DISCUSSION PAPER PI-0826

Stochastic Portfolio Specific Mortality and the Quantification of Mortality Basis Risk

Richard Plat

March 2009

ISSN 1367-580X

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

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

STOCHASTIC PORTFOLIO SPECIFIC MORTALITY AND THE QUANTIFICATION OF MORTALITY BASIS RISK

RICHARD PLAT [a]
University of Amsterdam, Eureko / Achmea Holding and Netspar

This version: March 13, 2009 [1]
First version: September 10, 2008

Abstract
The last decennium a vast literature on stochastic mortality models has been developed. However, these models are often not directly applicable to insurance portfolios because:

a) For insurers and pension funds it is more relevant to model mortality rates measured in insured amounts instead of measured in number of policies.
b) Often there is not enough insurance portfolio specific mortality data available to fit such stochastic mortality models reliably.

Therefore, in this paper a stochastic model is proposed for portfolio specific mortality experience. Combining this stochastic process with a stochastic country population mortality process leads to stochastic portfolio specific mortality rates, measured in insured amounts. The proposed stochastic process is applied to two insurance portfolios, and the impact on the Value at Risk for longevity risk is quantified. Furthermore, the model can be used to quantify the basis risk that remains when hedging portfolio specific mortality risk with instruments of which the payoff depends on population mortality rates.

Keywords: portfolio specific mortality, stochastic mortality models, mortality basis risk, longevity risk, Solvency 2

1. Introduction

In recent years there has been an increasing amount of attention of the insurance industry for the quantification of the risks that insurers are exposed to. Important drivers of this development are the increasing internal focus on risk measurement and risk management and the introduction of Solvency 2 (expected to be implemented around 2012).

Solvency 2 will lead to a change in the regulatory required solvency capital for insurers. At this moment this capital requirement is a fixed percentage of the mathematical reserve or the risk capital. Under Solvency 2 the so-called Solvency Capital Requirement (SCR) will be risk-based, and market values of assets and liabilities will be the basis for these calculations.

[1] The author likes to thank Antoon Pelsser, Erik Tornij, Michiel Janssen, Marwa Khalaf-Allah and David Epstein.
[a] University of Amsterdam, Dept. of Quantitative Economics, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands, e-mail: H.J.Plat@uva.nl

Also for pension funds, a new solvency framework will be developed, either as part of Solvency 2 or as a separate project (usually named IORP 2). It is expected that the general principles will be similar as Solvency 2, meaning market valuation of assets and liabilities and risk-based solvency requirements.

Important risks to be quantified are mortality and longevity risk. Not only are these important risks for most (life) insurers, the resulting solvency margin will also be part of the fair value reserve. Reason for this is that it is becoming best practice for the quantification of the Market Value Margin to apply a Cost of Capital rate to the solvency capital necessary to cover for unhedgeable risks, such as mortality and longevity risk (see for example CEOIPS (2007)).

There is a vast literature on stochastic modeling of mortality rates. Frequently used models are for example those of Lee and Carter (1992), Brouhns et al (2002), Renshaw and Haberman (2006), Cairns et al (2006a), Currie et al (2004) and Currie (2006). These models are generally tested on a long history of mortality rates for large country populations, such as the United Kingdom or the United States. However, the ultimate application is to quantify the risks for specific insurance portfolios. And in practice there is often not enough insurance portfolio specific mortality data to fit such stochastic mortality models reliably, since:

- The historical period for which observed mortality rates for the insurance portfolio are available is usually limited, often in a range of only 5 to 15 years.
- The number of people in an insurance portfolio is much less than the country population.

Also, for insurers it is more relevant to model mortality rates measured in insured amounts instead of measured in number of people, because in the end the insured amounts have to be paid by the insurer. Measuring mortality rates in insured amounts has two effects:

- Policyholders with higher insured amounts tend to have lower mortality rates[2]. So measuring mortality rates in insured amounts will generally lead to lower mortality rates.
- The standard deviation of the observations will increase. For example, the risk of an insurance portfolio with 100 males with average salaries will be lower then that of a portfolio with 99 males with average salaries and 1 billionaire.

So fitting the before mentioned stochastic mortality models to the limited mortality data of insurers, measured in insured amounts, will in many cases not lead to results that are sufficiently reliable. In practice, this issue is often solved by applying a (deterministic) portfolio experience factor to projected (stochastic) mortality rates of the whole country population. However, it is reasonable to assume that this portfolio experience factor is a stochastic variable.

In this paper a stochastic model is proposed for portfolio specific mortality experience. This stochastic process can be combined with the stochastic country population mortality process, leading to stochastic portfolio specific mortality rates. The proposed model is, amongst others, based on historical mortality rates measured in insured amounts, but can also be used when only historical mortality rates measured in number of policies are available.

The model can be used to quantify portfolio specific mortality or longevity risks for the purpose of determining the Value at Risk (VaR) or SCR, which could also be the basis for the

[2] See for example CMI (2004).

quantification of the Market Value Margin. Also, it gives more insight in the basis risk when hedging portfolio mortality or longevity risks with hedge instruments of which the payoff depends on country population mortality. The market for mortality or longevity derivatives is emerging (see Loeys et al (2007)) and one of the characteristics of these derivatives is that the payoff depends on country population mortality. While this certainly has advantages regarding transparency and market efficiency, the impact of the basis risk is unclear. This basis risk is the result of differences between country population mortality and portfolio specific mortality, which is exactly what the proposed model is able to quantify.

Measurement of (portfolio specific) mortality rates in insured amounts is already used for a long time, starting with CMI (1962) and more recently for example in Verbond van Verzekeraars[3] (2008) and CMI (2008). In these papers portfolio experience factors, measured in amounts, are determined based on portfolio data that is collected from a representative part of the insurance market. The results of this are frequently used by the market participants as part of an estimate of future mortality rates. Furthermore, Brouhns et al (2002) also determine deterministic portfolio experience factors for the Belgian annuity policyholders, based on 3 years of historical data.

The literature on stochastic modeling of portfolio specific experience and mortality basis risk is less developed, possibly because of a lack of historical insurance portfolio data. Van Broekhoven (2002) determines a Market Value Margin for portfolio specific mortality risk. However, the model is not set up to be easily combined with existing country population models and the structure of the model over ages is very restrictive. Since the pattern of the portfolio experience factor over ages can vary for different portfolios, there has to be enough flexibility in the assumed structure over ages.
A related paper is the one of Jarner and Kryger (2009) who set up a model for mortality in small (country) populations, using the concept of frailty. The model seems to be too complex though to be calibrated to the limited data of insurance portfolios. Sweeting (2007) focuses in a more qualitative way on basis risk in survivor swaps. More generally, Dahl and Møller (2006) look at hedge strategies for mortality risk in life insurance liabilities.

So the model proposed in this paper is the first stochastic model for portfolio specific mortality that:

- can be combined easily with any stochastic country mortality process
- has enough flexibility in the assumed structure over ages
- has a structure that is simple enough to be able to calibrate it to limited historical data of life insurance portfolios

The remainder of the paper is organized as follows. First, in Section 2 the general model for stochastic portfolio specific experience mortality is defined. In Section 3 a 1-factor version of this model is applied to two insurance portfolios. Then in Section 4 and 5 the impact on the VaR and on the hedge effectiveness is quantified. Section 6 gives conclusions.

