nor age. However, the impact of S_{trend} in terms of both mortality rates and life expectancy is larger for young ages than for old ages due to the longer time horizon. The stress thereby agrees with the general conception that longevity risk should increase with horizon.

The unsystematic longevity risk is interpreted as the variation in the estimated current level of company-specific mortality from year to year due to the randomness of deaths. It is described by a further reduction of the model mortality on top of the systematic stress above. The total longevity stress taking account of both systematic and unsystematic longevity risk takes the form

$$\tilde{\mu}_k^{model}(x, t) = (1 - S_{Poisson}) \exp\left(\beta_1^k r_1(x) + \beta_2^k r_2(x) + \beta_3^k r_3(x)\right) \tilde{\mu}_k^{FSA}(x, t), \quad (4)$$

where $S_{Poisson} = 2.6/\sqrt{5H}$, and where H is the total expected number of deaths (females and males combined) over the last 5 years in the company's portfolio assuming the FSA benchmark mortality (1). The reduction factor $S_{Poisson}$ depends on neither

gender nor age. Estimation uncertainty arises as a consequence of the underlying assumption that the number of deaths in a portfolio follows a Poisson distribution, hence the subscript.

The value of the systematic and unsystematic stress components calibrated to a stress level of 99.5% on a one-year horizon is shown in Table 1. The interpretation is that by stressing the current benchmark by reducing the initial level by 6% and by increasing the future improvement rates by 6% we obtain a new mortality surface which corresponds in “severity” to a one-in-two-hundred benchmark update. Similarly, the factor $2.6/\sqrt{5H}$ quantifies the change in the overall level of company-specific model mortality which will happen with a probability of one-in-two-hundred due to Poisson variation alone. In the calibration we do not allow the systematic and unsystematic risks to diversity each other, i.e. the two risks are simply aggregated.

Benchmark level stress (S_{level})	6%
Benchmark trend stress (S_{trend})	6%
Unsystematic stress ($S_{Poisson}$)	$2.6/\sqrt{5H}$

Table 1: Systematic and unsystematic 99.5% longevity stress components. The parameter H is the expected number of deaths in the insurance portfolio during a period of five years assuming benchmark mortality.

To exemplify the role played by the unsystematic risk component we show in Table 2 the size of $S_{Poisson}$ for various portfolio sizes. The stress associated with the unsystematic risk is linked to the size of data used for the statistical analysis of company-specific mortality. Note that although the analysis is performed for each gender separately, the unsystematic risk is calculated for the portfolio as a whole, i.e. for the two genders combined. For a portfolio with for example 1,000 expected deaths each year H is 5,000 and the stress becomes 1.6%. As the size of the portfolio (in terms of expected number of deaths) decreases the size of the stress increases. For very small portfolios the size of the stress becomes excessive; if there is only one expected death every five years the stress will in fact exceed 100%! This illustrates that based on the statistical evidence alone we can say virtually nothing about the underlying mortality. For such portfolios, we propose to apply alternative methods and introduce a priori information about the portfolio, for example by using the mortality experience from similar portfolios. However, in the present paper we do not pursue this issue further.

Expected number of deaths, one year	1	10	100	1,000	10,000
Expected number of deaths, 5 years (H)	5	50	500	5,000	50,000
Unsystematic stress ($S_{Poisson}$)	52.0%	16.4%	5.2%	1.6%	0.5%

Table 2: The unsystematic stress for various portfolio sizes. The size of the portfolio is measured by the expected number of deaths assuming benchmark mortality.

The effect of different longevity stress scenarios is illustrated in Table 3. For each gender and age considered the table shows the expected remaining life time in 2012 calculated under the longevity benchmark (1) together with the life expectancy increase for each stress scenario. We consider four different stress scenarios: The uniform stress of 20% from the Solvency 2 standard model; the combined systematic and unsystematic stress for portfolios with H equal to 500 and 5,000, respectively; and the systematic

stress on its own. The latter can be interpreted as the stress for a portfolio so large that the unsystematic risk is negligible.

We see from Table 3 that the proposed systematic stress is much less severe than the uniform 20% stress from the Solvency 2 standard model. In terms of life expectancy the systematic stress is approximately one third of the Solvency 2 standard model stress. We also see that for a portfolio with $H = 5,000$ corresponding to 1,000 expected deaths per year the additional effect of the unsystematic stress is modest, while for a portfolio with $H = 500$ the increase caused by the unsystematic stress is substantial. For the latter case the combined systematic and unsystematic stress is about two thirds the size of the Solvency 2 standard model stress. The portfolio sizes considered are typical for many Danish pension companies.

The size of the proposed stress reflects the fact that the Danish longevity benchmark is updated annually. The stress quantifies the effect of adding data for the most recent year and removing data for the most distant historical year from the data window used to set best-estimate mortality assumptions. The high frequency of updates and the extend of overlap in data both contribute to reduce the variability between consecutive updates, compared to updating assumptions only every five or ten years, say.

Gender	Age	Remaining life time Benchmark	Increase of life time			
			20%	$H=500$	$H=5,000$	Systematic
Males	20	66.3	1.7	1.1	0.8	0.7
	40	45.1	1.8	1.1	0.8	0.7
	60	24.7	1.7	0.9	0.7	0.5
	80	8.5	1.1	0.6	0.4	0.3
Females	20	68.6	1.8	1.2	0.9	0.7
	40	47.5	1.9	1.1	0.8	0.7
	60	27.4	1.8	1.0	0.7	0.6
	80	10.5	1.3	0.7	0.5	0.4

Table 3: Expected remaining life time in 2012 for males and females at various ages under the Danish longevity benchmark, and life expectancy increases for various stresses: Uniform 20% stress; combined systematic and unsystematic stress for portfolios with expected number of deaths during a five year period (H) equal to 500 and 5,000, respectively; systematic stress on its own.

3 Calibration methodology

In the framework of Solvency 2 the goal is to quantify the loss in basic own funds due to longevity that will arise with a probability of 0.5% on a one-year horizon. This can be calculated as the 99.5% percentile in the loss distribution (due to longevity) obtained from a large number of simulated scenarios for a stochastic model describing the dynamics of the company. In the Solvency 2 jargon such a model is termed a (partial) internal model. However, building and maintaining such a model is far from straightforward, and it is therefore desirable to derive a single mortality stress scenario and evaluate the impact on basic own funds for this scenario only. This is the idea of the longevity stress in the Solvency 2 standard model.

The purpose of the work presented in this paper is to derive an alternative stress

taking the specifics of the Danish regulatory regime into account. This is done by the use of a stochastic model calibrated to historical Danish data. The model describes the joint distribution of the three layers of data entering the estimation of mortality rates used by Danish life and pension companies: national data, data for insured lives (sector data), and company-specific data.

We use the model to simulate both the underlying evolution in mortality intensities and the realized death counts for the coming year. For each simulated data set, we calculate the corresponding longevity benchmark by following the Danish FSA longevity benchmark construction algorithm. We thus obtain the longevity benchmark distribution one year ahead from which we calculate the systematic stress.

In principle, we need detailed information on the liability structure to assess for which of the simulated longevity benchmarks the biggest losses would occur. However, rather than making assumptions about a generic company, we make the simplifying, but reasonable, assumption that the loss related to a given cohort is increasing in the cohort's life expectancy. Under this assumption, the task simplifies to that of finding a stress of a simple parametric form reproducing the 99.5% percentile of the life expectancy distribution for each cohort.

When calibrating the unsystematic stress, we make the assumption that the true underlying company-specific mortality intensity is proportional to the sector intensity. However, the constant of proportionality is unknown and needs to be inferred from company data. The constant is reestimated every year and this causes unsystematic variation from year to year in addition to benchmark variation. The unsystematic stress is calibrated such that it corresponds to the 99.5% percentile of (minus) the annual change in the estimate of the proportionality constant.

In our framework the systematic and unsystematic risks are independent by design. It could thus well be argued, that the two sources of risk should be allowed to diversify each other. In reality, however, the risk sources will not be fully independent, since sector data consist of pooled company data. Rather than trying to model the level of dependence, and in order not to underestimate the dependency, we propose a stress, which in effect treats the two risk sources as *fully dependent*; i.e. the systematic and unsystematic risks both enter the proposed stress at their (marginally) calibrated 99.5% level with no diversification discount.

