Netspar

Network for Studies on Pensions, Aging and Retirement

Netspar THESES

Menno van Wijk

Mortality in a Pensions Portfolio and the Application to an Internal Model for Solvency II

MSc Thesis 2012-054

Mortality in a pensions portfolio and the application to an internal model for Solvency II

Master Thesis

A thesis submitted in partial fulfillment of the requirements for the degree of Master in Quantitative Finance and Actuarial Science

Author:
Menno van Wijk BSc. (s791230)

Supervisor:
Prof. Dr. A.M.B. De Waegenaere

August 18, 2012

Abstract

This thesis aims at modeling the mortality experienced in the portfolio of a pension insurer. Not only is past mortality modeled, also projections about future mortality are made. The global model is split in two parts: a stochastic model for population mortality which allows for forecasts, combined with a point estimate from a portfolio mortality model. Future insurer-specific life tables are obtained for the period 2010-2060, along with 95% confidence intervals of the death probabilities. Results are compared to other projections used in Dutch (insurance) practice. The obtained, insurer specific, mortality model is then used to develop an internal model for assessing the longevity risk of the insurer in the Solvency II framework. The results for this model are compared to the results for the standard approach.

KEYWORDS: longevity risk, pensions, insurance, (portfolio) mortality, forecasting, life expectancy, Solvency II, internal model.

Contents

Chapter 1

Introduction

History shows that mortality, not only in the Netherlands, is rapidly decreasing over time. As a result, life expectancies increase. In 1850, Dutch new-born males had a period life expectancy (for explanation of this term, see Chapter 2) of 38.90 years. This number has risen to 78.53 years in 2009. For females, life expectancy at birth has grown from 40.81 to 82.64.[1] These life expectancies are however based on present death probabilities, whereas the actual realizations of life spans are based also on future, unknown, probabilities. It is thus known that people tend to get older, but it is not known by how much exactly. The uncertainty in the future distribution of life expectancies, or actually in the future distribution of the underlying death rates, is called longevity risk. Pension funds and pension insurers that offer lifelong annuities to their clients are very much affected by this type of risk. In the new financial regulation framework Solvency II, insurers are obliged to quantify their longevity risk and reserve extra capital for it. In order to be able to determine those reserves, projections for future mortality need to be made, with acknowledging the surrounding uncertainty.

In the past, many different people and organizations have come up with ways to forecast future mortality. Lee and Carter (1992) for instance introduced a relatively simple model in which they extrapolate the development of death rates to the future. Since then, many researchers have developed similar models. Likewise, governmental institutions perform research on future mortality. There is for instance Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS) that extrapolates mortality rates for specific causes of death. The National Institute for Public Health and the Environment (Rijksinstituut voor Volksgezondheid en Milieu, RIVM) makes a distinction between smoking related and non-smoking related mortality. In practice, models for future mortality have been developed as well. In the Netherlands there are two main organizations of actuaries who have come up with projections. The Dutch Actuarial Association (Actuarieel Genootschap, AG) have a history of projecting mortality. Their latest publication was issued in the summer of 2010. In that paper they presented forecasts for mortality up to 2060. In the same summer, the Workgroup PLT of the Dutch Association of Insurers (Verbond van Verzekeraars) came up with a forecast as well.

The models by AG and PLT do not differ that much. They largely implement the same idea of forecasting, based on a short term and a long term trend. There is however one fundamental difference. Whereas the AG only models future mortality for the Dutch population as a whole, the Association of Insurers also specifically looks at mortality rates of the group of people who have a pension insurance at one of their affiliates. In this way they do not only take into account the longevity of the population, but also the effects of adverse selection on mortality in the portfolio of these insurers. As widely expressed in the literature, people with a pension insurance are expected to live longer than people who do not have one. Therefore, it is wise for insurers to not

[1] Numbers via Human Mortality Database.

base their calculations solely on mortality rates of the general population but to also acknowledge that people in their portfolio will likely live longer than the average person in the population.

The goal of this thesis is twofold. First I will derive a model for the mortality in the pensions portfolio of ASR, which allows for a forecast of the future mortality in this portfolio. Basically I will do this by developing two separate models and then integrating them. I will start by assessing the longevity of the Dutch population as a whole, based on data of this group. For forecasting future mortality in the Dutch population I will use the Lee-Carter method. Subsequently I will implement the effects of adverse selection in the ASR pensions portfolio. These effects I will estimate by comparing past mortality in the portfolio to past population mortality. Combining the longevity projections and the effects of adverse selection I will forecast future mortality in the portfolio. The resulting mortality model is then specific for the pensions portfolio of ASR. Secondly, I will use this model as a basis for an internal model for quantifying the amount of capital ASR should hold for covering their longevity risk according to the Solvency II framework for insurers. The results for this internal model are then compared to those from the standard approach. On the basis of this comparison, I draw conclusions and make recommendations.

The structure of this thesis is as follows. Chapter 2 provides a short explanation of important terms in the scope of this thesis. Chapter 3 presents models regarding the projection of future mortality for general populations, both from the academic literature and from practice in the Netherlands. Furthermore I also present my approach regarding modeling future population mortality and the results of this. In Chapter 4 I do the same, but then for modeling portfolio specific mortality and the combining of both models. Chapter 5 first discusses how the Solvency Capital Requirement for longevity risk is determined according to Solvency II. In this Chapter I also present an internal model, which could be used to better approximate this capital requirement. Also the results of both approaches applied to a portfolio resembling ASR's pension portfolio are presented and compared. Conclusions and recommendations are then presented in Chapter 6.

Company profile
With over 2 million customers, ASR is one of the largest insurance companies in the Netherlands. The 4.5 billion euro revenue generated in 2011 makes ASR in terms of revenue one of the top four insurers in the Netherlands (excluding health insurers). Although ASR has only held that name since 2008, the organization consists of several brands - AMEV, Stad Rotterdam and Woudsend Verzekeringen among others - that date back to 1720. Through four different labels, ASR provides insurance products for private individuals and small and medium-sized enterprises. Core markets for ASR are non-life, occupational disability, individual-life and pensions. Besides these core markets, ASR also focuses on products related to these core markets (for instance saving and investment and mortgages) and on relatively independent markets such as the travel and leisure market and the funeral market. The main office of ASR is situated in Utrecht (ASR, 2011).

Chapter 2

Explanation of terms

In this chapter I will briefly explain some terms commonly used in mortality studies and which I will use throughout this thesis as well. More elaborate explanations can for instance be found in Pitacco et al. (2009).

- *Exposure-to-risk:* the exposure-to-risk $E_{x,t}$ denotes the number of person years lived during year t by people aged x at the start of the year. Assuming that people who die during a year have on average been alive during half of the year, the exposure-to-risk can be approximated by the number of survivors plus half the number of deaths in this group.

- *Central death rate:* the central death rate $m_{x,t}$ is defined as the total number of people aged x who have died during year t ($D_{x,t}$), divided by the exposure-to-risk of age group x during year t. In formula: $m_{x,t} = \frac{D_{x,t}}{E_{x,t}}$.

- *Death probability:* the probability $q_{x,t}$ that an individual aged x at the start of year t will die before having reached year $t+1$. This quantity is closely related to the central death rate $m_{x,t}$. Throughout this thesis, I use the relation $q_{x,t} = 1 - \exp(-m_{x,t})$. $q_{x,t}$ is a one-year death probability. One can however also define a s-year death probability (probability of dying within s years) after having reached age x in year t by ${}_sq_{x,t}$. Throughout this thesis, the term death probability will refer to a one-year death probability.

- *Survival probability:* the survival probability $p_{x,t}$, i.e., the probability that a person aged x will survive year t, is defined by $p_{x,t} = 1 - q_{x,t}$. Like for the death probabilities, one can also define the probability of surviving an additional s years after having reached year t by ${}_sp_{x,t} = 1 - {}_sq_{x,t}$. In applications s will typically be an integer, but it need not be.