[3] Dutch Association of Insurers

2. General model for stochastic portfolio specific mortality experience

The first step in stochastic modeling of portfolio specific mortality rates is determining the historical portfolio mortality rates, measured by insured amounts. There are different kinds of definitions for mortality rates which are calculated in a slightly different manner (see Coughlan et al (2007)), for example the initial mortality rate or the central mortality rate. Regardless which definition is used, it is important that the same mortality rate definition is used for setting the country population mortality rates and the portfolio specific mortality rates. In the remaining part of this paper, we use the following definition for the initial mortality rate (see for example Namboodiri and Suchindran (1987):

(2.1) $$q_{x,t} = \frac{D_{x,t}}{\frac{1}{2}(N^P_{x,t} + N^U_{x,t} + D_{x,t})}$$

where $D_{x,t}$ is the number of deaths and $N^P_{x,t}$ and $N^U_{x,t}$ are the primo and ultimo total populations. The related portfolio mortality rate, measured by insured amounts, is:

(2.2) $$q^A_{x,t} = \frac{A^D_{x,t}}{\frac{1}{2}(A^P_{x,t} + A^U_{x,t} + A^D_{x,t})}$$

where $A^P_{x,t}$ and $A^U_{x,t}$ are the insured amounts primo and ultimo for the total portfolio and $A^D_{x,t}$ the insured amount of the deaths, for age x and year t.

Now the aim is to define a stochastic mortality model for the so-called portfolio experience mortality factor $P_{x,t}$ for age x and year t:

(2.3) $$P_{x,t} = \frac{q^A_{x,t}}{q^{Pop}_{x,t}}$$

where $q^{Pop}_{x,t}$ is the specific country population mortality rate for age x and year t, determined using (2.1). So $P_{x,t}$ represents the relation between a portfolio specific mortality rate (measured by insured amounts) and a country population mortality rate. Multiplying stochastic country mortality rates with stochastic $P_{x,t}$'s will give stochastic portfolio specific mortality rates. In this context a portfolio is seen as a group of homogenous risks, or a product group. $P_{x,t}$ is specific for each product group, it behaves differently for annuities as it does for term insurance. For reasons of convenience, the product specific nature is left out of the notation in the remaining of the paper, but the reader should be aware that all of the following is product (group) specific.

2.1 The basic model

Given that the model will often be based on a limited amount of data, it is desirable that the model for $P_{x,t}$ is as parsimonious as possible. Furthermore, the conjecture is that the difference between portfolio mortality and country population mortality is expected to be less at the highest ages, since the remaining country population at the highest ages is expected to have a relatively high percentage of people that are insured and have relatively high salaries. This is corroborated by the results in CMI (2004), where the difference between portfolio mortality and country population mortality is decreasing in age. Therefore, the proposed model leads to an expectation of $P_{x,t}$ that approaches 1 for the highest ages.

Given the above, we propose to model the mortality experience factor $P_{x,t}$ as:

$$(2.4) \qquad P_{x,t} = 1 + \sum_{i=1}^{n} X^{i}(x)\beta_t^i + \xi_{x,t}$$

where n is the number of factors of the model, X^{i} (x) is the element for age x in the i^{th} column of design matrix X, β_t^i is the i^{th} element of a vector with factors for year t and $\xi_{x,t}$ the error term. Another way to define the model is in matrix notation:

$$(2.5) \qquad P_t = \mathbf{1} + X\beta_t + \xi_t$$

where P_t is the vector of mortality experience factors, β_t the vector with factors and ξ_t the vector of error terms for time t. Furthermore, to ensure $P_{x,t}$ approaches 1, we require:

$$(2.6) \qquad \sum_{i=1}^{n} X_t^{i}(\omega)\beta_t^i = 0$$

where ω is the closing age of the mortality table (usually 120).

Now given a design matrix X, the vector β_t has to be estimated for each year. The structure of X (and the corresponding β's) can be set in different ways, depending on what fits best with the data and the problem at hand. One could use for example:

a) principal components analysis to derive the preferred shape of the columns X^{i}.
b) a similar structure as the multi-factor model proposed by Nelson and Siegel (1987) for modeling of yield curve dynamics.
c) a more simple structure, for example using 1 factor where the vector X is a linear function in age.

a) Principal Components Analysis (PCA)
Principal components analysis is a statistical technique that linearly transforms an original set of variables into a substantially smaller set of uncorrelated variables that represents most of the information in the original set of variables. Its goal is to reduce the dimensionality of the original data set.

The (m x k) matrix P contains historical observations of $P_{x,t}$ for m years and k ages. Instead of assessing the $P_{x,t}$ process for each age individually, the goal of PCA is to derive r linear combinations (where $r < k$) that capture most of the information in the original variables:

$$
(2.7) \quad \begin{aligned}
Z_1 &= v_{11}P_1 + v_{21}P_2 + \ldots + v_{k1}P_k \\
Z_2 &= v_{12}P_1 + v_{22}P_2 + \ldots + v_{k2}P_k \\
&\vdots \\
Z_r &= v_{1r}P_1 + v_{2r}P_2 + \ldots + v_{kr}P_k
\end{aligned}
$$

where P_j is the vector of observations for age class j.

Or, in matrix notation:

$$(2.8) \quad Z = PV$$

It can be shown (see for example Jolliffe (1986)) that the difference between the original data set and the set of linear combinations can be minimized by taking the eigenvectors of the covariance matrix Σ_{P^*} of the de-meaned historical observation matrix P^* as the columns of matrix V. The corresponding eigenvalues λ_j indicate the proportion of variance that each eigenvector (principal component) accounts for. By ordering the eigenvalues in such a way that $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_k \geq 0$, the dimensionality of the problem can be reduced by selecting the r eigenvalues (and the corresponding eigenvectors) that explain most of the variance of the original data set. The selected eigenvectors can be used as the columns X^i in (2.4).

b) Similar structure as Nelson and Siegel (1987)
Nelson and Siegel (1987) developed a parsimonious multi-factor model for yield curves that has the ability to represent shapes generally associated with yield curves. They model the instantaneous forward curve as:

$$(2.9) \quad f_t(\tau) = \beta_{1t} + \beta_{2t}e^{-\lambda_t \tau} + \beta_{3t}\lambda_t e^{-\lambda_t \tau}$$

where the parameters β_{1t}, β_{2t}, β_{3t} and λ_t have to be estimated from the observed yield curves. In practice λ_t is often fixed at a pre-specified value, simplifying the estimation procedure.

Since we are interested in the curve of $P_{x,t}$ over the ages and this curve historically shows shapes that roughly resemble possible shapes of yield curves, a structure similar as (2.9) could be used for modeling the $P_{x,t}$'s. An example of a possible 2-factor structure is given in Appendix 1.

c) A more simple structure
An alternative for structure a) and b) is a more simple structure, for example one where it is assumed that $P_{x,t}$ is linear in age for each t. It depends on the size of the insurance portfolio and the historical period whether structure a) leads to usable results and structure b) leads to a better fit to the data than this simple structure. For very large portfolios, structure a) and b) could be the most appropriate solutions. However, for the insurance portfolios considered in this paper, with 14 years of history and respectively about 100.000 policies and about 45.000 policies, principle components analysis didn't lead to usable results, and structure b) did not fit the data better than a simple linear structure.