The proposed stress is calibrated under the current Danish regulatory regime. We are exploiting the well-described mechanism for the annual benchmark construction concerning, in particular, choice of data, length of estimation window and smoothing algorithm. The size of the resulting stress reflects the variability in mortality assumptions from year to year under this regime. If, for any reason, the Danish FSA decides to change the regulatory framework concerning longevity, we would need to recalibrate the stress.

3.1 Systematic and unsystematic longevity risk

We have phrased the proposed stress in terms of systematic and unsystematic risks. The systematic risk quantifies the annual variability in the underlying benchmark shared by all Danish life and pension companies, whereas the unsystematic risk quantifies the annual variability in the company-specific component of the company's mortality assumptions. The stress for the former risk source is common to all companies, while the stress for the latter risk source depends on the size of the company's portfolio, as

measured by the number of expected deaths.

In the Solvency 2 directive, quoted in the Introduction, longevity risk is described as the risk of loss resulting from changes in the level, trend, or volatility of mortality rates. The systematic stress can be viewed as a quantification of the risk resulting from changes in the level and trend, and the unsystematic stress can be viewed as a quantification of the risk resulting from volatility of observed (raw) death rates. In this sense the proposed stress complies with the Solvency 2 definition of longevity risk.

In the work presented, we calibrate the risk components within a given regulatory regime using a specific stochastic mortality model. One might argue that the Solvency 2 definition of longevity risk also aims at quantifying the effect of potential changes in regulation or fundamental changes in the view of future mortality arising from e.g. medical breakthroughs. We acknowledge the presence of these “risk” sources, but we also believe that they are very hard to quantify by statistical methods. In our view structural changes like these constitute one aspect of model risk, which is better assessed by varying the underlying assumptions (sensitivity analysis) than by including them in the model itself. This, however, is outside the scope of the present work.

3.2 Terminology and model framework

The benchmark trend and benchmark level are estimated each year by the FSA; the trend is estimated from the mortality experience of the Danish population over a 30 year period, while the level is estimated from a large pool of Danish insured lives over a 5 year period. The model mortality used by the company is also estimated each year. The estimation is based on the mortality experience over a 5 year period of the company’s portfolio. The actual longevity assumptions used by a given company thus depends on the mortality experience of both the Danish population, the sector (pool of insured lives), and its own portfolio.¹

In order to assess the longevity risk faced by a company, we need to model jointly the mortality experience of these three populations. For a given population we denote by $\mu(x, t)$ the true underlying, but unobservable, mortality intensity for people of age x in year t , and we denote by $D(x, t)$ and $E(x, t)$, respectively, the number of deaths and the exposure (“the number of people at risk of dying”) of age x in year t . To distinguish between the different populations we superscript these quantities with N for national data (Danish), S for sector data (insured lives), and C for company data. Further, we make the standard assumption that death counts are Poisson distributed with mean μE . The model is summarized below.

$$\begin{aligned} \text{National data:} \quad & D_k^N(x, t) \sim \text{Poisson}(\mu_k^N(x, t)E_k^N(x, t)). \\ \text{Sector data:} \quad & D_k^S(x, t) \sim \text{Poisson}(\mu_k^S(x, t)E_k^S(x, t)). \\ \text{Company data:} \quad & D_k^C(x, t) \sim \text{Poisson}(\mu_k^C(x, t)E_k^C(x, t)). \end{aligned}$$

The three underlying intensities, μ_k^N , μ_k^S , and μ_k^C , are stochastic and dependent, while death counts are assumed to be independent and Poisson distributed conditioned on the underlying intensities.

¹The procedures for calculating the benchmark trend and benchmark level as well as the company-specific analysis are described in detail on the following homepage of the Danish FSA: <http://www.ftnet.dk/levetider> [in Danish].

One could in principle apply a binomial model for the number of deaths instead of the Poisson distribution. This might appear to be more natural, in particular if the populations are small. However, the number of deaths is binomially distributed only if each individual in an age cell (combination of age x and year t) has the same probability of dying. For persons born in the same year but at different dates, this is not the case. In addition, we note that the underlying data are given in terms of the number of deaths and exposures rather than individual data. In particular, the number of people in each age cell is not available.

We think of the benchmark intensity, μ_k^{FSA} , as an estimate of μ_k^S , and the model intensity μ_k^{model} as an estimate of μ_k^C . Within this framework we can study the effect of the uncertainty in the underlying intensities and death counts on the benchmark and model mortality.

3.3 Longevity stress within the Danish regulatory regime

In order for the longevity stress to be easy to use in practice we need it to conform with the benchmark parametrization, see Appendix A. As the analysis shows, the systematic longevity risk can be adequately captured by a stress of the form (3). That is, by reducing the current observed level by the factor S_{level} and increasing the future rates of improvement by the factor S_{trend} . Similarly, the unsystematic longevity risk is quantified by a reduction of the model mortality by the factor $S_{Poisson}$, which is allowed to depend on the size of the portfolio. The aim of the analysis is hence to quantify the stress factors S_{level} , S_{trend} and $S_{Poisson}$ in (3) and (4).

3.4 Choice of stochastic mortality model

A vast number of stochastic mortality models varying in complexity and area of application have been proposed in the literature. However, the most widely used methodology for making mortality projections is still the method proposed by Lee and Carter (1992). The method has gained widespread popularity due to its simplicity and ease of interpretation. There has been a wealth of applications, see e.g. Tuljapurkar et al. (2000); Booth et al. (2006); Pitacco et al. (2009) and references therein. A number of extensions and improvements have been proposed, e.g. Brouhns et al. (2002); Lee and Miller (2001); Renshaw and Haberman (2006, 2003); de Jong and Tickle (2006); Currie et al. (2004); Cairns et al. (2006), but the original Lee-Carter method still serves as the industry standard. The Danish FSA longevity benchmark is essentially a Lee-Carter model forecast, although the estimation algorithm differs.

In this paper we also employ the Lee-Carter model as a workhorse. Specifically, we will use the Poisson variant of Brouhns et al. (2002) to model the evolution in Danish mortality. Sector and company mortality intensities will be assumed to move in parallel with Danish mortality, i.e. no additional randomness will be introduced. In some sense these simplifying assumptions represent a worst case scenario for risk assessment since they preclude diversification benefits. On the other hand several authors have argued that the rigid structure of the Lee-Carter model might lead to excessively narrow forecasting confidence intervals, see e.g. Li et al. (2009). A description of the chosen model and comments on the model's suitability can be found in Appendix B.

In principle, we could have used any of the models proposed for joint modeling of two or more related populations, e.g. Li and Lee (2005); Plat (2009); Biatat and Currie (2010); Cairns et al. (2011); Li and Hardy (2011); Jarner and Kryger (2011). In one

form or another all of these models assume that the joint evolution is given by a shared trend in combination with components pertaining to each individual population. The aim of this structure is to ensure coherence for long-term projections while allowing for short term deviations from the common trend. However, as our focus is on risk assessment on a one-year horizon and not on long-term mortality forecasting, we are primarily concerned with not underestimating the short term dependence.

4 Systematic longevity risk

The analysis of the systematic longevity risk proceeds as follows. Firstly, we fit a Poisson Lee-Carter stochastic mortality model for Danish national mortality. We then use this model to generate the joint distribution of national and sector data one year ahead. We assume that the two populations evolve in parallel to guard us from underestimating the double impact of simultaneous changes in both benchmark trend and level. Secondly, we calculate the (remaining) life expectancy distribution for each age and gender one year ahead. Thirdly, we calibrate a longevity stress of the form (3) to reproduce the 99.5% quantiles of these life expectancy distributions. The calibration is initially done for each gender separately, and afterwards averaged to arrive at a unisex stress.

4.1 Simulating the longevity benchmark

The benchmark published by the Danish FSA in August 2012 has reference year $T = 2011$. The trend is estimated from Danish data available on the Human Mortality Database (HMD) at www.mortality.org. The level for the sector is estimated from data for a pool of insured lives gathered by The Danish Centre of Health and Insurance.

The trend is estimated on the basis of Danish data for the 30 year period 1980–2009, while the level is estimated from sector data for the 5 year period 2007–2011. We let $T_{trend} = 2009$ and $T_{level} = 2011$ denote the last data year used in estimation of respectively the trend and the level.

To describe the evolution in Danish mortality, we use the Poisson variant of the classical Lee-Carter model. We fit this model to Danish data from 1980-2009 for ages 0-105. Selected parameter estimates and plots of fit can be found in Appendix B.