- *Force of mortality:* the force of mortality $\mu_{x,t}$ is defined by $\mu_{x,t} = \lim_{s \downarrow 0} \frac{{}_sq_{x,t}}{s}$. It is also referred to as the instantaneous rate of mortality at the age x in the year t. A typical assumption in the literature is that the force of mortality is piecewise constant: $\mu_{x+\tau_1,t+\tau_2} = \mu_{x,t}$ for $0 \leq \tau_1, \tau_2 < 1$. I will also adopt this assumption, which is an essential one when the analysis is done for age groups with widths of one year. This assumption implies that the force of mortality becomes equal to the central death rate $m_{x,t}$.

- *Period life expectancy:* the (remaining) life expectancy calculated for a person in year t, based on mortality rates which hold for year t. If this person is aged 40 in year t, his survival probability to reach the age 50 in year $t+10$ will entirely be based on the survival probabilities for people aged 41, 42, etc. in the year t. For a person aged x in year t, the remaining period life expectancy $e_{x,t}$ is defined by

$$e_{x,t} = \sum_{\tau=1}^{\omega-x} {}_\tau p_{x,t} + \frac{1}{2}, \tag{2.1}$$

where ω denotes the maximum age an individual can reach. The first term calculates the number of complete years lived, the half is to compensate for the fact that, on average, a person dies half a year after his last and half a year before his next birthday. Note that the maximum age to be obtained is of course unknown, so typically an assumption is made. Throughout this thesis I will use that ω equals 120.

- *Cohort life expectancy:* as above, but then cohort mortality rates are used. This means that the survival probability to reach the age 50 for the person discussed above is based on the probability to reach age 41 in year t, age 42 in year $t+1$, age 43 in year $t+2$ etcetera. These future death probabilities are not deterministic at time t, but stochastic (i.e., there exists longevity risk). Therefore the cohort life expectancy can only be calculated if assumptions about the future death probabilities are made, for instance that they are according to a best estimate. When one does not assume that mortality rates are constant over time, the cohort life expectancy differs from the period life expectancy. If mortality rates decline over time, cohort life expectancy is higher than period life expectancy.

Chapter 3

Projecting population mortality

This chapter first of all presents an overview of the models used to predict future mortality. Section 3.1 focuses on models developed in the academic literature. Section 3.2 presents the models used and results from projections that are made by influential organizations in the Netherlands. The following section 3.3 presents the approach that I have used for modeling future population mortality. Section 3.4 concludes this Chapter and presents the results of my population mortality model.

3.1 Models in the literature

In the past, mortality patterns were parameterized in the form of different laws. One of them, and probably the most important one, is the Gompertz-Makeham law. This assumes that the logarithm of mortality approximates a straight line when viewed over age. This law, already developed in the 19th century, seems to hold well for ages between 30 and 100 (Peters et al., 2012). As mortality rates have rapidly declined in the 20th century, the need was felt to come up with not only models for present mortality, but also with models that aim to predict future mortality rates. This meant that models with only an age component needed to be expanded with a time component as well. Lee and Carter (1992) have developed probably the most influential and most widely used model in this respect. In short, the Lee-Carter method assumes the existence of a time effect in log mortality rates, meaning that death rates in a population have a strong tendency to move up or down together over time. This indeed seems to apply to low-mortality countries (Pitacco et al., 2009). The Lee-Carter method does not only allow for the calculation of point estimates of future rates of mortality and life expectancies, but also for the determination of confidence intervals. I will discuss the Lee-Carter method extensively in section 3.3.

The Lee-Carter model is somewhat simpler than other models, but not less accurate (see for instance Dowd et al., 2010). Lee (2000) showed that the original Lee-Carter method performed well in explaining the rise in life expectancy in the US in the period 1989-1997. However, many alterations have been proposed in the literature to either improve the statistical soundness of the model or to come to a better fit. Brouhns et al. (2002) for instance propose the Poisson model. This model is very similar to the Lee-Carter model, but models the number of deaths $D_{x,t}$ conditional on the exposure-to-risk as a Poisson random variable, whereas Lee and Carter model the central death rates $m_{x,t}$ as a random variable. In formula, Brouhns et al. (2002) consider

$$D_{x,t} \sim \text{Poisson}(E_{x,t} m_{x,t}), \tag{3.1}$$

where $m_{x,t} = \exp(a_x + b_x \kappa_t)$, as in the Lee-Carter method. The use of a Maximum Likelihood estimation process allows for the quantification of parameter risk (the risk that the parameter values found after estimation are not the right ones), for which in the Lee-Carter method more effort

is required (one can for instance use a bootstrapping method). Note that reformulating the Lee-Carter model in this way also allows the estimation process to be continued for data cells in which there are no deaths observed. In the original Lee-Carter method this is not possible, see section 3.3.

Some empirical analyses suggest that the relation $\log \frac{q_{x,t}}{p_{x,t}}$ is approximately linear across age for fixed time. This is why Cairns et al. (2006) propose a model (CBD-model) based on this relation, including one age component and two time components. They use the equation

$$\log \frac{q_{x,t}}{p_{x,t}} = \kappa_t^{(1)} + \kappa_t^{(2)} x, \tag{3.2}$$

where $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ are stochastic processes and $\log(\cdot)$ denotes the natural logarithm (as in the remainder of this thesis). The advantage of this model compared to the Lee-Carter method is that it does not impose perfect correlation of changes in mortality at different ages from one year to the next (Pitacco et al., 2009). A disadvantage is however that the model only seems to hold for higher ages (usually $x > 40$, see Peters et al., 2012). If one would want to forecast future mortality patterns for all ages, which is for instance what a demographer would do, the CBD-model is thus not appropriate. Plat (2009a) proposes a model which aims to combine the strong points of the Lee-Carter model and the CBD-model and which seems to provide a better fit than those models, however being slightly more complex (Haberman and Renshaw, 2011).

Besides including an extra time component, one could also include a cohort effect in a mortality model. Renshaw and Haberman (2006) do this for instance for the Lee-Carter method, coming up with an Age Period Cohort (APC) model. They use the equation

$$\log m_{x,t} = a_x + b_x^{(0)} i_{i-x} + b_x^{(1)} \kappa_t, \tag{3.3}$$

where $b_x^{(0)} i_{i-x}$ is the additional term (compared to the Lee-Carter method) that is supposed to model the specific cohort effects. Estimating the parameters has now become significantly more difficult than in the standard Lee-Carter setting, but it also makes the model more flexible. Cairns et al. (2009) also propose an alteration of the CBD-model to include a cohort effect. When there is a cohort effect present in mortality for a specific population, people born in a certain year or period experience significantly different changes in mortality than other people in the population. A cohort effect thus differs from a time effect, as the latter holds for the entire population. Haberman and Renshaw (2011) argue that including a cohort parameter can lead to more accurate projections, but only for countries for which there indeed exists a significant cohort effect. Typically the UK is considered as such a country. For the Netherlands, there is as far as I know no conclusive evidence for the existence of a cohort effect.

It is important to note that all models discussed in this section do not only provide point estimates of future death rates and life expectancies, but also allow for the construction of confidence intervals. This can be seen as a statistically valid measure of the uncertainty incorporated in the model. As the problem of longevity risk specifically concerns the uncertainty in future death rates and life expectancies, the confidence intervals are essential in assessing this risk. In section 3.2 one can see that this is an important property that is not present in the models that are used in practice in the Netherlands.

3.2 Projections from practice

In this section I present models and results from projections made in practice in the Netherlands. In general, the methods presented in this section are mathematically slightly less advanced than the models from the academic literature. For insurers in the Netherlands these projections are however important and often used. This is why I will discuss them in short, so I can later compare my results to theirs. As ASR now uses life tables based on these projections, especially those

by AG and PLT, they would like to know how my projection method and results relate to those approaches.