2.2 Fitting the basic model

The structure of the model is such that it could be fitted with Ordinary Least Squares (OLS). However, the observations $P_{x,t}$ are all based on different exposures to death and observed deaths, so there is generally significant heteroskedasticity. Therefore Generalized Least Squares (GLS) should be used (Verbeek (2008)). When applying GLS in case of heteroskedasticity, each observation is weighted by (a factor proportional to) the inverse of the error standard deviation. Fitting this transformed model with OLS gives the GLS estimator, which accounts for the heteroskedasticity in the data.

When the available data are a cross-section of group averages with different group sizes and the observations are homoskedastic at individual level, the variance of the error term of the group averages is inversely related to the number of observations per group. In that case the square root of the number of observations in the group can be used as weights (Verbeek (2008)). For the problem in this paper this means that the square root of the number of deaths can be used as weights. So using a diagonal weight matrix W_t with these weights and applying it to (2.5) leads to a transformed model:

$$(2.10)\quad W_t\left(P_t-\mathbf{1}\right)=W_tX\beta_t+W_t\xi_t \quad or \quad \left(P_t-\mathbf{1}\right)^*=X^*\beta_t+\xi_t^*$$

Where the vectors or matrices labeled with an * are weighted with W_t. Now applying OLS to (2.10) gives the GLS estimator for β_t:

$$(2.11)\quad \hat{\beta}_t=\left(X^{*\prime}X^*\right)^{-1}X^{*\prime}\left(P_t-\mathbf{1}\right)^*=\left(X'W_t'W_tX\right)^{-1}X'W_t'W_t\left(P_t-\mathbf{1}\right)$$

This procedure can be repeated for each historical observation year, leading to a time series of vector $\hat{\beta}_t$.

2.3 Adding stochastic behavior

Now using the time series of the fitted β_t's, a Box-Jenkins analysis can be performed to determine which stochastic process fits the historical data best[4]. However, an important requirement in this case is biological reasonableness. For example, when assuming a non-stationary process such as a Random Walk for the β_t's, in certain scenario's the $P_{x,t}$'s could be 0 for all ages for some time, which is not biologically reasonable. Since the difference between country population mortality and portfolio mortality is dependent on factors that in our experience are normally relatively stable (size, composition and relative welfare of the portfolio), it doesn't seem reasonable to assume that this difference can increase unlimitedly. Therefore, a stationary process seems the most appropriate in this case. Given the often limited historical period of observations and the requirement of parsimoniousness, in most cases the most appropriate model will then be a set of correlated first order autoregressive (*AR*(1)) processes or equivalently, a restricted first order Vector Autoregressive (*VAR*) model:

[4] This is possible under the assumption that the historical fitted parameters are certain. Another possible approach would be to fit the parameters and the stochastic process at once, for example using a state space method combined with the Kalman filtering technique.

(2.12) $\beta_t = \delta + \Theta_1 \beta_{t-1} + \varepsilon_t$

where Θ_1 is a n x n diagonal matrix, δ is a n-dimensional vector and ε_t is a n-dimensional vector of white noise processes with covariance matrix Σ.

Possible alternatives are an unrestricted *VAR*(1) model or a first order (restricted) Vector Moving Average (*VMA*) model. In some cases an even simpler process than (2.12) is possible, being the so-called *ARIMA*(0,0,0)[5] process:

(2.13) $\beta_t = \delta + \varepsilon_t$

Model (2.12) and (2.13) can be fitted using OLS equation by equation. From the residuals e of the n equations the elements (i,j) of Σ can be estimated as:

(2.14) $\hat{\sigma}_{ij} = \frac{1}{(T-K)} \sum_{t=1}^{T} e_{it} e_{jt}$

where K is the maximum number of parameters used in either equations i or j (that is 2 when both processes are *AR*(1) processes).

An alternative is to estimate this simultaneously with the stochastic processes of the country population mortality model, which is the subject of the next paragraph.

When the insurance portfolio has developed significantly over the years, the fitted parameters over time are subject to heteroskedasticity. In this case GLS could be used, using either the results from table A2.1 in Appendix 2 or the square root of the number of deaths (see paragraph 2.2) as weights. When the portfolio has grown significantly and the current size of the portfolio is believed to be more representative for the future, the *relative* weights can also be applied to the residuals, weighting the earlier residuals less than the more recent ones.

2.4 Combine the process with the stochastic country population model

To end up with a stochastic process for portfolio specific mortality rates, the correlation between country population mortality rates and the portfolio mortality experience factors has to be taken into account. Therefore, the processes of the drivers of these have to be estimated simultaneously. Let's assume that the country population mortality is driven by m factors of which the processes α_t can be written as:

(2.15) $\alpha_t^k = X_k^\alpha \eta_k^\alpha + \varepsilon_k^\alpha \qquad k = 1, ..., m$

Now when the historical observation period is equal for the country mortality rates and the portfolio mortality experience factors, Seemingly Unrelated Regression (SUR, see Zellner (1963))

---

[5] Note that various names are used for this process in literature. Since the name ARIMA(0,0,0) seems to be the most widely used, we have adopted this name in this paper.

can be applied to fit all processes simultaneously. The processes don't have to be similar, so *AR*(1), Random Walk or other *ARIMA* models can be combined.

Re-writing (2.12) for each element *i* in a more general form as $\beta_t^i = X_i^\beta \eta_i^\beta + \varepsilon_i^\beta$ and combining all processes gives:

$$(2.16)\quad \begin{bmatrix} \beta^1 \\ \vdots \\ \beta^n \\ \alpha^1 \\ \vdots \\ \alpha^m \end{bmatrix} = \begin{bmatrix} X_1^\beta & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & \ddots & & & & 0 \\ \vdots & 0 & X_n^\beta & & & \vdots \\ \vdots & \vdots & & X_1^\alpha & & \vdots \\ \vdots & \vdots & & & \ddots & 0 \\ 0 & 0 & \cdots & \cdots & 0 & X_m^\alpha \end{bmatrix} \begin{bmatrix} \eta_1^\beta \\ \vdots \\ \eta_n^\beta \\ \eta_1^\alpha \\ \vdots \\ \eta_m^\alpha \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ \varepsilon_{n+m} \end{bmatrix}$$

which can be written more compactly as:

$$(2.17)\quad Y = X^{\alpha,\beta}\eta^{\alpha,\beta} + \varepsilon^{\alpha,\beta}$$

Now these processes can be fitted with SUR using the following steps:

1) Fit equation by equation using OLS
2) Use the residuals to estimate the total covariance matrix $\hat{\Omega}$ with (2.14)
3) Estimate $\hat{\eta}$ using GLS

To be more specific, the resulting estimator in step 3) is determined as:

$$(2.18)\quad \hat{\eta} = \left(X^{\alpha,\beta\prime}\,\hat{\Omega}^{-1}X^{\alpha,\beta}\right)^{-1}\left(X^{\alpha,\beta\prime}\,\hat{\Omega}^{-1}Y\right)$$

As mentioned earlier, in most cases the historical data period for portfolio mortality will be shorter than of country population mortality. In this case an alternative is only to do steps 1) and 2). In step 1) all available historical observations can be used for the different processes. In step 2) for the country population mortality the same historical data period should be used as is available for the portfolio mortality.