With this model at hand, we are able to simulate from the distribution of the Danish mortality one year ahead, i.e. we can simulate from the distribution of $\mu_k^N(x, T_{trend} + 1)$ for ages $x = 0, \dots, 105$. In order to calculate the benchmark level, we also need to simulate the sector mortality one year ahead. We assume that the sector experiences the same rates of improvement as in the national data. Specifically, we set

$$\mu_k^S(x, T_{level} + 1) = \mu_k^{FSA}(x, T_{level}) \frac{\mu_k^N(x, T_{trend} + 1)}{\mu_k^N(x, T_{trend})}, \quad (5)$$

where the numerator in the last term is the simulated value from the Poisson Lee-Carter model, and the denominator is the fitted value of the model to the Danish data at the last data year.

Assumption (5) represents in some sense the worst case, since improvements affect both benchmark trend and level simultaneously. On the other hand it can be argued that rates of improvements for insured lives might be higher than for the population at large. Given the short time series of sector data, we cannot verify this hypothesis. Also, if that were indeed to be the case, it would be reasonable to assume a less than

perfect correlation between national and sector data (which would offset this effect to some degree). Overall, we find that the chosen model captures the main effects of the longevity risk in a simple way.

Having the joint distribution of national and sector intensities one year ahead, we then simulate the actual number of deaths as independent Poisson variates. As we do not yet know the exposures, we use the exposures for the last data year as proxy. That is, we simulate national and sector deaths for ages $x = 0, \dots, 105$ as²

$$D_k^N(x, T_{trend} + 1) \sim \text{Poisson}(\mu_k^N(x, T_{trend} + 1)E_k^N(x, T_{trend})), \quad (6)$$

$$D_k^S(x, T_{level} + 1) \sim \text{Poisson}(\mu_k^S(x, T_{level} + 1)E_k^S(x, T_{level})). \quad (7)$$

Finally, we use the FSA algorithm for estimating the benchmark trend and level. The trend is estimated from Danish data for the 30 year period from 1981 to $T_{trend} + 1 = 2010$, where we use the historic data for the first 29 years and the simulated data for the last year. Similarly, we estimate the level from sector data for the 5 year period from 2008 to $T_{level} + 1 = 2012$, where we use historic data for the first 4 years and simulated data for the last year.

To sum up, we generate a new benchmark by:

- Simulate $\mu_k^N(x, T_{trend} + 1)$ for ages $x = 0, \dots, 105$ from the Poisson Lee-Carter model
- Calculate $\mu_k^S(x, T_{level} + 1)$ for ages $x = 0, \dots, 105$ using relation (5) with the simulated $\mu_k^N(x, T_{trend} + 1)$
- Simulate national and sector data using the simulated intensities and the last known exposures by (6)–(7)
- Compute a new benchmark trend from 29 years of historic data and the simulated national data for year $T_{trend} + 1$
- Compute a new benchmark level from 4 years of historic data and the simulated sector data for year $T_{level} + 1$

We repeat this procedure 10,000 times which gives us the benchmark distribution with reference year $T_{level} + 1 = 2012$.

4.2 Calibration of systematic longevity stress

The next step in the analysis is to use the benchmark distribution to find a longevity stress corresponding to a 99.5% confidence level. Here, we are faced with the usual problem when dealing with multidimensional distributions: it is not obvious which of the 10,000 simulated benchmarks belong to the “worst” 0.5%. One way to proceed would be to calculate the marginal 99.5% quantiles separately for the trend and for the level for each age (and each gender). This however clearly exaggerates the stress since the probability that all of these events happen at once is less than 0.5%.

Since we are looking at longevity risk, it is safe to assume that the worst benchmarks (in terms of economically most expensive for the companies) are those with the highest

²Technical note. For ages 0-25 the benchmark level is based on Danish data. We therefore use the exposure for Denmark from year T_{trend} for these ages when simulating sector data for year $T_{level} + 1$.

(remaining) life expectancy. We will therefore calibrate the longevity stress of the form (3) such that the life expectancy of $\tilde{\mu}_k^{FSA}$ equals the 99.5% percentile in the life expectancy distribution of the simulated benchmarks for a range of (economically) relevant age groups.

First we calculate for each of the simulated benchmarks (pair of trend and level) the cohort life expectancy for ages $x = 0, \dots, 100$ in 2012

$$e_k^i(x, 2012) = e(x, 2012; \mu_k^i, R_k^i) \text{ for } i = 1, \dots, 10,000, \quad (8)$$

where μ_k^i and R_k^i denote respectively the simulated benchmark level and trend in 2012 for gender k . We now have the life expectancy distribution in 2012 for each age and gender. In each of these distributions we calculate the 99.5% percentile, $q_k(x, 2012)$.

We also calculate the life expectancy in 2012 under the current FSA benchmark for ages $x = 0, \dots, 100$

$$e_k^{FSA}(x, 2012) = e(x, 2012; \mu_k^{FSA}(\cdot, 2011), R_k), \quad (9)$$

where $\mu_k^{FSA}(\cdot, 2011)$ and R_k denote respectively the current benchmark level and trend for gender k . The difference,

$$\Delta_k(x) = q_k(x, 2012) - e_k^{FSA}(x, 2012), \quad (10)$$

can be interpreted as the increase in the benchmark life expectancy for age x , which can happen with probability 0.5% by the next annual update of the benchmark. By calculating both life expectancies in 2012 we only look at the increase in longevity in excess of what is already anticipated by the current benchmark.

Figure 1 shows the life expectancy increase Δ_k (solid black line) together with the life expectancy increase resulting from a uniform reduction in the current benchmark of 5%, 10%, 15% and 20% (dotted lines). The figure also features two dashed green and blue lines to which we will return shortly. We can see from the figure that Δ_k corresponds to a uniform stress in the range from 5% to 10%. We also note that compared to a uniform stress the profile of Δ_k is different. The life expectancy increase under Δ_k is more rapidly decreasing with age which seems to conform better with intuition than a uniform stress. (The peculiar bump for old women is likely due to parametric extrapolation used after age 80 when constructing the benchmark level. The bump is also visible for men, but less pronounced.)

The last step in the analysis is to calibrate a longevity stress of the form (3) such that the “stressed” life expectancy corresponds to the 99.5% quantiles $q_k(x, 2012)$. To formalize the calibration procedure we define the difference

$$\tilde{\Delta}_k(x; S_{trend}, S_{level}) = \tilde{e}_k^{FSA}(x, 2012; S_{trend}, S_{level}) - e_k^{FSA}(x, 2012), \quad (11)$$

where $\tilde{e}_k^{FSA}(x, 2012; S_{trend}, S_{level})$ denotes the life expectancy in year 2012 for age x and gender k calculated from the stressed benchmark (3). Using this terminology the dotted lines in Figure 1 correspond to $\tilde{\Delta}_k(x; 0, S_{level})$ for $S_{level}=5\%$, 10%, 15% and 20%.

To find the combination of S_{trend} and S_{level} we minimize the squared distance

$$\sum_{x=x_{min}}^{x_{max}} \left(\tilde{\Delta}_k(x; S_{trend}, S_{level}) - \Delta_k(x) \right)^2 \quad (12)$$

for the economically relevant age span $(x_{min}, x_{max}) = (30, 90)$. We minimize (12) over S_{trend} and S_{level} in steps of 0.5%. The gender-specific calibrated values are shown in Table 4. Note that the stress is of the same magnitude for females and males, with the stress for females being 1%-point higher for both parameters. The table also contains a column labeled “Unisex” of the average values.

Parameter	Females	Males	Unisex
S_{trend}	6.5%	5.5%	6.0%
S_{level}	6.5%	5.5%	6.0%

Table 4: Calibrated values of gender-specific and unisex parameters S_{trend} and S_{level} .

The green and blue dashed lines in Figure 1 show the life expectancy increase for the calibrated gender-specific and unisex values, respectively. That is, the green line is $\tilde{\Delta}_k(x; 6.5\%, 6.5\%)$ for females and $\tilde{\Delta}_k(x; 5.5\%, 5.5\%)$ for males, while the blue line is $\tilde{\Delta}_k(x; 6.0\%, 6.0\%)$ for both females and males. The fit of the green line is remarkably good from age 40 to 80 for both genders, while deviating somewhat above age 80 in particular for females. This demonstrates that a stress of the form (3) is able to capture in a simple way the systematic longevity risk for the most economically important age groups.