3.2.1 CBS

On December 17, 2010, CBS launched their latest projection for future mortality (Van Duin and Garssen, 2010 and Van Duin et al., 2010). In this publication they presented a forecast of mortality for the period 2010-2060. Their analysis is based on the extrapolation of past mortality for different causes of death, with opinions of (mostly medical) experts taken into account. They perform this analysis for men and women separately and for five-year age intervals up to the age of 85. They distinguish the following causes of death:

- heart and coronary diseases
- cancer (divided in lung cancer and other smoking related cancers, breast cancer, prostate cancer, colorectal cancer and 'other cancers')
- COPD (Chronic Obstructive Pulmonary Disease, a chronic lung disease)
- non-natural causes of death
- other causes of death.

For the age intervals above 85 only lung cancer, COPD and other causes of death are distinguished.

The progress of mortality due to these specific causes of death is predicted by looking at past trends and determining whether these can be seen as representative for the future, using expert opinions. By aggregating the different predictions, CBS finds the projection of total future mortality. Figure 3.1 shows the projected development of the death probabilities for different age groups and sexes. [1]

CBS expects that in 2060, the period life expectancy at birth for men is 84.5 years. For women they expect a life expectancy at birth of 87.4 years. When compared to 2009, this means a gain of 6.0 years for men and a gain of 4.7 years for women. Life expectancies for the genders thus converge. One should notice that these numbers are mere point estimates, but CBS also provides confidence intervals. For both genders they use in 2060 a symmetric 95% confidence interval with a width of in total eleven years. This can be interpreted as that they expect the period life expectancy at birth for men in 2060 to lie, with a probability of 0.95, within the interval $[79.0; 90.0]$. It should be noted that these confidence intervals are not based on statistical features of the model, but on the accuracy of past estimates of life expectancies compared to the realized values and on the uncertainty present in other, extrapolative models (Carolina and Van Duin, 2010). Figure 3.2 shows the evolution of life expectancies at birth, together with the 67% and 95% confidence intervals.

Pension funds and insurers are also very interested in the remaining life expectancies at the age of 65, as this is (at this moment) the age at which pension benefits start to pay out. In 2060, CBS expects those to be 21.7 years for men and 24.2 years for women. In 2009, these numbers were 17.4 for males and 20.8 for females. One can thus again see convergence for the sexes. The total width of the 95% confidence interval is 9 years. Note that these numbers are period life expectancies, not cohort life expectancies.

As put forward before, CBS partially bases its projections on the opinion of (medical) experts. I will not fight the statement that medical experts know more about mortality than statisticians, econometricians or actuaries, but there is no conclusive evidence that the use of expert opinions

[1] It also shows the former projection of CBS, but I will not discuss that forecast.

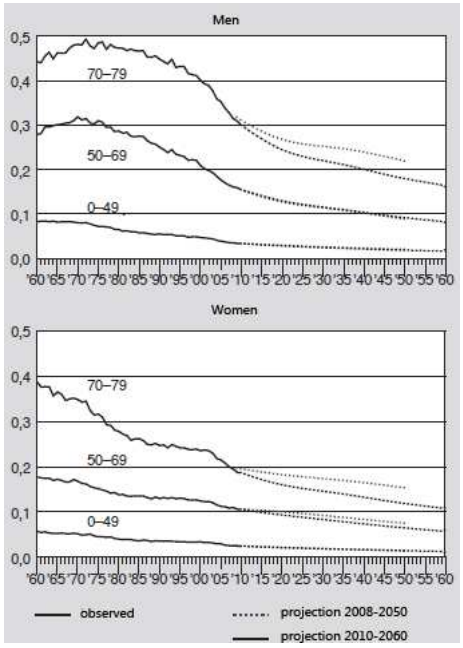

Figure 3.1: *(Projected) death probabilities, aggregated for all causes of death, according to CBS. Source: Van Duin et al. (2010). Translated legend.*

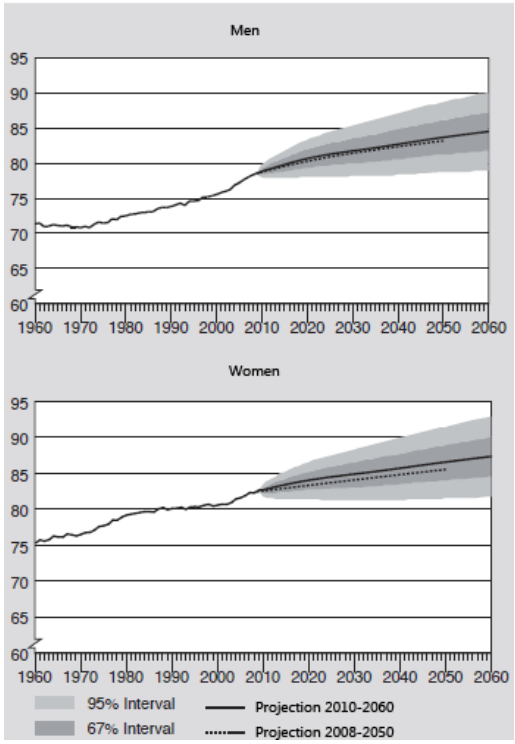

Figure 3.2: *(Projected) period life expectancies at birth according to CBS. Source: Van Duin et al. (2010). Translated legend.*

leads to better forecasts than mere extrapolations (see Wilmoth, 2000). Furthermore, their analysis is based on the projection of future mortality per cause of death. The cause-specific death rates are then aggregated to come to a total projected death rate. In the literature there is "consensus that projections based on projections of causes of death in general is worse than the direct projection of all-cause mortality" (Peters et al., 2012, p. 8). This comes mostly from misdiagnoses and cases where the deceased suffered from multiple illnesses (multi-morbidity).

3.2.2 RIVM

RIVM (Luijben and Kommer, 2010) uses a very similar approach in forecasting future mortality. Like CBS they do not model total mortality at once, but split it in different causes of death. Their division is smaller however, they only distinguish between smoking related and non-smoking related mortality. They argue that they can accurately predict smoking related mortality, by observing cohort patterns and smoking related mortality in similar countries. Also, the decline in smoking related mortality for women is lagged when compared to men. In the long run smoking related mortality is decreasing for both males and females, but in the short run it is increasing for women. Non-smoking related mortality decreases for both men and women according to approximately the same pattern. For non-smoking related mortality, also trends in ten comparable countries (Germany, England and Wales, France, Denmark, Norway, Sweden, Finland, Switzerland, Spain and Italy) are taken into account. Modeling of this type of mortality is done via the methodology by Li and Lee (2005). This method is based on the use of information on mortality trends in similar populations when modeling future mortality. Note that this method is thus extrapolative and not based on expert opinions.

These effects lead at first to an expected convergence of life expectancies for men and women, later followed by an expected divergence. In 2050, RIVM predicts period life expectancies at birth of 83.8 years for men and 88.1 years for women. Period life expectancies at age 65 are 21.1 years for men and 24.2 years for women. Note that no confidence intervals are provided. Figure 3.3 shows the development of these quantities. For men, the results are comparable to the ones found by the CBS. For women however, RIVM predicts lower mortality (higher life expectancies) than CBS. This is mainly because by distinguishing between smoking related and non-smoking related mortality, RIVM observes a steeper decreasing trend in mortality rates for women than CBS.

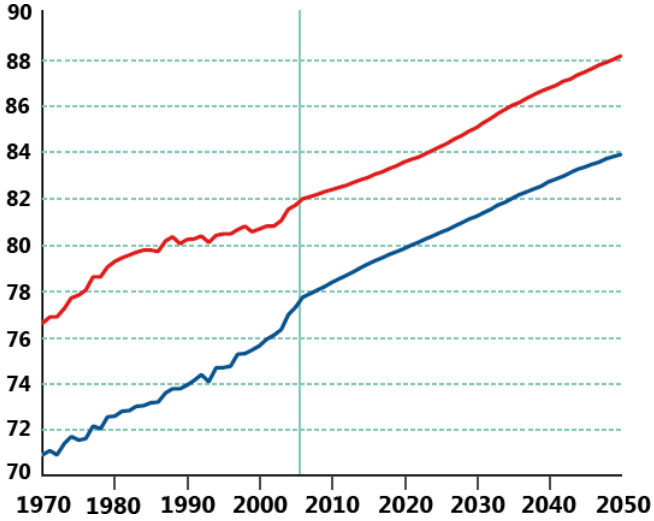

Figure 3.3: *(Projected) period life expectancies at birth according to RIVM (2010). Blue line concerns men, red line concerns women.*

Like the forecast by CBS, RIVM uses some sort of expert opinion in their projections. They do however distinguish between less different causes of mortality, making the model somewhat simpler. Also, they have empirical evidence to back their assumptions for the development of smoking-related mortality. An important drawback of their model is however the lack of quantification of uncertainty. Whereas CBS provides confidence intervals, RIVM only comes up with point estimates. For insurers the margins of uncertainty are arguably just as important as the point estimates. This comes forward for instance in the new Solvency II regulation framework for insurers, where insurers are not only obliged to reserve capital for the best estimate of future liabilities, but also for scenarios in which things turn out worse for them.