3. Application to example insurance portfolios

As mentioned in section 2, $P_{x,t}$ is specific for each product group or portfolio of homogeneous risks. In this section the general model described in section 2 is applied to two insurance portfolios[6]. The portfolios are respectively large and medium sized, and only data for males from age 65 on is taken into account. The large portfolio is a collection of collective pension portfolios of the Dutch insurers and contains about 100.000 male policyholders aged 65 or older. The medium portfolio is an annuity portfolio with about 45.000 male policyholders aged 65 or older. Note that this medium portfolio has developed significantly over time, so had less policyholders in earlier years. For both portfolios 14 years of historical mortality data is available.

For both portfolios, we examined a collection of 1-, 2-, and 3-factor models and concluded that the 2- and 3- factor models did not fit the data much better than a 1-factor model[7]. Since the 1-factor model uses less parameters, the Bayesian Information Criterion (*BIC*)[8] is more favorable for this structure. Therefore, the model we use is model (2.4) with $n = 1$ and:

$$(3.1) \qquad X^1(x) = 1 - \frac{x-\delta}{\omega-\delta} \qquad \qquad \delta \le x \le \omega$$

where δ is the start age (in this case 65) and ω is the end age (120). So in this formulation of model (2.4), the vector X is a linear function in age and, as required, $X^1(\omega) = 0$.

The reason why the 1-factor model fits the data as well as 2- or 3-factor models is that the data shows an upward slope for increasing ages, but the pattern along the ages is very volatile. For example, figure 1 shows two fits for the years 2006 and 2000. Fitting a more complex model through this data will not reduce the residuals significantly. Of course, this observation depends on the characteristics of the specific portfolio to which the model is fitted. For larger portfolios a 2- or 3-factor could give better results, since such a model is able to capture more shapes of the portfolio experience mortality factor curve.

Figure 1: example fit of model to actual observations for years 2006 and 2000

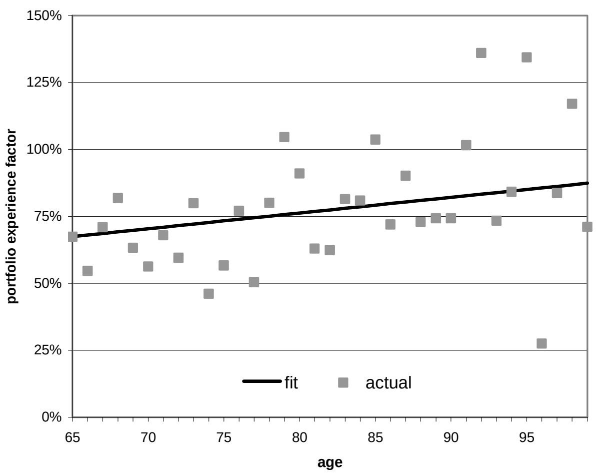

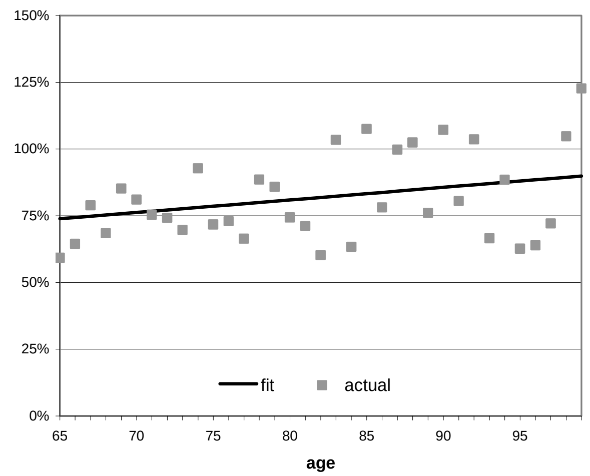

[6] The author thanks the Centrum voor Verzekeringsstatistiek (CVS) and Erik Tornij for the data of the large portfolio, and Femke Nawijn and Christel Donkers for the data of the medium portfolio.
[7] The fitting results for the 2-factor and 3-factor models are available upon request.
[8] *BIC* is a criterion that provides a trade-off between goodness-of-fit and the parsimony of the model.

The model is fitted using the procedure described in paragraph 2.2, where we have used the square root of the number of deaths as weights. The fitted β's are shown in figure 2[9]. Further results are given in table A2.1 in Appendix 2.

Figure 2: fitted β's for historical years 1993-2007

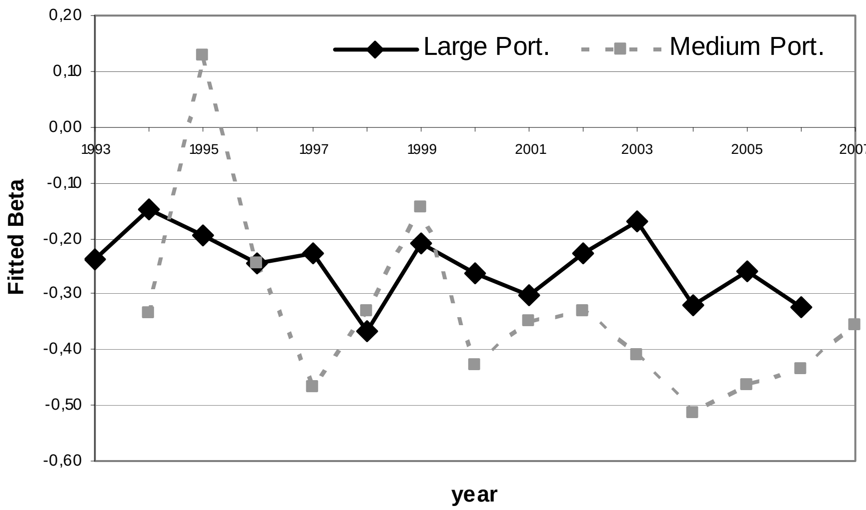

For both portfolios the results show an autoregressive pattern for the β's. Now a stochastic process for the future β's has to be selected. As mentioned in paragraph 2.3, a stationary process will be most appropriate. Also, since the historical data period is limited, the model should be as parsimonious as possible. We have fitted an *AR*(1), *AR*(2) and *ARIMA*(0,0,0) process to the data shown in figure 2. For both portfolios the *ARIMA*(0,0,0) process led to a more favorable *BIC* compared to the other processes.

Because of the significant development of the medium sized portfolio over the historical years, GLS is used for fitting the *ARIMA*(0,0,0) process. The square roots of the *relative* number of deaths in a year are used as weights. Relative means relative to the average number of deaths. These weights are also applied to the residuals, giving less weight to years where the portfolio was relatively small. Since the large portfolio was relatively stable over time, OLS is used for fitting the *ARIMA*(0,0,0) process for this portfolio.