The higher stress for females than for males is driven by a more volatile development of female mortality than of male mortality in the historic period used to estimate the underlying Lee-Carter model. It might be argued that we have no reason to believe that this is an intrinsic feature of female mortality, and that looking forward it would be more reasonable to assume the same level of uncertainty for both genders. This can be obtained by using the unisex stress in Table 4. The unisex stress is shown as the blue dashed line in Figure 1. By construction, the unisex stress is slightly lower for females and slightly higher for males than the gender-specific stress, but it still provides a very good description of the systematic longevity risk overall.




Figure 1: Increase in life expectancy for ages 0–100 in 2012 relative to current FSA benchmark for females (top) and males (bottom): 99.5% percentile in simulated life expectancy distribution (solid black), gender-specific stress (dashed green) and unisex stress (dashed blue), see Table 4 for values of S_{trend} and S_{level} . Also shown, life expectancy increase for stress with $S_{trend}=0$ and $S_{level}=5\%$, 10% , 15% and 20% (dotted).

5 Unsystematic longevity risk

We interpret unsystematic longevity risk as the risk associated with the variation in the estimated current level of company-specific mortality from year to year due to the randomness of deaths (Poisson variation). Like in the analysis of systematic longevity risk we analyze the *change* in assumptions from year to year taking account of the overlap in data used in consecutive years. We derive the unsystematic longevity stress by theoretical means based on the general framework of Section 3.2.

Under the current Danish regulatory regime companies can only use a model mortality differing from the benchmark if the deviations are statistically significant. This implies that small companies will tend to use the benchmark every year and therefore — in the strict sense of the definition above — faces *no* unsystematic longevity risk. This is clearly not the intention of the regulators. In order to arrive at a stress which is meaningful for all companies regardless of size, we therefore disregard this aspect of the regulation, and perform the analysis as if all companies used an estimated model mortality. In other words, we quantify the change in the annual estimate of the level of company-specific mortality whether or not this is actually used. This gives the true picture of uncertainty in a given population, and gives rise to a stress that scales (inversely) with the size of the population.

5.1 Company data

We use the following model for a company with a single parameter α , independent of age and gender, to measure the overall excess mortality relative to sector mortality

$$D_k^C(x, t) \sim \text{Poisson}(\mu_k^C(x, t)E_k^C(x, t)), \text{ where} \quad (13)$$

$$\mu_k^C(x, t) = \alpha\mu_k^S(x, t). \quad (14)$$

This model allows us to derive a simple formula for unsystematic longevity risk which depends only on a single summary statistic of the member base. The model, although simple, captures the effect of the Poisson variation in a succinct way.

In the actual regulation, companies are required to perform a Poisson regression with a specific parametrization as described in Section A.1. It would be possible to take this parametric form as our starting point and study the distribution of the β -estimates under this model. The results, however, will depend on the entire age-profile of the company which leaves us with little hope of arriving at a general formula. On the other hand, companies may well want to perform this analysis on their own.

The rationale behind linking the company specific mortality to the sector mortality is the same as in Section 4.1. We assume that the mortality of the member base of the company evolves in parallel to the sector mortality such that the unsystematic risk is indeed an add-on risk on top of the systematic risk. As in Section 4.1 it could well be argued that the evolution of company specific mortality may differ from that of the sector, but then it would also be reasonable to assume a less than perfect dependency and thereby a diversification discount. Like before we find that the chosen model captures the main effect of unsystematic longevity risk in a simple way.

5.2 Estimation uncertainty

Before turning to the variation in the estimate of α for two successive years we first look at the estimate of α itself. Assuming μ_k^S is known, the maximum likelihood estimator for α is given by

$$\hat{\alpha} = \frac{\sum_{x,t,k} D_k^C(x,t)}{H}, \quad \text{with } H = \sum_{x,t,k} \mu_k^S(x,t) E_k^C(x,t), \quad (15)$$

where the summation extends over the same ages and years used to estimate the model mortality. Under Danish regulation, the model mortality is estimated on the basis of data from a 5 year period, and hence we will assume an estimation period of 5 years. The quantity H is the number of deaths expected in the data period, had the company mortality been equal to the sector mortality.

One could also consider to allow the level of excess mortality to depend on sex, i.e. to have α_k instead of α in (14). At first sight this would perhaps appear more natural considering that the company-specific mortality is estimated for each sex separately. However, if we assume that the economic impact is, at least approximately, proportional to the estimated level multiplied by the size of the population for which that level applies, it follows that the combined effect is proportional to the estimated level of *average* excess mortality multiplied by the size of the total population. Therefore, we can assess the combined effect in the simpler one-parameter model above.

The argument can be made more mathematically rigorous as follows. In the two-parameter model, $\hat{\alpha}_F$ and $\hat{\alpha}_M$ are estimated by (15), where the summations extend only over t and x . Denoting by H_F and H_M the denominator in the fraction defining $\hat{\alpha}_F$ and $\hat{\alpha}_M$ respectively, we have the following relation between the estimators: $\hat{\alpha}H = \hat{\alpha}_F H_F + \hat{\alpha}_M H_M$. Hence, if the combined effect is a function of $\hat{\alpha}_F H_F + \hat{\alpha}_M H_M$ it is also a function of $\hat{\alpha}H$, and we need only consider the one-parameter model to assess the risk.

Under the one-parameter model we have that

$$\sum_{x,t,k} D_k^C(x,t) \sim \text{Poisson}(\alpha H), \quad (16)$$

and thereby

$$\hat{\alpha} \sim \frac{\text{Poisson}(\alpha H)}{H}, \quad \text{E}(\hat{\alpha}) = \alpha, \quad \text{Var}(\hat{\alpha}) = \frac{\alpha}{H}, \quad \text{Std}(\hat{\alpha}) = \frac{\sqrt{\alpha}}{\sqrt{H}}. \quad (17)$$

We see that the standard deviation of the estimator depends on the true value of α (and H). Since the true value of α is unknown, a possible solution would be to use the estimated value of α instead. However, for small companies the estimation uncertainty is substantial and this approach will therefore lead to variation in the perceived uncertainty from year to year.

A more stable estimate of the standard deviation, which is applicable to all companies regardless of size, can be obtained by noting that α and thereby $\sqrt{\alpha}$ will (on average) be close to 1, since α measures the level of mortality for a specific group of insured lives relative to other insured lives. We can therefore approximate the standard deviation by

$$\text{Std}(\hat{\alpha}) \sim \frac{1}{\sqrt{H}}. \quad (18)$$

Note that the approximation does not imply that we assume α to equal 1. The approximation can be used for all α , but of course it is more precise the closer α is to 1.

5.3 Uncertainty of successive estimates

We now extend the analysis of the preceding section to study the variability in the change of the α estimate from one year to the next. To present the analysis we define the partial sums

$$D_t = \sum_{x,k} D_k^C(x, t) \quad \text{and} \quad H_t = \sum_{x,k} \mu_k^S(x, t) E_k^C(x, t), \quad (19)$$

which represent actual and expected number of deaths during year t . We also need to distinguish between different estimation periods, and for this purpose we will subscript the estimator with the last data year of the estimation period, i.e. we write $\hat{\alpha}_T$ for the estimate of α based on data from the 5 year period from $T - 4$ to T . In the notation introduced above we have

$$\hat{\alpha}_T = \frac{D_{T-4} + \dots + D_T}{H_{T-4} + \dots + H_T}, \quad (20)$$

and for the ratio of two successive estimators we have

$$\frac{\hat{\alpha}_{T+1}}{\hat{\alpha}_T} = \frac{H_{T-4} + \dots + H_T}{H_{T-3} + \dots + H_{T+1}} \left(\frac{D_{T-3} + \dots + D_T}{D_{T-4} + \dots + D_T} + \frac{D_{T+1}}{D_{T-4} + \dots + D_T} \right). \quad (21)$$

We are interested in the variability of (21) caused by inclusion of the new data D_{T+1} . In statistical terms we want to find the standard deviation of (21) conditioned on data up to and including time T

$$\begin{aligned} \text{Std} \left(\frac{\hat{\alpha}_{T+1}}{\hat{\alpha}_T} \middle| \{D_t\}_{t \leq T} \right) &= \frac{H_{T-4} + \dots + H_T}{H_{T-3} + \dots + H_{T+1}} \text{Std} \left(\frac{D_{T+1}}{D_{T-4} + \dots + D_T} \middle| \{D_t\}_{t \leq T} \right) \\ &\approx \frac{\text{Std}(D_{T+1})}{\alpha (H_{T-4} + \dots + H_T)} \\ &= \frac{\sqrt{\alpha H_{T+1}}}{\alpha (H_{T-4} + \dots + H_T)} \\ &\approx \frac{1}{\sqrt{\alpha} \sqrt{5H}} \\ &\approx \frac{1}{\sqrt{5H}}, \end{aligned} \quad (22)$$

where $H = H_{T-4} + \dots + H_T$ (as in Section 5.2). In the first approximation above, $D_{T-4} + \dots + D_T$ is replaced by its expectation $\alpha (H_{T-4} + \dots + H_T)$, and the H -ratio is replaced by 1. The latter rests on the assumption that the exposure and thereby H_t is almost constant over time. This assumption is also used in the second approximation to replace H_{T+1} by $H/5$. Finally, we approximate $\sqrt{\alpha}$ by 1, by the same arguments as in Section 5.2, to arrive at a stable, easily implementable formula.