3.2.3 AG

The forecast of AG (Actuarieel Genootschap, 2010) is performed by a committee of actuaries active for different companies in the field. Where CBS and RIVM used forecasts (partially) based on expert opinions, AG only extrapolates historical data. They observe a long term decreasing trend in mortality in the Netherlands for the period 1988-2008. Besides this long term effect, they also claim the existence of an ongoing short term trend which started in 2001. They use both trends in the forecasting of future mortality for the period 2010-2060. The long term trend is used to come up with a 'goal table' for the final year 2060. This goal table is the table that is reached when the trend between 1988 and 2008 continues until 2060. The short term trend is a determinant of the path which is followed to reach the goal table. Since the short term trend shows a more rapid decrease of mortality than the long term trend, the resulting path of the progress of life expectancy is a concave one. Figure 3.4 demonstrates this. It also shows that the short term trend does not determine the level of the goal table. AG have chosen for 2060 as the end year of the projection period to be able to cover a large part of the run-off of present portfolios of pension funds and life insurers.

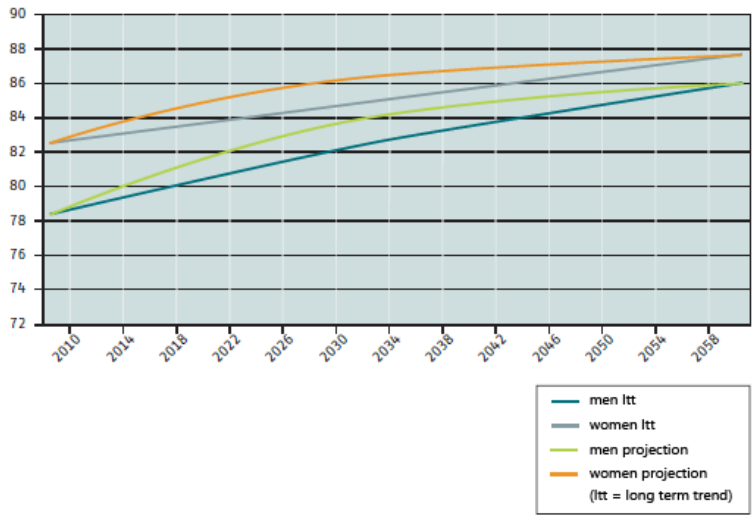

Figure 3.4: *Projected period life expectancies at birth according to AG: only long term trend and full model. Source: Actuarieel Genootschap (2010). Translated legend.*

Another peculiarity of the forecast by AG, is that they use mortality rates which are averaged over two years.[2] This reduces the irregularity in the progression of mortality rates, but might also cause a (small) lag in monitoring the trend in mortality:

$$q_{x,t,t+1} = 0.5(q_{x,t} + q_{x,t+1}). \tag{3.4}$$

[2]In earlier projections they averaged mortality rates over five years.

The trends are then modeled according to:

$$f_{\text{short}}(x) = \sqrt[6]{\frac{q_{x,2007,2008}}{q_{x,2001,2002}}} \tag{3.5}$$

$$f_{\text{long}}(x) = \sqrt[20]{\frac{q_{x,2007,2008}}{q_{x,1987,1988}}}. \tag{3.6}$$

The start table and goal table are formulated as

$$q^{\text{start}}_{x,2008} = q_{x,2007,2008}\sqrt{f_{\text{short}}(x)} \text{ for } x < 94 \tag{3.7}$$

$$q^{\text{start}}_{x,2008} = q_{x,2003,2008}(f_{\text{short}}(x))^2 \text{ for } x > 95 \tag{3.8}$$

$$q^{\text{goal}}_{x,2060} = q_{x,2008}(f_{\text{long}}(x))^{52}. \tag{3.9}$$

Furthermore it is imposed that in every year, death probabilities for women are lower than death probabilities for men of the same age. The results show first of all, as the trend of the past suggests, a decrease in mortality over time for both men and women. What is also observable, is the phenomenon of rectangularization: mortality is concentrated in an increasingly smaller interval. Figures 3.5 and 3.6 show this. Both these effects are more prominently visible for men than for women, meaning that their life expectancies converge over time. The period life expectancy at birth in 2060 is found to be 85.9 years for men and 87.6 years for women. Remaining life expectancies at age 65 are 22.2 years for men and 23.9 years for women. No confidence intervals are provided. If I compare this with the results found by CBS and RIVM, I see that AG especially predicts lower mortality for men. Predictions for mortality of women at higher ages are higher, which can be seen from the remaining life expectancy at age 65.

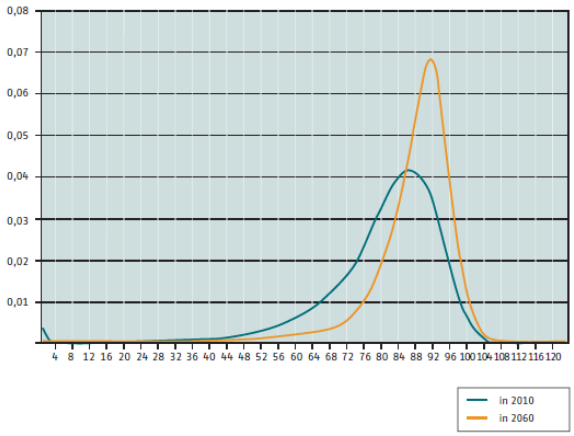

Figure 3.5: *Present (2010) and projected (2060) density of male mortality according to AG. Source: AG (2010).*

Unlike the previous two models, the forecast by AG is entirely based on historical data. Actuarieel Genootschap (2010) make a bald statement in saying that a forecast in fact should solely be based on historical data (p. 10). Furthermore, they distinguish between a short and a long term trend. They argue that this makes the model robust in the long run, but leaves room for some volatility in the short run. I believe however that one should be cautious in implementing such a short run trend, as it is only based on a few number of observations. Another shortcoming of the model is that they, like RIVM, do not assess the uncertainty in the estimates. Note that in September 2012, AG will come with a new projection, possibly based on an entirely different model.

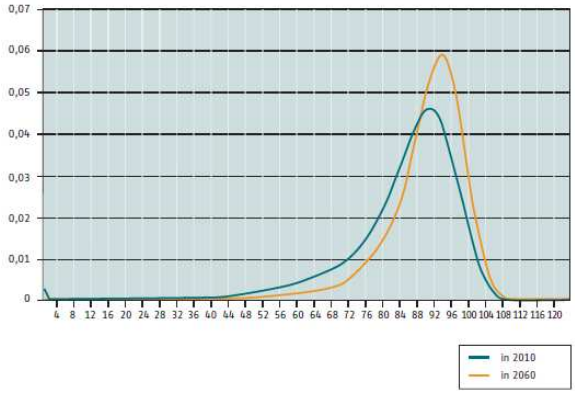

Figure 3.6: *Present (2010) and projected (2060) density of female mortality according to AG. Source: AG (2010).*

3.2.4 PLT

The forecast of the Dutch Association of Insurers (in Dutch: Verbond van Verzekeraars) was made by their workgroup PLT (Pensioen- en Lijfrentetafels). Like the committee that performed the forecast of AG, this workgroup consists of a group of actuaries active for different companies in the field. Their approach taken is also very similar to that of AG. As the latter did, PLT (Workgroup PLT, 2010) extrapolates historical data, using a (implied) short term and a long term trend. The forecast is done for the period 2009-2058, so this is a fifty-year forecast. An important difference with the AG forecast is the data used. Where AG uses data starting in 1988, PLT starts observations in 1958. It is questionable whether mortality rates from such a long time ago should still be used, especially since the trend in the sixties and seventies (increasing mortality rates) seems to differ a lot from today's trend. End year of the data used is the same for both models, namely 2008. AG bases its forecast on mortality data of the general population. PLT does not only use these data, but also mortality data on people with a pension insurance at one of the companies that have delivered mortality data to the Association of Insurers. This covers about eighty percent of the Dutch pension insurance market. In this way PLT is able to model both future population mortality and the future mortality of people with a pension insurance. These models are developed apart from each other, so that PLT is also able to come up with estimates for future mortality of the general population.