The fitted processes for the portfolios are:

(3.2) Large portfolio: $\beta_t = -0,2497 + \varepsilon_t$, $\hat{\sigma} = 0,0625$

(3.3) Medium portfolio: $\beta_t = -0,3798 + \varepsilon_t$, $\hat{\sigma} = 0,1130$

The estimated error standard deviation $\hat{\sigma}$ is significantly larger for the medium sized portfolio, which is mainly the result of having less policyholders. The result of this is shown in figure 3,

[9] Note that although we have 14 years of data for both portfolios, the periods are slightly different, having data from 1993-2006 for the large portfolio and from 1994-2007 for the medium portfolio.

where the best estimates and the 99,5% / 0,5% percentiles are given for the portfolio experience mortality factors in the year 2016[10]. These specific percentiles are shown because the SCR of Solvency 2 is based on a 99,5% percentile.

Figure 3: best estimates and 99,5% / 0,5% percentiles for both portfolios - 2016

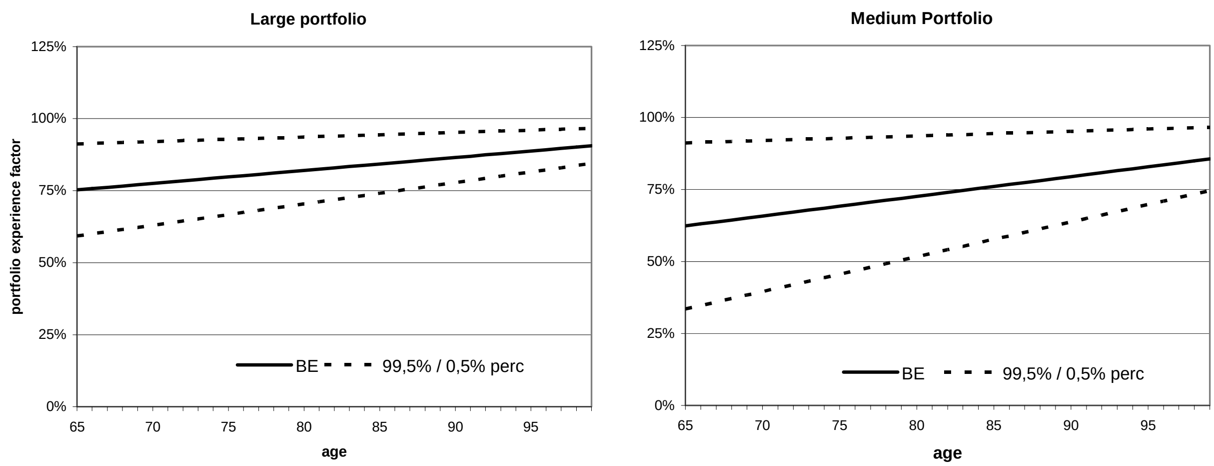

The figure shows that for the large portfolio the difference between the best estimate and the percentile(s) is in the range 10-15 %-point for ages 65-80. So taking this stochastic behavior of the portfolio experience mortality factor into account can have a reasonable impact on for example the Value at Risk. As expected, the impact is larger for the medium portfolio, where the difference between the best estimate and the percentile(s) is almost 30 %-point at its maximum.

4. Numerical example 1: Value at Risk

An important application of the presented model is the quantification of the Value at Risk (VaR) or SCR for longevity or mortality risk. In this paragraph the VaR is determined for the two portfolios, for different definitions / horizons of the VaR. First the model has to be combined with a model for country population mortality risk.

4.1 Stochastic country population mortality model

For the stochastic country population model we use the model of Cairns et al (2006a):

(4.1) $$\text{logit}\; q_{x,t}^{Pop} = \kappa_t^1 + \kappa_t^2 \left(x - \overline{x} \right)$$

Where $\overline{x}$ is the mean age in the sample range and κ_t^1 and κ_t^2 the two (stochastic) factors. We fitted this model to data of the Dutch population for the years 1950 – 2007. Using the resulting time series of parameter estimates, a 2-dimensional random walk process is fitted for the factors.

[10] Since a stationary process is assumed, the figure will be similar for other projection years.

The fitted parameters and the covariance matrices, including the covariance's with the portfolio experience mortality process of both portfolios, are given in Appendix 2.

Now combining the stochastic process above and the process described in section 3 leads to stochastic portfolio specific mortality rates. Figure 4 gives the best estimate mortality rates and percentiles for age 65. The percentiles are based on respectively deterministic and stochastic $P_{x,t}$'s.

Figure 4: best estimates and percentiles, with stochastic or deterministic $P_{x,t}$

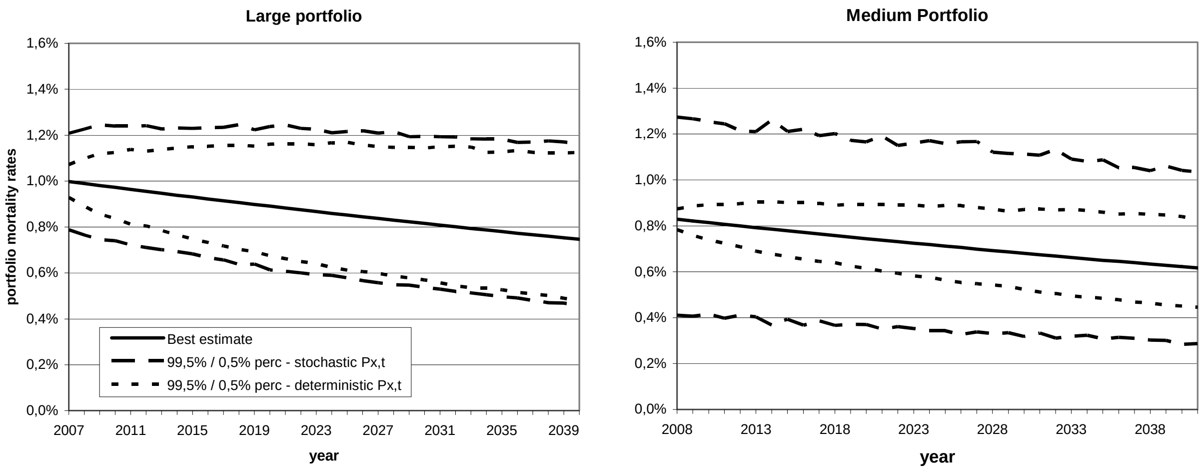

The figure shows that the additional risk of including stochastic $P_{x,t}$'s is highest at the start of the projection and decreases slowly in time. The reason for this is that the country population mortality rate risk is gradually increasing over time, resulting in a higher diversification effect between country population mortality rates and the $P_{x,t}$'s over time.

The percentiles for the medium portfolio seem quite dramatic. However, note that the shown percentiles are a result of picking the particular percentile every year, and not picking 1 scenario that represents the *x*%-percentile for the whole projection. Because of the assumed *ARIMA*(0,0,0) process the extremely low outliers will normally be (partially) compensated somewhere in time by high outliers. This is shown in figure 5, where two random (simulated) scenarios of the β's are given as an example.