Comparing formulas (18) and (22) we see that the effect of reusing 4 years of data in successive estimations is to reduce the standard deviation of the ratio of to successive α estimates by a factor of $\sqrt{5}$ relative to the standard deviation of α itself.

5.4 Calibration of unsystematic longevity stress

The theoretical developments above show that the standard deviation of the change in the estimate of the level of company mortality from one year to the next attributable to unsystematic (Poisson) variation can be approximated by $1/\sqrt{5H}$. From this we conclude that a 99.5% unsystematic longevity stress can be obtained by setting $S_{Poisson}$ of (4) equal to

$$S_{Poisson} = \frac{2.6}{\sqrt{5H}}. \quad (23)$$

The stress is obtained by utilizing that the distribution of the ratio follows a (scaled and translated) Poisson distribution which is well approximated by a normal distribution. The stress $S_{Poisson}$ corresponds to the 99.5% percentile in the approximating normal distribution.

As seen in the derivations in Sections 5.2–5.3 leading to (23), we approximate α by 1 to achieve a simple and stable formula. We justified the approximation by the observation that on average company mortality ought to be close to sector mortality. On the other hand, for a specific company the value of α will deviate to a smaller or larger extent from 1, and the approximation will therefore (from a theoretical perspective) introduce a bias in the estimate of the unsystematic risk of the given company. For $\alpha = 0.8$ and $\alpha = 0.6$ the stress in (23) should be multiplied by respectively 1.12 and 1.29, while for $\alpha = 1.2$ and $\alpha = 1.4$ the stress should be multiplied by 0.91 and 0.85.

If instead we approximate α by its estimator $\hat{\alpha}$, we can reduce the bias at the expense of introducing variability in the estimate of the risk. This is the usual variance-bias tradeoff. For large populations the use of $\hat{\alpha}$ will tend to improve the estimate of the risk, while for small populations the use of $\hat{\alpha}$ will tend to degrade the estimate of the risk. Overall we judge that the benefits of simplicity and stability outweigh the potential bias introduced by the approximation.

In the theoretical derivations, H is expressed in terms of the underlying sector mortality μ_k^S . In practice we will use the FSA benchmark instead and calculate H by the formula

$$H = \sum_{x,t,k} \mu_k^{FSA}(x,t) E_k^C(x,t), \quad (24)$$

where the summation extends over the same 5 year period used to estimate the model mortality of the company.

The magnitude of the stress depends on the size of the population through H . Table 5 shows the stress for various values of H ranging from a very small population with only 1 expected death per year ($H = 5$) to a very large population with 10,000 expected deaths per year ($H = 50,000$). The expected number of deaths are calculated under the FSA benchmark. We note that for very small populations the size of the stress might be excessive, and one might consider to introduce a cap.

Expected number of deaths per year	1	10	100	1,000	10,000
Expected number of deaths in 5 years (H)	5	50	500	5,000	50,000
Unsystematic longevity stress ($S_{Poisson}$)	52.0%	16.4%	5.2%	1.6%	0.5%

Table 5: Unsystematic longevity stress, $S_{Poisson}$ for various population sizes. The size of the population is measured by the expected number of deaths under the FSA benchmark mortality.

Let us finally illustrate the combined effect of the systematic and unsystematic longevity stress. We illustrate the effect by the increase in life expectancy relative to the current benchmark for different population sizes. To make the results comparable to those of Section 4 we use a company with all β 's equal to zero, and we use the unisex systematic stress derived in Section 4.2 with $S_{trend} = S_{level} = 6.0\%$. In Figure 2, the life expectancy increase in 2012 for the systematic stress relative to the current benchmark is shown as the solid blue line (as in Figure 1).

The additional effect of including unsystematic risk for populations with respectively 10, 100 and 1,000 expected number of deaths per year are shown as dashed blue lines. For comparison the figure also shows the life expectancy increase for a uniform reduction in the benchmark of 5%, 10%, 15% and 20% as dotted black lines (as in Figure 1). For a population with only 10 expected deaths per year (uppermost dashed blue line) the combined effect is more severe than a uniform reduction of 20%, while for a population with 1,000 expected deaths per year the combined effect corresponds to a uniform reduction of approximately 10%.




Figure 2: Increase in life expectancy for ages 0–100 in 2012 relative to current FSA benchmark for females (top) and males (bottom): unisex stress without unsystematic risk (solid blue) and with unsystematic risk (dashed blue) for populations with respectively 10, 100 and 1,000 expected number of deaths per year. Also shown, life expectancy increase for stress with $S_{Poisson} = S_{trend} = 0$ and $S_{level} = 5\%, 10\%, 15\%$ and 20% (dotted).

6 Refinements of the unsystematic longevity stress

The unsystematic longevity stress derived in the previous section rests on a number of simplifying assumptions. In this section we consider various refined approaches relaxing some of these assumptions. The alternative approaches generally, but not always, come at the price of greater complexity.

We first pursue the idea mentioned in Section 5.4 of approximating α by its estimator $\hat{\alpha}$ (instead of by 1). Substituting α by its estimate $\hat{\alpha}_T$ in the last approximation leading to (22) of Section 5.3 yields an alternative approximation to the standard deviation of two successive estimators

$$\text{Std} \left(\frac{\hat{\alpha}_{T+1}}{\hat{\alpha}_T} \middle| \{D_t\}_{t \leq T} \right) \approx \frac{1}{\sqrt{5D}}, \quad (25)$$

where D denotes the observed, total number of deaths in the portfolio over the 5 year period used to estimate the model mortality of the company. From this we obtain the following alternative expression for the unsystematic longevity stress

$$S_{Poisson} = \frac{2.6}{\sqrt{5D}}. \quad (26)$$

The only difference to formula (23) is that the expected number of deaths under the benchmark is replaced by the observed number of deaths.

The purpose of the alternative expression is to reduce bias of the stress. Table 6 shows the probability of α being worse approximated by $\hat{\alpha}$ than by 1, i.e. the probability that bias is *not* reduced. Only for very small portfolios or values of α close to 1 is the risk of an unintentional increase in bias substantial, but generally bias will indeed be reduced. For portfolios with, say, 100 deaths or more per year (26) therefore offers an appealing alternative to the stress of Section 5.4.

Size of portfolio	Mortality level relative to benchmark (α)							
	0.6	0.7	0.8	0.9	1.1	1.2	1.3	1.4
$H = 5$	13%	28%	45%	64%	67%	54%	43%	34%
$H = 50$	0%	1%	10%	41%	46%	17%	5%	1%
$H = 500$	0%	0%	0%	2%	3%	0%	0%	0%

Table 6: The table shows the probability of α being worse approximated by $\hat{\alpha}$ than by 1 for different portfolio sizes and levels of mortality relative to benchmark, i.e. $P(|\hat{\alpha} - \alpha| > |1 - \alpha|)$ where $\hat{\alpha} \sim \text{Poisson}(\alpha H)/H$.