PLT uses a simple model to assess the decline in future population mortality. For each year, the present death probabilities are obtained by multiplying the death probabilities of the past year by a reduction factor $\alpha_{x,t}$:

$$q_{x,t} = \alpha_{x,t} q_{x,t-1} \tag{3.10}$$

$$\widehat{\alpha}_{x,t} = \sqrt[k']{\frac{Q_{x,N}}{Q_{x,N-k'}}} \text{ where } N = 2008, k = t - N, k' = \max(k, 5), \tag{3.11}$$

where $Q_{x,t}$ is the death probability for 'whole ages', $x = 0, 1, \ldots, 94$, after applying the Van Broekhoven algorithm (Van Broekhoven, 2002). The initial $q_{x,t}$'s from the data are only defined for 'half ages', $x = 0, \frac{1}{2}, \ldots, 98\frac{1}{2}$, but after performing the algorithm they are replaced by the $Q_{x,t}$'s . Now there are estimated reduction factors $\widehat{\alpha}_{x,t}$ for both genders, which are used to project future death probabilities for the Dutch population:

$$q_{x,N+k} = \widehat{\alpha}_{x,N+k} q_{x,N+k-1} \text{ for } k \in [1, 2, ..., 50]. \tag{3.12}$$

In the way that the reduction factors are designed, there appear both a short run and a long run trend in the model.

PLT find a period life expectancy at birth in 2058 of 86.45 years for men and 88.17 years for women. No confidence intervals are given. Remaining life expectancies at age 65 are 22.86 years for men and 24.97 years for women.[3] It is important to notice that these numbers are still based on the Dutch population as a whole and not on people with a pension insurance. Already the numbers are higher than in the other forecasts, except for the life expectancy for women which is higher in the forecast by RIVM. Figure 7 shows the progression of life expectancies for men and women according to PLT. In section 4.2 I will discuss how PLT models portfolio mortality for people with a pension insurance at one of the companies providing data.

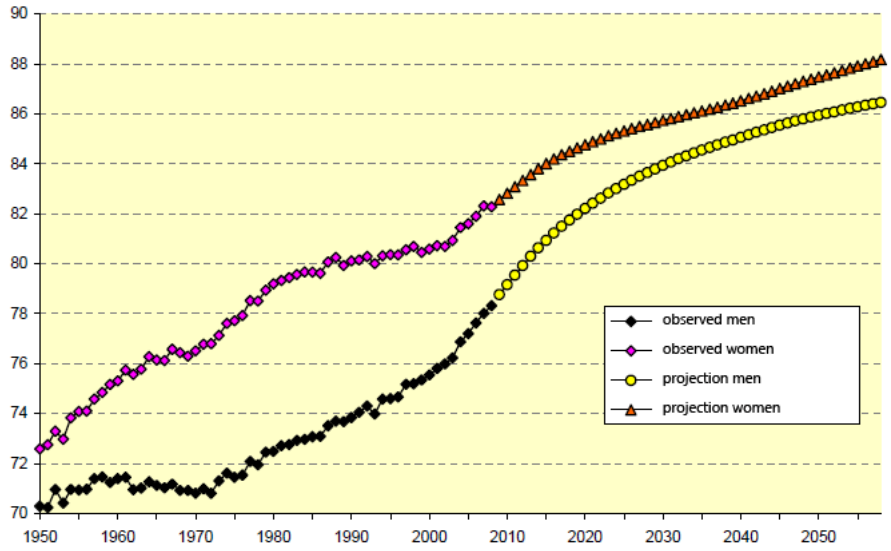

Figure 3.7: *(Projected) period life expectancies at birth according to PLT. Source: Workgroup PLT (2010). Translated legend.*

To me, the longevity model of PLT is very similar to the one developed by AG. Both have a short term and a long term trend and both provide no measure for the uncertainty in their forecasts. The forecast of PLT leads to slightly higher life expectancies and thus slightly higher reserves for the insurer. I can however find no scientific reason to prefer either one of the models.

If I compare the models presented in this section to the models from the academic literature, I can first of all say that there is far less attention for the uncertainty surrounding the point estimates of future life expectancies. Three of the four models do not give confidence intervals, the fourth (CBS) bases the confidence intervals not on statistically sophisticated procedures but on estimation errors obtained in the past. Furthermore, this model is not entirely based on extrapolative methods, but also on opinions of (medical) experts about the development of mortality for multiple causes of death. These are important reasons for me to prefer one of the scientific models. In longevity risk the uncertainty in future death rates and life expectancies is essential for issuers of annuities. Note that this is also the main point of criticism by Van de Poel et al. (2010), who were asked to evaluate the report of Actuarieel Genootschap (2010). The use of expert opinions is no option for me, as I have no medical background. I therefore choose to use one of the academic methods for modeling the (future) population mortality.

[3]Source: own calculations.

3.3 Approach taken

As already mentioned in Chapter 1, I will model population mortality and portfolio specific mortality separately. This is because I simply do not have enough data to use a model that allows for simultaneous modeling. See also section 4.3 of this thesis.

I have decided to use the Lee-Carter method (Lee and Carter, 1992) to predict future mortality of the general Dutch population. Over the years this method has evolved following proposals by other scientists. The reason that I have chosen for the Lee-Carter method, is that it is a relatively simple model, but not less accurate compared to other models. In section 3.1 I have already discussed some of the other models developed in the past. The CBD-model is for instance also a model that has empirically proven its validity. A strong disadvantage of this method is however that it limits the analysis to ages over 40. Since I would eventually want to use the population mortality model for mortality in the ASR pensions portfolio, which also has a significant number of insureds with age below 40, this model is no option for me. The APC variant of the Lee-Carter method including a cohort effect would also perform well, but as I can (in the literature) not find evidence of a significant cohort effect for the Netherlands it would most likely only complicate the analysis, without significantly improving the outcomes. The Poisson model does in essence not really deviate from the Lee-Carter method, but statistically it is more sophisticated. I do however not need to use this variant, since I can quantify the parameter risk by means of a bootstrap method and there are no empty data cells (combinations of age and year for which no deaths are observed in my data). Furthermore, no model from the academic literature has shown to systematically outperform the (original) Lee-Carter method. See for instance Cairns et al. (2009) and Dowd et al. (2010). The models described in section 3.2 that are used in practice do not provide ways to quantify uncertainty (AG, PLT, RIVM) or are (partially) based on expert opinions rather than extrapolation (CBS, RIVM), which makes them not attractive for me. The (statistical) weak sides to the Lee-Carter model that do exist, I will try to limit by using a slightly different approach than Lee and Carter (1992) originally did. These alterations are based on proposals by other scientists and have throughout the years become commonly used. I will carefully mention throughout this section where the approach used here deviates from the original Lee-Carter method.