Figure 5: two random (simulated) scenarios for β - medium portfolio

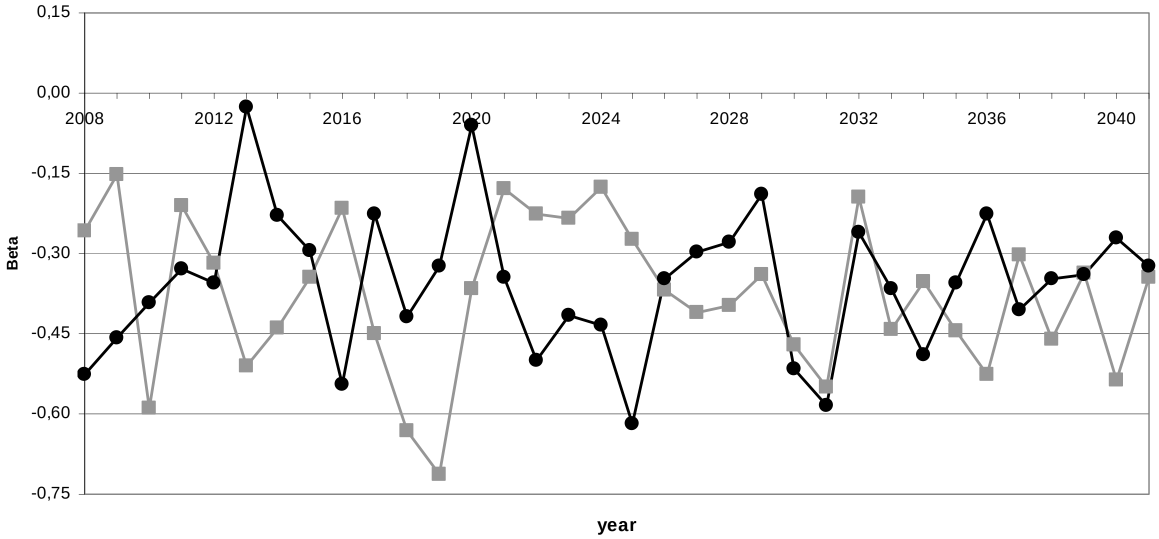

4.2 Impact on Value at Risk

Now using the described stochastic processes the impact on the VaR of stochastic (instead of deterministic) $P_{x,t}$'s is determined for both portfolios. The (present) value of liabilities is calculated for all simulated mortality rate scenarios[11]. The VaR is then defined as the difference between the *x*%-percentile and the average value of the liabilities. The impact is determined for three different definitions / horizons, which are all being used in practice:

1) 1-year horizon, 99,5% percentile, including effect on best estimate after 1 year
2) 10-year horizon, 95% percentile, including effect on best estimate after 10 years
3) Run-off of the liabilities, 90% percentile

So for definitions 1) and 2), at the 1-year or 10-year horizon all parameters are re-estimated using the (simulated) observations in the first 1 or 10 years, for each simulated scenario. The impact of the new parameterization on the best estimate of liabilities (for each scenario) is taken into account in the VaR. The results for the large and medium portfolio are given in respectively table 1 and table 2.

Table 1: impact of stochastic $P_{x,t}$ on VaR – large portfolio (in millions of Euros)

VAR definition	Deterministic $P_{x,t}$	Stochastic $P_{x,t}$	% difference
1-year, 99,5%	126,4	138,2	+ 9,3%
10-year, 95%	182,3	194,3	+ 6,6%
Run off, 90%	136,5	145,9	+ 6,8%

[11] For convenience we assumed that the portfolios only contain pension or annuity payments, so no spouse pension or annuities on a second life.

Table 2: impact of stochastic $P_{x,t}$ on VaR – medium portfolio (in millions of Euros)

VAR definition	Deterministic $P_{x,t}$	Stochastic $P_{x,t}$	% difference
1-year, 99,5%	45,1	73,0	+ 61,8%
10-year, 95%	69,2	95,1	+ 37,4%
Run off, 90%	54,1	75,1	+ 38,8%

Table 1 shows that for the large portfolio stochastic $P_{x,t}$'s lead to a VaR that is about 7%-9% higher compared to the VaR calculated with deterministic $P_{x,t}$'s. Table 2 shows that the impact for the medium portfolio is very high. The increase in VaR is between 37% and 68%, depending on the definition for VaR used. The reason for this is the large increase in volatility due to the addition of the stochastic $P_{x,t}$'s, which is mainly related to the size of the portfolio. Since a large part of the insurance portfolios in practice are of this size or smaller, this should be a point of attention when developing or reviewing internal models for mortality and longevity.

5. Numerical example 2: hedge effectiveness / basis risk

Because of the increasing external requirements and focus on risk measurement and risk management, the interest in hedging mortality or longevity risk is also increasing. A result of this is that a market for mortality and longevity derivatives is emerging (see Loeys et al (2007)). One of the main characteristics of these derivatives is that the payoff depends on country population mortality. While this certainly has advantages regarding transparency and market efficiency, the impact of the basis risk is unclear. Basis risk is the risk arising from difference between the underlying of the derivative and the actual risk in the liability portfolio. The model presented in this paper can be used to quantify this basis risk. In the example below the basis risk will be quantified for the two portfolios, where the longevity risk is (partly) hedged with the so-called q-forwards.

A q-forward is a simple capital market instrument with similar characteristics as an interest rate swap. The instrument exchanges a realized mortality rate in a future period for a pre-agreed fixed mortality rate. This is shown in figure 6. The pre-agreed fixed mortality rate is based on a projection of mortality rates, using a freely available and well documented projection tool[12].

Figure 6: mechanics of a q-forward

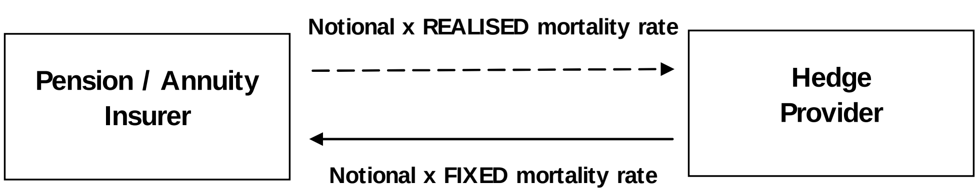

[12] For more information, see http://www.jpmorgan.com/pages/jpmorgan/investbk/solutions/lifemetrics

For example, when the realized mortality rate is lower than expected, the pension / annuity insurer will receive a payment which (partly) compensates the increase of the expected value of the insurance liabilities (caused by the decreasing mortality rates).

The basis for the instrument is the (projected) mortality of a country population, not the mortality of a specific company or portfolio. This makes the product and the pricing very transparent compared to traditional reinsurance.

For both insurance portfolios we determined a minimum variance hedge, based on deterministic $P_{x,t}$'s. The hedge is determined for a horizon of 10 years, but including the effect on the best estimate after 10 years (conform definition 2 of VaR in paragraph 4.2) . The hedge is determined for age-buckets of 5 years. For every bucket *i*, the impact of small shocks of the two factors of the country population model on the value of the liabilities and the value of an appropriate q-forward contract are calculated. The required nominal a_i^* for the q-forward of bucket *i* is then determined as:

$$(5.1) \qquad a_i^* = \frac{l_1 h_1 + l_2 h_2}{h_1^2 + h_2^2}$$

where l_i and h_i are the impact of the shock of the i^{th} factor on respectively the liabilities (*l*) and the hedge instrument (*h*).