6.1 Theoretical elaborations

In the derivation of the unsystematic longevity stress we implicitly assumed that the (conditional) expectation of the ratio of two successive estimates of α was 1. More generally, we could replace the term $1 - S_{Poisson}$ in (4) representing the 0.5% percentile in the ratio distribution by a term of the form

$$\mathbb{E} \left(\frac{\hat{\alpha}_{T+1}}{\hat{\alpha}_T} \middle| \{D_t\}_{t \leq T} \right) - 2.6 \text{Std} \left(\frac{\hat{\alpha}_{T+1}}{\hat{\alpha}_T} \middle| \{D_t\}_{t \leq T} \right) = e_{T,T+1} - 2.6sd_{T,T+1}, \quad (27)$$

and use more elaborate approximations to the conditional expectation ($e_{T,T+1}$) and conditional standard deviation ($sd_{T,T+1}$).

Using the notation and framework of Sections 5.1–5.3 we have from (21)

$$e_{T,T+1} = \frac{H_{T-4} + \dots + H_T}{H_{T-3} + \dots + H_{T+1}} \left(\frac{D_{T-3} + \dots + D_T}{D_{T-4} + \dots + D_T} + \frac{\alpha H_{T+1}}{D_{T-4} + \dots + D_T} \right),$$

$$sd_{T,T+1} = \frac{H_{T-4} + \dots + H_T}{H_{T-3} + \dots + H_{T+1}} \frac{\sqrt{\alpha H_{T+1}}}{D_{T-4} + \dots + D_T},$$

and hence we only need to approximate the unknown value of α to be able to calculate a stress of the form (27). In the spirit of using as few approximations as possible we substitute α by its estimate $\hat{\alpha}_T$ yielding

$$e_{T,T+1} \approx \frac{H_{T-4} + \dots + H_T}{H_{T-3} + \dots + H_{T+1}} \frac{D_{T-3} + \dots + D_T}{D_{T-4} + \dots + D_T} + \frac{H_{T+1}}{H_{T-3} + \dots + H_{T+1}}, \quad (28)$$

$$sd_{T,T+1} \approx \frac{\sqrt{H_{T+1}} \sqrt{H_{T-4} + \dots + H_T}}{H_{T-3} + \dots + H_{T+1}} \frac{1}{\sqrt{D_{T-4} + \dots + D_T}}. \quad (29)$$

This, in a sense, represents the best we can do within the theoretical framework. If we are willing to make the additional assumption that H_t is constant over time we retrieve the approximation (25) for the standard deviation while the expression for the mean simplifies to

$$e_{T,T+1} \approx \frac{D_{T-3} + \dots + D_T}{D_{T-4} + \dots + D_T} + \frac{1}{5}. \quad (30)$$

Note that the mean will generally be different from one, unless D_{T-4} accounts for exactly one fifth of the total number of deaths over the past 5 years.

Relaxing the assumptions is instructive and in principle the risk is more precisely assessed by using formula (27) with the mean and standard deviation approximated by (28) and (29), or (30) and (25), than by the simpler stress of Section 5.4. However, this is only true in a strictly mathematical sense within the given theoretical model. Further, the stress of Section 5.4 has the benefit of being more robust since it does not rely on realized death counts. It also better serves the desired objectives of simplicity and stability and we therefore maintain this as our main recommendation.

6.2 Simulation approach

The unsystematic longevity stress of Section 5.4 was derived in a stylized theoretical model. This made it possible to derive a general formula describing the unsystematic risk as a function of a single parameter. The main advantage of this approach is the ease of application of the resulting formula, but the disadvantage is that portfolio-specific characteristics other than the total number of (expected) deaths are ignored as is the estimation and testing procedure used to determine the β -parameters of the actual model mortality.

An alternative, tailored approach to assessing unsystematic risk is to simulate the model mortality on a one-year horizon, in much the same way as the longevity benchmark was simulated in Section 4.1, and calculate the value of liabilities for each simulated mortality. Specifically, we could calculate the unsystematic risk by the following steps

- Let μ_k^{model} denote the current model mortality and assume it is estimated on the basis of company data for the 5-year period from $T - 4$ to T .

- Simulate death counts for year $T + 1$, $D_k^C(x, T + 1)$, from independent Poisson-distributions with mean $\mu_k^{model}(x, T + 1)E_k^C(x, T)$. Note that we assume that the exposures, $E_k^C(x, T + 1)$, for year $T + 1$ are unknown, and we therefore use the known exposures from year T as substitutes.
- Perform the estimation and testing procedure described in Section A.1 with 4 years of historic data and the simulated data to find the β -parameters of the new model mortality (relative to the current longevity benchmark).
- Calculate the value of liabilities for the simulated, new model mortality.
- Repeat the simulation, say, 10,000 times and compute the unsystematic risk as the 99.5% percentile in the simulated liability value distribution with the current value of liabilities subtracted.

For small portfolios the hypothesis that the model mortality equals the benchmark will typically be accepted and the unsystematic risk calculated by the above procedure will therefore be close to zero. In this situation the risk might be more properly assessed by omitting the tests from the model mortality construction algorithm. Similar remarks were made in the opening paragraphs of Section 5.

The simulation approach provides a more accurate, company-specific assessment of the unsystematic risk than the theoretical stress of Section 5.4. It is of course more time consuming to implement and run, but even companies opting for the theoretical stress might find it worth the effort to perform the simulation based analysis at least once to verify the appropriateness of the theoretical stress.

Acknowledgements

Much of the work presented in this paper originates from a working group consisting of Kristoffer Andre Bork (chairman), Ivan Toftegaard Carlsen, Gregers Østervig Frank, Kristian Hasløv, Søren Fiig Jarner, Lars Sommer Hansen, Aage Møller Holst, Nichlas Abel Korsgaard, Liselotte Milting, Thomas Møller and Merete Lykke Rasmussen, under the Danish Society of Actuaries, which was active during the Summer of 2012. In addition to the members of the working group we would also like to thank Lars Hougaard Nielsen, Kristian Buchardt, Esben Masotti Kryger and Finn Larsen for valuable comments and suggestions improving the paper.

A Danish FSA longevity benchmark

The Danish FSA benchmark consists of annual rates of improvements of age- and gender-specific mortality intensities and a current observed level of mortality of insured lives. We refer to the former as the (benchmark) *trend* and the latter as the (benchmark) *level*.

The gender-specific, benchmark mortality intensity for a person of age x in year t of gender k (F for females, M for males) takes the form

$$\mu_k^{FSA}(x, t) = \mu_k^{FSA}(x, T) (1 - R_k(x))^{t-T}, \quad (31)$$

where T is the reference year for the benchmark, i.e. the year of the current observed level of mortality, and x ranges from 0 to 110. The Danish FSA provides the level, $\mu_k^{FSA}(x, T)$, and the trend, $R_k(x)$, from which the users of the benchmark can construct the full benchmark intensity surface by use of formula (31). The benchmark is updated annually by the Danish FSA. At the time of writing the latest benchmark, published in August 2012, has reference year $T = 2011$.

The reduction factors, R_k , are estimated from the latest available 30-year period of Danish mortality data at the Human Mortality Database (www.mortality.org), i.e. the period 1980–2009 for the current benchmark. The estimation is based on the empirical (crude) death rates obtained as the ratio of deaths to exposure.

First, a logistic curve is fitted for ages 90 to 110 for each year and gender and the empirical death rates for age 100 and above are replaced by their fitted values. This is done in order to reduce the fluctuations in old age death rates due to missing or sparse data. Second, for each age and gender the raw annual reduction factor is obtained by log-linear regression of the death rates (the death rate is set to 10^{-7} for cells with no deaths). Third, the R_k 's are calculated by smoothing the raw reduction factors and maximizing with zero. For ages 40 to 80 the R_k 's are in the order of 1 to 2 per cent.

The benchmark level, $\mu_k^{FSA}(\cdot, T)$, is estimated from a 5-year period of mortality data for Danish insured lives provided by a long list of Danish life and pension companies. The current benchmark level is based on a dataset for the period 2007–2011 containing approximately one third of the entire Danish population. Data are collected by the Danish Centre for Health and Insurance.

The estimation of the benchmark level uses the same methodology as the estimation of the benchmark trend. First, a logistic curve is fitted for ages 80 to 110 for each year and gender and the empirical death rates for age 90 and above are replaced by their fitted values. Second, for each age and gender a log-linear regression of the death rates is performed and the fitted value at the final year ($T = 2011$) is recorded. Since the dataset contains primarily adults national data from HMD are used for ages 0–25. Third, the fitted values for the final year are smoothed to obtain the benchmark level.