3.3.1 Data

Mortality data that I have used for the population longevity model comes from Human Mortality Database (HMD), which makes use of numbers provided by CBS. I have used separate data for men and women. Both samples contain the period starting in 1970 ($t = t_1$) and ending in 2009 ($t = t_n$), which is the latest year for which data from HMD is available. I choose to go back not further than 1970, as the trend of increasing mortality rates in earlier years does, in my opinion, not seem to be a reasonable indicator for the future. As the Lee-Carter method is quite sensitive to the selection of the starting year, I will also run the model for some other starting periods. This I will discuss at the end of subsection 3.3.2. Death rates are provided for one-year age intervals and one-year period intervals. The first age group ($x = x_1$) is that of persons aged 0, the last age group ($x = x_m$) is that of persons aged 98. HMD provides data up to age group 110+, but I do not consider this last part. From the age of 98 onwards, mortality rates are based on some smoothing procedure performed by the Netherlands Interdisciplinary Demographic Institute (NIDI). The exact procedure that NIDI has used is however not specified (see Jasilonis, 2011). Since I want to develop future life tables ending at the age of 120, I would have to extend the tables anyhow. I therefore choose to only consider the non-smoothed data up to age $x = 98$, which is entirely based on numbers provided by CBS. Resulting life tables I will extend to higher ages, up to 120, by means of a procedure explained in subsection 3.3.2.

3.3.2 Lee-Carter model

Lee and Carter (1992) aim to estimate the natural logarithm of central death rates by

$$\log m_{x,t} = a_x + b_x \kappa_t + \varepsilon_{x,t}, \tag{3.13}$$

where a_x and b_x are age-dependent parameters and κ_t is time-dependent. Note that throughout this thesis I will use $x = x_1, ..., x_m$ and $t = t_1, ..., t_n$. It is easily seen that equation (3.13) is not yet uniquely defined. Therefore, Lee and Carter (1992) impose the restrictions $\sum_x b_x = 1$ and $\sum_t \kappa_t = 0$. When these restrictions hold, the role of the parameters can be interpreted as follows. a_x is referred to as the 'general shape of mortality across age'. The parameter b_x indicates the sensitivity to changes in κ_t. In interpretation, b_x can be seen as the percentage of 'profit' in total mortality, gained by people in age group x when medical care advances. It tells us at which ages mortality tends to decline rapidly and at which ages the decrease is slower. A negative value for b_x indicates that mortality rates for age group x increase as the general mortality pattern (represented by $kappa_t$) decreases. Note that this parameter is assumed to be constant over time. In reality this assumption might very well not hold, see Lee (2000). He however argues that it is not clear whether altering the method such that b_x can vary over time yields more accurate forecasts. The parameter κ_t tries to capture the changes in the general mortality pattern over time. The error term $\varepsilon_{x,t}$ is assumed to have mean 0 and variance σ_ε^2. Furthermore, it is assumed that the $\varepsilon_{x,t}$'s are independent over age and time.

To find the solution of this problem is to solve

$$\min_{\widehat{a},\widehat{b},\widehat{\kappa}} \sum_x \sum_t \left(\log m_{x,t} - \widehat{a}_x - \widehat{b}_x \widehat{\kappa}_t \right)^2. \tag{3.14}$$

Since there is no observable quantity on the right hand side of this equation, Ordinary Least Squares cannot be used. Lee and Carter (1992) propose to first of all estimate a_x by

$$\widehat{a}_x = \frac{1}{n} \sum_{t=t_1}^{t_n} \log m_{x,t}. \tag{3.15}$$

This follows from the restrictions on b_x and κ_t and from setting the derivative of (3.14) with respect to a_x equal to zero. The parameters b_x and κ_t can then be estimated by performing a Singular Value Decomposition on the matrix A,

$$A_{xt} = \log m_{x,t} - \widehat{a}_x \quad \forall x, t, \tag{3.16}$$

into

$$A = USV. \tag{3.17}$$

Here, U is an orthogonal, $(m \times m)$ matrix. S is a diagonal $(m \times n)$ matrix and V is an orthogonal matrix of size $(n \times n)$. One then gets

$$\widehat{b}_x = \frac{1}{c} s_{11} u_1 \tag{3.18}$$

$$\widehat{\kappa}_t = c v_1, \tag{3.19}$$

with u_1 the first column vector of U, v_1 the first column vector of V and with c equal to the sum of all elements of the vector $s_{11}u_1$. This is the factor which ensures that all $\widehat{b}_x$'s sum to 1. Combining this restriction with the way $\widehat{a}_x$ is estimated, it is ensured that all $\widehat{\kappa}_t$'s sum up to zero. It should be noted that problems occur in case the number of deaths $D_{x,t}$ would equal 0 for some pair (x, t). This would lead to $m_{x,t}$ being equal to 0, which would mean that $\log m_{x,t}$ is not defined. In the data I have used, this does however not occur. Note that in the way the Lee-Carter model is

defined, each data year is treated as equally important. This is opposed to what AG and PLT do, as they incorporate a short term trend that implicitly weights recent observations more heavily.

In their paper, Lee and Carter (1992) continue with a second step of estimation, which is meant to exactly fit $\widehat{\kappa}_t$ to the observed death rates in each year. This step is highly criticized (see for instance Girosi and King, 2007) and I have decided to drop it. Other methods of estimating are possible as well, for instance via Maximum Likelihood or Weighted Least Squares. The interested reader can consult Wilmoth (1993) for detailed information on these alternative estimation techniques.

These ways of estimating typically let one obtain estimates that follow a quite irregular pattern, especially for the parameters b_x and κ_t. Intuitively this is not desirable, since one would want smooth patterns for future mortality. It does not seem logical for instance that persons aged 50 have profited a lot from decreasing mortality (i.e., high b_{50}), persons aged 51 profit significantly less (lower b_{51}) and persons aged 52 again profit much (high b_{52}). Therefore, most of the time smoothing procedures are used to obtain a smooth pattern for the parameter b_x. This is what I will do as well. One can roughly distinguish between two ways: ordinary polynomial regression on b_x after estimation and a penalized least squares approach during estimation. The advantage of the former is that it is easily implemented and does not cost much computing power. The advantage of the latter is that it also takes into account possible changes in κ_t resulting from the smoothing of b_x. Beforehand one cannot say which method is better in terms of 'value for money'. Therefore it is recommended to use both and compare the results. Unfortunately I was not able to correctly implement the penalized least squares approach, which forced me to stick to the polynomial regression after smoothing. In Appendix E of this thesis I present the algorithm on which I have based my efforts for the penalized least squares smoothing. Delwarde et al. (2007) presents a more elaborate discussion on this approach. Note that to obtain a smooth mortality pattern in the future, smoothing is not necessary for κ_t, since the way of forecasting already ensures a smooth pattern.

Another point of criticism with regard to the original Lee-Carter method is that parameter risk is neglected in Lee and Carter (1992). I will include parameter risk in my model, by using a bootstrap method. After estimating the parameters based on the data, as described above, I use these parameters to draw 10,000 sample central death rate trajectories of the past. This is done by means of the obtained parameter values and a sample from the empirical distribution of $\varepsilon_{x,t}$. This method is called *residual bootstrap*.

An alternative bootstrapping approach would be to not take the empirical distribution of $\varepsilon_{x,t}$, but to assume a parametric distribution and to sample from that one. This distribution would typically be the Normal distribution, as this is the distribution under which the model is defined. For this *parametric bootstrap*, σ_ε^2 is estimated by

$$\widehat{\sigma}_\varepsilon^2 = \frac{1}{m(n-1)} \sum_{x=x_1}^{x_m} \sum_{t=t_1}^{t_n} \left(\log m_{x,t} - \widehat{a}_x - \widehat{b}_x \widehat{\kappa}_t \right)^2. \tag{3.20}$$

Analysis of my observed residuals however told me that the errors in fact did not seem to follow a Normal distribution. The peak in the empirical distribution was higher than in the Normal and the tails were thinner. Via a Jarque-Bera test, the null hypothesis that the residuals followed a Normal distribution was rejected at the 95% confidence level. This is why I decided to adopt the residual bootstrapping method, even though this slightly limits uncertainty (not every possible value around 0 can be reached, but only the ones which have occurred). An overview of the possible bootstrap methods can be found in Pitacco et al. (2009).