The resulting hedge portfolio consists of 5 q-forwards for age-buckets of 5, from age 65 till age 89. The payoff of such a q-forward depends on the average mortality rate for the 5 ages in the bucket. The exact composition of both the hedge portfolios is given in Appendix 3.

Tables 3 and 4 show the impact on the hedge effectiveness when the $P_{x,t}$'s are assumed to follow the stochastic process described in section 3.

Table 3: impact of stochastic $P_{x,t}$ on hedge effectiveness – large portfolio

	VAR unhedged	VAR hedged	% reduction
Deterministic $P_{x,t}$	182,3	64,0	64,9%
Stochastic $P_{x,t}$	194,3	81,9	57,8%

Table 4: impact of stochastic $P_{x,t}$ on hedge effectiveness – medium portfolio

	VAR unhedged	VAR hedged	% reduction
Deterministic $P_{x,t}$	69,2	23,6	65,8%
Stochastic $P_{x,t}$	95,1	47,9	49,7%

The tables show that given deterministic $P_{x,t}$'s, the hedge reduces the VaR with about 65%. The risk is not fully hedged, because the hedge is based on small shocks of the two country population factors, while the factors in the tails of the distributions (which are relevant for VaR) are often more extreme.

For the large portfolio, table 3 shows that the hedge quality is decreasing, but is still reasonable. The basis risk for this portfolio is therefore limited. The reason for this is that on a longer horizon the impact of stochastic $P_{x,t}$'s levels out because of the assumed autoregressive process.

For the medium portfolio the hedge effectiveness is reduced to a larger extent. The effectiveness of the hedge can be improved by periodically adjusting the hedge portfolio. For smaller portfolios than this, it is probably questionable whether it is sensible to set up such hedge constructions.

6. Conclusions

In this paper a stochastic model is proposed for stochastic portfolio experience. Adding this stochastic process to a stochastic country population mortality model leads to stochastic portfolio specific mortality rates, measured in insured amounts. The proposed stochastic process is applied to two insurance portfolios. The results show that the uncertainty for the portfolio experience factor $P_{x,t}$ can be significant, mostly depending on the size of the portfolio.

The impact of the VaR for longevity risk is quantified. Depending on the definition used, the VaR increases by about 7%-9% for the large portfolio. The impact for the medium portfolio is very high, with an increase in VaR of 37%-68%. The reason for this is the high increase in volatility due to the addition of the stochastic $P_{x,t}$'s. Since a large part of the insurance portfolios in practice are of this size or smaller, this should be a point of attention when developing or reviewing internal models for mortality and longevity.

Furthermore, the basis risk is quantified when hedging portfolio specific mortality risk with q-forwards, of which the payoff depends on country population mortality rates. For the large portfolio the hedge quality is decreasing, but is still reasonable. The reason for this is that on a longer horizon the impact of stochastic $P_{x,t}$'s levels out because of the assumed autoregressive process. For the medium portfolio hedge effectiveness is reduced to a larger extent. For smaller portfolios than this, it is probably questionable whether it is sensible to set up such hedge constructions.

Appendix 1: example 2-factor model based on Nelson & Siegel

Nelson and Siegel (1987) proposed a parsimonious model for yield curves, which allows for different shapes of the curve. The Nelson-Siegel forward curve can be viewed as a constant plus a Laquerre function, which is a polynomial times an exponential decay term. It has three elements, respectively for the short, medium and long term. The model is very often used for yield curves and could serve as a basis for thinking for the P_t curves that are the subject of this paper. However, the Nelson-Siegel curve cannot directly be used for the P_t curves because $P_{x,t}$ should approach 1 near the closing age. Also, another requirement mentioned in section 2 is that the model is as parsimonious as possible, so a 2-factor model might be more appropriate in most cases.

Many variations on the Nelson-Siegel curve are possible. An example of such a model is the following model:

(A.1) $$P_t(\tau) = 1 + \beta_{1t} e^{-\lambda_1 \tau} + \beta_{2t} w_\tau \left(e^{-\lambda_1 \tau} - e^{-\lambda_2 \tau} \right)$$

$$\text{where } w_\tau = \varphi\left(\alpha \left[\frac{\tau - \tau_m}{\tau_m} \right] \right) \phi$$

The variable τ is 0 for the starting age of the data (in this case 65 years), τ_m is a strategically set middle point of the age interval (in this case 20, representing age 85), φ is the density of a standard normal distributed variable, α is a variable that arranges the shape of w_τ and can be set at 2 for example, and ϕ is a scale variable. The variable λ_1 can be solved in such a way that the second term of (A.1) approaches 0 for the closing age. The variable λ_2 can be solved in such a way that the third term of (A.1) is at its maximum somewhere between $\tau = 0$ and τ_m (in this case 75 years). The factors are shown in figure A.1, where x_1 represents the second term and x_2 the third term of (A.1).

Figure A.1: factors for model (A.1)

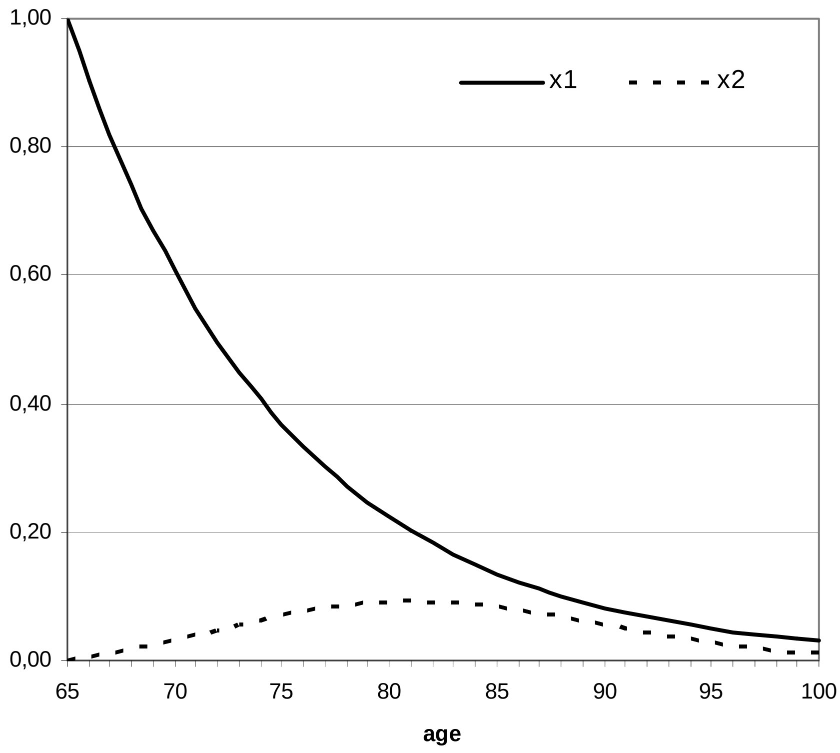

As can be seen from the figure and (A.1), the curve starts at age 65 at 1 + β_{1t} (where β_{1t} will be negative in general) and ends at 1 at higher ages. With the model (A.1) different shapes of the curve can be fitted, and the requirements in section 2 are met. A disadvantage of the model is the large number of parameters, of which some are set more or less arbitrary.