A.1 Company-specific mortality

The benchmark described above serves as the default mortality basis for calculating the value of liabilities for Danish life and pension companies. This basis must be used unless the mortality experience of the company shows statistically significant differences to the benchmark. In this case the benchmark level should be offset to reflect the company-specific level. The benchmark reduction factors, however, apply to all companies.

The company-specific mortality is termed the *model* mortality and has the form³

$$\mu_k^{model}(x, t) = \exp\left(\beta_1^k r_1(x) + \beta_2^k r_2(x) + \beta_3^k r_3(x)\right) \mu_k^{FSA}(x, t), \quad (32)$$

where the regressors are given by

$$r_m(x) = \begin{cases} 1 & \text{for } x \leq x_{m-1}, \\ (x_m - x)/(x_m - x_{m-1}) & \text{for } x_{m-1} < x < x_m, \\ 0 & \text{for } x \geq x_m, \end{cases} \quad (33)$$

for $m = 1, 2, 3$ and $(x_0, x_1, x_2, x_3) = (40, 60, 80, 100)$. The β -parameters are estimated by the company and subject to a sequential test procedure in which non-significant parameters are set to zero. Note that since both the benchmark mortality and the regressors are continuous the model mortality will also be continuous. However, the overlap of neighboring regressors imply that the corresponding β -coefficients will be (highly) correlated. Also note that the model mortality is always equal to the benchmark mortality from age 100 onwards.

The analysis of the model mortality is done within a Poisson model for company data over the last 5 years

$$D_k^C(x, t) \sim \text{Poisson}\left(\mu_k^{model}(x, t)E_k^C(x, t)\right), \quad (34)$$

where D_k^C and E_k^C denote, respectively, the number of deaths and exposure, and μ_k^{model} is given by (32). The death counts are assumed independent. For the purpose of the analysis the Danish FSA publishes a benchmark mortality surface for the preceding 5 years. The analysis is performed annually and for each gender separately.⁴

First, the hypothesis $\beta_1^k = \beta_2^k = \beta_3^k = 0$ is tested against the full model. The test is performed as a likelihood ratio test with 3 degrees of freedom (for the approximating χ^2 -distribution). If the hypothesis is accepted at the 5% significance level no further tests are conducted. In this case the company must use the benchmark mortality (31) to value its liabilities.

If the hypothesis for full equality is rejected, the three (sub)hypotheses $\beta_3^k = 0$, $\beta_2^k = \beta_3^k = 0$ and $\beta_1^k = \beta_2^k = \beta_3^k = 0$ are considered in the order stated. Each hypothesis is tested against the previous as a likelihood ratio test with 1 degree of freedom. If a hypothesis is accepted at the 5% significance level the next hypothesis is tested, otherwise the test sequence stops. Note that if the first two hypotheses for partial equality are accepted the hypothesis for full equality is being tested again, but this time against a different alternative. It can thus happen that the hypothesis for full equality is eventually accepted although initially rejected.

The mortality to be used by the company is obtained by inserting the β -parameters estimated under the last hypothesis accepted, or the full model if all hypotheses were rejected, into (32).

³For the analysis performed by the companies a centralized version of the benchmark mortality is used, i.e. $\mu_k^{FSA}(x, t)$ in (32) is replaced by $\bar{\mu}_k(x, t) = (\mu_k^{FSA}(x, t) + \mu_k^{FSA}(x + 1, t)) / 2$. This is done because the FSA benchmark $\mu_k^{FSA}(x, t)$ is intended to reflect the mortality for exact age x , while the model mortality $\mu_k^{model}(x, t)$ represents the mortality for age $[x, x + 1)$. In this paper we use for simplicity the latter convention for all quantities, including the FSA benchmark.

⁴Adapted procedures exist for unisex mortality, but these will not be considered in this paper.

B Poisson Lee-Carter model

We use the Poisson variant of the classical Lee-Carter model to describe the evolution in Danish mortality, see Brouhns et al. (2002) and Lee and Carter (1992). The model assumes that

$$D_k^N(x, t) \sim \text{Poisson}(\mu_k^N(x, t)E_k^N(x, t)), \text{ with} \quad (35)$$

$$\mu_k^N(x, t) = \exp(a_k(x) + b_k(x)h_k(t)), \quad (36)$$

where $a(x)$ and $b(x)$ are respectively an age-dependent level and “rate of improvement”, while $h(t)$ is an index common to all ages of accumulated “improvement” evolving in time. Note that it is the product of b and h that determines the actual improvements, and that b and h themselves are only measures of improvement (hence the quotation marks in the previous sentence). In order to identify the parameters constraints must be imposed. We use the usual parameter constraints,

$$\sum_t h_k(t) = 0 \quad \text{and} \quad \sum_x b_k(x) = 1, \quad (37)$$

and the maximum likelihood fitting procedure described in Brouhns et al. (2002).

The Lee-Carter model is a standard model used in numerous mortality studies. It is a one factor model in which all improvements are fully correlated. From a longevity risk perspective this is in a sense the worst that can happen, and therefore this model although simple seems well-suited to study longevity risk. The actual number of deaths are assumed independent (conditioned on μ) and will therefore make only a small contribution to the total risk since random excess mortality is unlikely to occur in many age groups at once. The Poisson version is chosen because it complies with our general framework and because it handles cells with no deaths or no exposure better than the original Lee-Carter model. Also, with the original Lee-Carter model one typically adjusts the fit in the jump-off year to reproduce the actual number of deaths in order to avoid bias in the forecast. This is not needed with the Poisson version since it is fitted to the number of deaths, ref. Brouhns et al. (2002).

We have fitted the model to Danish data for the period 1980-2009 for ages 0-105 for each sex separately. The data are available at the Human Mortality Database (HMD) at www.mortality.org. These are the same data used by the Danish FSA to estimate the trend in the current benchmark. To save space we have not included the parameter estimates of a and b . These can be found in DSA (2012).

Figure 3 shows the raw death rates (solid lines), the Lee-Carter fit (dashed lines) and 95% confidence intervals obtained from the model (dotted lines) for selected ages. The model provides a decent fit to data and, importantly, the fluctuations in observed death rates are within the confidence intervals predicted by the model. Figure 4 shows the fit for all ages in the first and last data year. This confirms that data are well described by the model.

Having estimated the model we need to specify the dynamics of the index of mortality, h_k , in order to simulate from the model. Again we will follow the standard route and treat the estimated index of mortality as an observed time series and assume that it follows a random walk with drift

$$h_k(t + 1) = h_k(t) + \epsilon_k(t + 1), \quad \text{where } \epsilon_k \text{ are independent } N(\xi_k, \sigma_k^2). \quad (38)$$

The mean and standard deviation of the innovations, ϵ , are estimated from the “observed” part of h_k by the usual estimators

$$\hat{\xi}_k = \frac{1}{n} \sum_{t=1981}^{2009} \Delta h_k(t), \quad \hat{\sigma}_k = \left\{ \frac{1}{n-1} \sum_{t=1981}^{2009} \left(\Delta h_k(t) - \hat{\xi}_k \right)^2 \right\}^{1/2}, \quad (39)$$

where $\Delta h_k(t) = h_k(t) - h_k(t-1)$ and $n = 29$ equals the number of differences of the series. The estimated drift and standard deviation of the innovations are shown in Table 7. The table also shows the value of the index in the last data year, $h_k(2009)$, used as jump-off value when forecasting.

The estimated standard deviation is higher for females than for males reflecting a more volatile historic evolution in female mortality. The higher standard deviation for females gives rise to a higher stress for females than for males when we perform the gender-specific analysis in Section 4.2. Although we have no reason to believe that female mortality is intrinsically more volatile than male mortality we have chosen to preserve the estimated quantities in the analysis. Instead we will propose a unisex stress on the basis of the results of the gender-specific analysis to reflect the point of view that the level of uncertainty should be the same for the two genders going forward.

With the estimates in place it is now straightforward to simulate from the model one or more years into the future

- Draw independent innovations, $\epsilon_k(2010), \dots$, from a normal distribution with mean and standard deviation given by (39)
- Forecast the mortality index, h_k , by repeated use of (38) starting from the value of the index in the last data year, $h_k(2009)$
- Calculate the future mortality intensity, $\mu_k^N(x, t)$, by the relation in (36) using the forecasted value of the mortality index and $a_k(x)$ and $b_k(x)$ replaced by their estimated values

The assumption that the mortality index evolves like a random walk rules out the possibility of a structural break. For long term projections this might be a questionable assumption. However, since we will be using the model only to quantify the uncertainty one year ahead and since this uncertainty is dominated by the standard deviation of the innovations this is deemed not to pose a problem.