After having compiled the 'simulations of the past', I again estimate the parameters for each simulation. In this way I end up with 10,000 sets of parameters. So essentially I have constructed

10,000 Lee-Carter models, each of which I can extrapolate to the future. By means of sampling with replacement from these models, I can come up with simulations for the future that include parameter risk. Later in this section I will elaborate on this.

The extrapolative nature of the model is something which can be criticized as well. It can be dangerous to base predictions for the future on results from the past, since there is no guarantee that the observed trends will continue. Furthermore, in this way any possible information already known about future patterns cannot be incorporated (see Lee, 2000). However, as Wilmoth (2000) argues, there are no forecasting methods that by definition perform better than extrapolative methods.

At first I will project future mortality rates up to the year 2060, which is the same year for which the projection by Actuarieel Genootschap (2010) ends. Later, when I apply the mortality model to the determination of the Solvency Capital Requirement, I will use projections for a longer time horizon. This is done in Chapter 5. Starting year of the predictions will for now be 2010, as the last year for which HMD has data available is 2009. The prediction of future mortality is done by extrapolating the time series κ_t. Lee and Carter (1992) found for their dataset on US mortality, sexes combined, that it followed an ARIMA(0,1,0) model, i.e. a random walk with drift:

$$\kappa_t = c + \kappa_{t-1} + \eta_t, \tag{3.21}$$

where η_t follows a Normal distribution with mean 0 and homoskedastic variance σ_κ^2. For other datasets, standard Box and Jenkins methods could suggest a slightly different structure for κ_t. For instance, Brouhns et al. (2002) found for their Belgian mortality data that an ARIMA(0,1,1) model would be the best fit. However, usually the random walk with drift model is adopted, as it has proven to be a good fit for at least 10 data sets of countries with low mortality (Lee and Miller, 2001). Based on these empirical facts, I will adopt the ARIMA(0,1,0) model as well. Still my time series model will be a bit different than the one by Lee and Carter (1992). This is because in my analysis I make the distinction between men and women, which Lee and Carter (1992) did not. As one would expect, the time trends in mortality rates for men and women in the same population are (strongly) correlated. I therefore need to take this correlation into account when modeling the separate time series κ_t^m for males and κ_t^f for females. Li and Lee (2005) propose the modeling of a common and a separate trend, by applying SVD on first the common and then the separate data. I however choose for the Vector version of the ARIMA model, which is the approach typically adopted in the literature:

$$\kappa_t^m = c_1 + \kappa_{t-1}^m + \eta_{1,t} \tag{3.22}$$

$$\kappa_t^f = c_2 + \kappa_{t-1}^f + \eta_{2,t}, \tag{3.23}$$

where η_t has mean vector 0 and variance-covariance matrix Σ. I can denote the difference between the values of κ_t and κ_{t-1} by $\Delta\kappa_t$, with

$$\Delta\widehat{\kappa}_t = \kappa_t - \kappa_{t-1}. \tag{3.24}$$

Estimation of drift parameters c_1 and c_2 in the Vector ARIMA(0,1,0) model then goes as in a normal ARIMA(0,1,0) model, where

$$\widehat{c} = \frac{1}{n-1}\sum_{t=t_2}^{t_n} \Delta\widehat{\kappa}_t \quad \Leftrightarrow \tag{3.25}$$

$$\widehat{c} = \frac{\widehat{\kappa}_{t_n} - \widehat{\kappa}_{t_1}}{n-1}. \tag{3.26}$$

The variance-covariance matrix Σ has the following structure:

$$\Sigma = \begin{bmatrix} \sigma_m^2 & \sigma_{mf} \\ \sigma_{fm} & \sigma_f^2 \end{bmatrix}. \tag{3.27}$$

The parameters σ_m^2 and σ_f^2 indicate the variance in residuals for men and women respectively. σ_{fm} $(= \sigma_{mf})$ indicates the covariance in the residuals for men and women. Estimation is done according to:

$$\Sigma = \frac{1}{n-2} e \cdot e', \tag{3.28}$$

where e is the matrix consisting of the observed residuals e_t:

$$e_{1,t} = \widehat{\kappa}_{t-1}^m - \widehat{\kappa}_t^m - \widehat{c}_1 \tag{3.29}$$

$$e_{2,t} = \widehat{\kappa}_{t-1}^f - \widehat{\kappa}_t^f - \widehat{c}_2, \tag{3.30}$$

for $t = t_2, \ldots, t_n$. In the original Lee-Carter method, one would come up with point estimates for the future values of κ_t, combined with a analytically derived confidence interval that is based on the assumed Normal distribution of η_t. I will however use a different approach. First I sample with replacement out of the 10,000 sets of parameters obtained[4], to include the parameter risk in the projections. For the specific 'Lee-Carter model' drawn, I estimate the future values for κ_t by means of the equation

$$\widehat{\kappa}_{t_n+\tau} = \widehat{c} + \widehat{\kappa}_{t_n+\tau-1} + \gamma_{t_n+\tau} \tag{3.31}$$

for both genders, where $\gamma_{t_n+\tau}$ denotes a draw from the multivariate Normal distribution of η and accounts for the process risk in the model. By repeating this procedure 4,999 times I obtain 5,000 paths of future developments of κ_t.

For each path of future values of κ_t, I can now derive future death probabilities. Typically one can revert to equation (3.13) and obtain an estimate of the logarithm of central death rates in the year $t_n + \tau$ by filling in the estimates for a_x, b_x (the smoothed values) and $\kappa_{t_n+\tau}$. This method however incorporates a so-called jump-off bias when future mortality is estimated. This bias arises because $\widehat{a}_x$ and $\widehat{b}_x$ are only partially based on the mortality rates experienced in the last sample year t_n. Lee and Miller (2001) therefore propose to use the central death rates of the last sample year as jump-off point for future mortality:

$$\widehat{m}_{x,t_n+\tau} = m_{x,t_n} \exp\left(\widehat{b}_x(\widehat{\kappa}_{t_n+\tau} - \widehat{\kappa}_{t_n})\right). \tag{3.32}$$

This method of estimating future central death rates ensures that a jump-off bias is avoided. I therefore adopt this method in favor of the original one. The central death rates can be used to compute future death probabilities, by means of the general relation

$$q_{x,t} = 1 - \exp(-m_{x,t}). \tag{3.33}$$

If I take these steps for each of the 5,000 simulated paths of κ_t, I cannot only construct a best estimate path for the future death probabilities, but also a confidence interval based on the empirical distribution. In the standard Lee-Carter approach, one would do this based on the (assumed) normality of the error terms γ_t.

At this point I have life tables up to and including age 98. It is however certainly not unlikely that an individual will reach the age of 99, not to mention the fact that today there is already a significant number of people aged above 98 in the population. In the future it might even become more probable that someone survives until this age. Therefore I need to extend the obtained life tables up to the age group 120 (so the maximal age to be reached by an individual, ω, equals 120). Many methods for doing this are available; I choose one proposed by Denuit and Goderniaux (2005). They perform a constrained log-quadratic regression on mortality rates for the old, which they then extrapolate to the very old. In their original method, Denuit and Goderniaux (2005)

[4] $\widehat{a}, \widehat{b}$ and the parameters for the process κ_t.

use this approach on the input data (mortality rates of the past). I will apply their method not to my input data, but to the death probabilities for future years generated by the Lee-Carter method.

The regression Denuit and Goderniaux (2005) consider is

$$\log \widehat{q}_{x,t} = \lambda_t^{(1)} + \lambda_t^{(2)} x + \lambda_t^{(3)} x^2 + \delta_{x,t}, \tag{3.34}$$

where $\widehat{q}_{x,t}$ is the observed (in my approach the fitted) death probability for age group x in year t and $\delta_{x,t} \sim N(0, \sigma_q^2)$. This relation is then fitted to each year t and for a select group of high ages x of which the data is deemed to be still reliable. The two constraints imposed to make sure that indeed no one can survive after age 120 are

$$q_{\omega,t} = 1 \tag{3.35}$$

$$q'_{\omega,t} = 0, \tag{3.36}$$

where ω denotes the maximum age an individual can reach (120 in this thesis). The first constraint makes sure that no one survives after age ω, whereas the second constraint makes sure that the increase over age of the mortality rates slows down at old ages, which is in line with empirical evidence.