Appendix 2: further results

Table A.2.1 shows the fitting results for the β's in each year, for the large and medium sized portfolio.

Table A2.1: yearly fitting results for β's

Results large portfolio			
Year	β	**s.e.**	**t-ratio**
1993	-0,239	0,036	-6,55
1994	-0,149	0,041	-3,67
1995	-0,194	0,030	-6,55
1996	-0,246	0,033	-7,43
1997	-0,228	0,032	-7,20
1998	-0,368	0,023	-16,12
1999	-0,208	0,036	-5,77
2000	-0,261	0,029	-8,91
2001	-0,304	0,032	-9,46
2002	-0,226	0,033	-6,88
2003	-0,168	0,046	-3,62
2004	-0,321	0,048	-6,71
2005	-0,259	0,042	-6,11
2006	-0,325	0,040	-8,18

Results medium portfolio			
Year	β	**s.e.**	**t-ratio**
1994	-0,333	0,103	-3,23
1995	0,127	0,201	0,63
1996	-0,243	0,127	-1,92
1997	-0,467	0,091	-5,14
1998	-0,330	0,056	-5,92
1999	-0,143	0,065	-2,21
2000	-0,427	0,057	-7,56
2001	-0,349	0,089	-3,94
2002	-0,331	0,052	-6,32
2003	-0,408	0,050	-8,11
2004	-0,515	0,035	-14,77
2005	-0,462	0,047	-9,81
2006	-0,434	0,046	-9,52
2007	-0,355	0,079	-4,52

Table A.2.2 shows the fitted parameters for the 2-dimensional random walk model of section 4, and the covariance matrix including the covariance's with the process of section 3. Note that the country population parameter estimates slightly differ for the large and medium portfolio, because for the medium portfolio the year 2007 is also taken into account.

Table A.2.2: fit of country population model and covariance matrices

Fit - large portfolio		
	μ	σ
κ^1	-0,006206	0,0312395
κ^2	0,000182	0,00140485

Fit - medium portfolio		
	μ	σ
κ^1	-0,006722	0,022663
κ^2	0,000175	0,001436

Covariance matrix - large portfolio			
	β	κ^1	κ^2
β	0,003911	0,000445	0,000041
κ^1	0,000445	0,000976	0,000022
κ^2	0,000041	0,000022	0,000002

Covariance matrix - large portfolio			
	β	κ^1	κ^2
β	0,013808	0,002174	0,000053
κ^1	0,002174	0,000514	0,000020
κ^2	0,000053	0,000020	0,000002

Appendix 3: hedge portfolios

Table A.3.1: hedge portfolios for large and medium insurance portfolio

Characteristics hedge portfolio - large portfolio				
q-forward	Start age	End age	Nominal	Tick Size
1	65	69	117.865.528	100
2	70	74	34.221.141	100
3	75	79	10.420.145	100
4	80	84	2.213.481	100
5	85	89	315.640	100

Characteristics hedge portfolio - medium portfolio				
q-forward	Start age	End age	Nominal	Tick Size
1	65	69	74.346.033	100
2	70	74	24.226.106	100
3	75	79	3.203.979	100
4	80	84	143.135	100
5	85	89	8.224	100

References

BROUHNS, N., M. DENUIT, AND J.K. VERMUNT (2002): A Poisson log-bilinear regression approach to the construction of projected life tables, *Insurance: Mathematics and Economics 31*, 373-393

CAIRNS, A.J.G., D. BLAKE, AND K. DOWD (2006a): A two-factor model for stochastic mortality with parameter uncertainty: Theory and Calibration, *Journal of Risk and Insurance 73*, 687-718

CAIRNS, A.J.G., D. BLAKE, AND K. DOWD (2006b): Pricing death: Frameworks for the valuation and securitization of mortality risk, *ASTIN Bulletin 36*, 79-120

CEIOPS (2007): QIS 4 Technical Specifications

CONTINUOUS MORTALITY INVESTIGATION (1962): Continuous investigation into the mortality of pensioners under life office pension schemes, *available at:* *http://www.actuaries.org.uk/knowledge/cmi*

CONTINUOUS MORTALITY INVESTIGATION (2004): working paper 9, *available at:* *http://www.actuaries.org.uk/knowledge/cmi*

CONTINUOUS MORTALITY INVESTIGATION (2008): working paper 31, *available at:* *http://www.actuaries.org.uk/knowledge/cmi*

COUGHLAN, G. ET AL (2007): Lifemetrics Technical Document, *available at:* *http://www.jpmorgan.com/pages/jpmorgan/investbk/solutions/lifemetrics*

CURRIE, I.D., M. DURBAN, AND P.H.C. EILERS (2004): Smoothing and forecasting mortality rates, *Statistical Modelling 4*, 279-298

CURRIE, I.D. (2006): Smoothing and forecasting mortality rates with P-splines, *Talk given at the Institute of Actuaries, June 2006, available at: http:www.ma.hw.ac.uk/~iain/research/talks.html*

DAHL, M AND T. MØLLER (2006): Valuation and hedging of life insurance liabilities with systematic mortality risk, *Insurance: Mathematics and Economics 39*, 193-217

DIEBOLD, F.X., AND C. LI (2006): Forecasting the term structure of government bond yields, *Journal of Economics 130*, 337-364

JOLLIFFE, I.T. (2002): Principal Component Analysis, *Springer-Verlag New York, Inc.*

LEE, R.D., AND L.R. CARTER (1992): Modelling and forecasting U.S. mortality, *Journal of the American Statistical Association 87*, 659-675

LOEYS ET AL (2007): Longevity: a market in the making, *available at:* *http://www.jpmorgan.com/pages/jpmorgan/investbk/solutions/lifemetrics*

NAMBOODIRI, K. AND C.M. SUCHINDRAN (1987): Life Table Techniques and Their Applications, *Academic Press, Inc.*

NELSON, C.R., AND A.F. SIEGEL (1987): Parsimonious modeling of yield curve, *Journal of Business 60*, 473-489

RENSHAW, A.E., AND S. HABERMAN (2006): A cohort-based extension to the Lee-Carter model for mortality reduction factors, *Insurance: Mathematics and Economics 38*, 556-570

VAN BROEKHOVEN, H. (2002): Market Value of Liabilities Mortality Risk: A Practical Model, *North American Actuarial Journal 6*, 95-106

VERBEEK, M. (2008): Modern Econometrics, 3th edition, *John Wiley & Sons, Ltd*

VERBOND VAN VERZEKERAARS (2008): Generatietafels Pensioenen 2008, *publication in dutch*

ZELLNER, A. (1963): Estimators of Seemingly Unrelated Regressions: Some Exact Finite Sample Results, *Journal of the American Statistical Association 58*, 977-992