Gender	$\hat{\xi}_k$	$\hat{\sigma}_k$	$h_k(2009)$
Females	-1.8953	4.0983	-35.1604
Males	-1.9511	2.4744	-35.4906

Table 7: Estimated drift and standard deviation for innovations together with the value of the index in the last data year for females ($k = F$) and males ($k = M$).



Figure 3: Data and model fit for Danish females (top) and males (bottom) for the period 1980-2009. The plots show observed death rates (solid), Lee-Carter fit (dashed) and 95% confidence intervals (dotted) for ages 10, 20, 30, 40, 50, 60, 70, 80, 90 and 100 years.


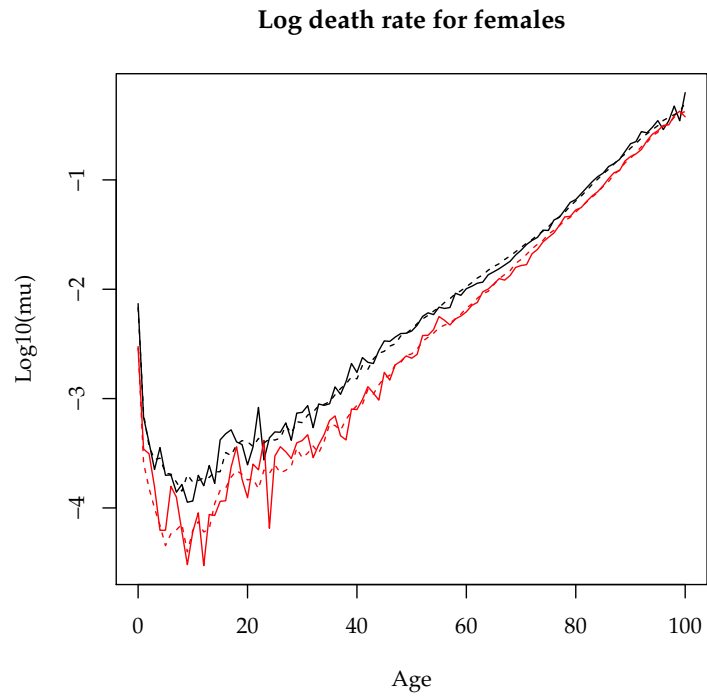


Figure 4: Data and model fit for Danish females (top) and males (bottom) for ages 0–100 years. The plots show observed death rates (solid) and Lee-Carter fit (dashed) in 1980 (black) and 2009 (red).

References

- D. Bauer, M. Börger, and J. Russ. On the pricing of longevity-linked securities. *Insurance: Mathematics and Economics*, 46:139–149, 2010.
- V. D. Biatat and I. D. Currie. Joint models for classification and comparison of mortality in different countries. Proceedings of 25rd International Workshop on Statistical Modelling, Glasgow, pages 89–94, 2010.
- H. Booth, R. J. Hyndman, L. Tickle, and P. de Jong. Lee-Carter mortality forecasting: A multi-country comparison of variants and extensions. *Demographic Research*, 15: 289–310, 2006.
- M. Börger. Deterministic shock vs. stochastic value-at-risk — an analysis of the Solvency II standard model approach to longevity risk. *Blätter DGVFM*, 31:225–259, 2010.
- N. Brouhns, M. Denuit, and J. K. Vermunt. A Poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics and Economics*, 31:373–393, 2002.
- A. J. Cairns. Modelling and management of longevity risk: approximations to survival functions and dynamic hedging. *Insurance: Mathematics and Economics*, 49:438–453, 2011.
- A. J. Cairns, D. Blake, and K. Dowd. A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk and Insurance*, 73: 687–718, 2006.
- A. J. Cairns, D. Blake, K. Dowd, G. D. Coughlan, and M. Khalaf-Allah. Bayesian stochastic mortality modelling for two populations. *ASTIN Bulletin*, 41:29–59, 2011.
- CEIOPS. Solvency II Calibration Paper, CEIOPS-SEC–40–10, 2010.
- I. D. Currie, M. Durban, and P. H. C. Eilers. Smoothing and forecasting mortality rates. *Statistical Modelling*, 4:279–298, 2004.
- M. Dahl. Stochastic mortality in life insurance: market reserves and mortality-linked insurance contracts. *Insurance: Mathematics and Economics*, 35:113–136, 2004.
- M. Dahl and T. Møller. Valuation and hedging of life insurance liabilities with systematic mortality risk. *Insurance: Mathematics and Economics*, 39:193–217, 2006.
- P. de Jong and L. Tickle. Extending Lee-Carter mortality forecasting. *Mathematical Population Studies*, 13:1–18, 2006.
- DSA. Solvency 2: Longevity Stress and the Danish Longevity Benchmark. Published on the homepage of the Danish Society of Actuaries, 2012. URL <http://www.aktuarforeningen.dk/English/Publications/tabid/92/Default.aspx>.
- EIOPA. EIOPA Report on the fifth Quantitative Impact Study (QIS 5) for Solvency II, EIOPA–TFQIS5–11/001, 2011.

- EIOPA. Revised Technical Specifications for the Solvency II valuation and Solvency Capital Requirements calculations (Part I), EIOPA-DOC-12/467, 2012. URL <https://eiopa.europa.eu/consultations/qis/insurance/long-term-guarantees-assessment>.
- EIOPA. Technical Specification on the Long Term Guarantee Assessment (Part I), EIOPA-DOC-13/061, 2013. URL <https://eiopa.europa.eu/consultations/qis/insurance/long-term-guarantees-assessment>.
- European Commission. Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II), 2009. URL <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:335:0001:0155:EN:PDF>.
- European Commission. Draft Implementing measures Solvency II. unpublished, 2011.
- Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). URL <http://www.mortality.org>.
- S. F. Jarner and E. M. Kryger. Modelling adult mortality in small populations: the SAINT model. *ASTIN Bulletin*, 41:377–418, 2011.
- R. D. Lee and L. R. Carter. Modeling and Forecasting of U.S. Mortality. *Journal of the American Statistical Association*, 87:659–675, 1992.
- R. D. Lee and T. Miller. Evaluating the performance of the Lee-Carter method for forecasting mortality. *Demography*, 38:537–549, 2001.
- J. S.-H. Li and M. R. Hardy. Measuring basis risk in longevity hedges. *North American Actuarial Journal*, 15:177–200, 2011.
- J. S.-H. Li, M. R. Hardy, and K. S. Tan. Uncertainty in mortality forecasting: An extension to the classical Lee-Carter approach. *ASTIN Bulletin*, 39:137–164, 2009.
- N. Li and R. Lee. Coherent mortality forecasts for a group of populations: An extension of the Lee-Carter method. *Demography*, 42:575–594, 2005.
- L. H. Nielsen. Assessment of longevity risk under Solvency II. *Life & Pension Risk, November issue*, pages 41–44, 2010.
- A. Olivieri and E. Pitacco. Stochastic mortality: The impact on target capital. *ASTIN Bulletin*, 39:541–563, 2009.
- E. Pitacco, M. Denuit, S. Haberman, and A. Olivieri. *Modelling longevity dynamics for pensions and annuity business*. Oxford University Press, Oxford, 2009.
- R. Plat. Stochastic portfolio specific mortality and the quantification of mortality basis risk. *Insurance: Mathematics and Economics*, 45:123–132, 2009.
- R. Plat. One-year Value-at-Risk for longevity and mortality. *Insurance: Mathematics and Economics*, 49:462–470, 2011.
- A. E. Renshaw and S. Haberman. Lee-Carter mortality forecasting with age-specific enhancement. *Insurance: Mathematics and Economics*, 33:255–272, 2003.

- A. E. Renshaw and S. Haberman. A cohort-based extension to the Lee-Carter model for mortality reduction factors. *Insurance: Mathematics and Economics*, 38:556–570, 2006.
- S. J. Richards, I. D. Currie, and G. P. Ritchie. A Value-at-risk framework for longevity trend risk. *Discussion paper presented to The Institute and Faculty of Actuaries*, 2012.
- S. Tuljapurkar, N. Li, and C. Boe. A universal pattern of mortality decline in the G7 countries. *Nature*, 405:789–792, 2000.