Implementing the two constraints yields that (3.34) can be rewritten to (cf. Coelho et al., 2007)

$$\log \widehat{q}_{x,t} = \left(\omega^2 - 2\omega x + x^2\right) \lambda_t^{(3)} + \delta_{x,t}. \tag{3.37}$$

For the years $t = 2010, \ldots, 2060$ I estimate $\lambda_t^{(3)}$ using OLS on equation (3.37) based on the simulated death probabilities for age groups $x = 85, \ldots, 98$. For ages $x = 99, \ldots, 120$ I then estimate the natural logarithm of future death probabilities by filling in the estimate for $\lambda_t^{(3)}$ in (3.37). Inspection of the results suggested that to get a smooth transition, some smoothing around $x = 98$ was needed. For this I adopted the geometric averaging procedure proposed by Denuit and Goderniaux (2005):

$$\widehat{q}_{x,t}^{smooth} = \left(\widehat{q}_{x-2,t} \cdot \widehat{q}_{x-1,t} \cdot \widehat{q}_{x,t} \cdot \widehat{q}_{x+1,t} \cdot \widehat{q}_{x+2,t}\right)^{\frac{1}{5}}, \tag{3.38}$$

for ages $x = 94, \ldots, 103$. Doing this for every simulated path of future death probabilities, I have finally come up with 5,000 population cohort life tables for the age range $x = 0, 1, \ldots, 120$ and years $t = 2010, 2011, \ldots, 2060$.

In the Lee-Carter setting, one can roughly distinguish between three types of risk: process risk, parameter risk and model risk. Process risk is the risk arising from the uncertainty in the future development of the time series κ_t. This process risk is thus quantified by means of the error terms $\gamma_{t_n+\tau}$ in equation (3.31). Parameter risk is the risk that the estimates found for the parameters are not correct and is quantified by means of the bootstrap method described earlier in this section. Model risk is the risk that the model chosen, here the Lee-Carter model, is in fact not the right one. Note that since I have only used the Lee-Carter method, this form of model risk is not included in my analysis. One could include this additional source of risk by assessing projections for future death probabilities using different models. Other candidates could for instance be the Cairns-Blake-Dowd model (Cairns et al., 2006) or the Poisson model developed by Brouhns et al. (2002). Another source of model risk in the Lee-Carter method comes from the period considered. As the Lee-Carter method extrapolates linear trends in mortality patterns, it is very sensitive to the start and end year considered. I have therefore run the model not only for starting year 1970, but also for starting years 1975, 1980, 1985 and 1990. So essentially I have performed the steps described in this section 5 times. By considering all these simulations (in total I have obtained 25,000 future paths of death probabilities), I weaken the sensitivity of the results towards this characteristic. Results of this approach for modeling population mortality are discussed in the next section.

3.4 Results

In section 3.3 I have described the original Lee-Carter method and the commonly used adaptations that I have adopted as well. In this section, the results of this approach for modeling future mortality in the Dutch population are presented and discussed. I will also compare my findings to the results found by CBS, RIVM, AG (Actuarieel Genootschap, 2010) and Workgroup PLT (2010).
In Appendix A some figures are displayed, depicting the results obtained from the first stage of estimation based on the actual data. Since the 10,000 simulations performed thereafter are based on these initial estimates, the average values for the parameters are approximately equal to the first estimates (convergence according to Weak Law of Large Numbers). All figures present the results for men in the upper panel and the results for women in the lower panel. Furthermore it should be noted that all figures are based on starting year of the observation period $t_1 = 1980$. The results obtained for the other starting years are incorporated later, when the future death probabilities etc. are presented. The estimates for a_x and b_x are obtained for the ages 0-98. Figure A.1 indeed presents a pattern for a_x which is consistent with previous results in the literature, see for instance De Waegenaere et al. (2010). At age 0 mortality is quite high due to infant mortality, after which it is decreasing until the age of 10. Afterwards it is approximately linearly increasing, except for the 'accident hump' noticeable for adolescents. Note that this hump is more prominent for males than for females.

Figure A.2 depicts the pattern for b_x together with the smoothed values coming from a polynomial regression. The pattern shows that young children have profited most (high b_x) from the decrease in mortality over time, whereas for the very old mortality rates have tended to increase a bit (negative b_x). Again the pattern shows close resemblance with previous results in the literature, see for instance Stevens et al. (2010a). Note also that variation in the value is higher for lower ages, showing that the mortality rates over time have varied more for the young. It can furthermore be seen that variation is slightly higher for females than for females, which leads to a fit of the smoothed b_x which is a bit worse. The global patterns of b_x for males and females are roughly the same.

Estimates for κ_t are initially obtained for years 1980-2009 and displayed in figure A.3. In the same figure also the predicted future pattern of κ_t for the years 2010-2060 can be seen, together with 95% confidence intervals when only process risk is included and when also parameter risk is included (but not model risk). Note that I choose for the 95% interval because this is the typical quantity considered in the literature. The graph for males shows a steeper trend than for females, indicating that in the period 1980-2009 the drop in mortality rates for men was more prominent than the drop in mortality rates for women. Furthermore, the variation around the trend was in the past higher for women than for men. This results in a higher variance of the process κ_t, which explains why the confidence interval for women is wider than for men. As expected, confidence intervals are wider when also parameter risk is included.

For men, the model explains 81.91% of the total variance in the data. For women, this number is equal to 65.01%. These numbers are lower than typically found in the literature. One should however note that this is mainly because there the age interval considered is much smaller. Brouhns et al. (2002) for instance find numbers slightly under 90%. They however only treat the ages 65-98. If I would limit the analysis to these age groups, I find similar quantities.

The period life expectancy at birth in 2060 for men derived from the best estimate table equals 85.52 years, whereas the 2.5% quantile and the 97.5%quantile (based on the simulated life tables for 2060) lead to period life expectancies of 83.13 and 88.48 years (including the model risk in the way specified before). For women, the best estimate equals 87.25 and the bounds on the 95% confidence level are 84.43 and 89.98. These intervals are not particularly wide when compared to the interval obtained for instance by CBS, which is a typical feature of the Lee-Carter method.

Lee (2000) argues that this arises from the small effects of uncertainty in κ_t on life expectancy when mortality is already low. In 2009, the period life expectancy at birth for men was 78.54 years and for women it was 82.65 years. This means that I find an expected increase of 6.98 years for men and 4.60 years for women, again showing that mortality is expected to decrease more for men than for women when the Lee-Carter approach is used. The point estimates of period life expectancies at birth obtained via the Lee-Carter method are slightly lower than the values obtained by AG. CBS found slightly higher life expectancy for women and a slightly lower life expectancy for men. Results of both forecasts fall well within the 95% confidence intervals. Remaining life expectancies at age 65 are 21.93 years for men and 24.15 years for women. For men, this number is lower than the estimate by AG, but higher than the estimate by CBS. For women the number is higher than the estimate by AG, but (slightly) lower than the one by CBS. The bounds on the 95% confidence intervals are 20.11 and 24.47 for men and 22.01 and 26.39 for women. Again, the point estimates by AG and CBS fall well in these intervals. The forecast by PLT (without their factors for portfolio mortality) cannot directly be compared to the numbers presented here, as it only runs until 2058. For this year, the confidence interval for the Lee-Carter model is slightly narrower, but the values found by PLT lie within the confidence interval of the Lee-Carter method for 2058. The forecast by RIVM even ends in 2050. For both genders the (remaining) life expectancies still fall into the 95% confidence interval for 2050, even though the values found for women are quite higher than expected when the Lee-Carter method is used.

Model	e_0	**Lower**	**Upper**	e_{65}	**Lower**	**Upper**
AG	85.90			22.21		
CBS	84.5			21.7		
Lee-Carter	85.52	83.13	88.48	21.93	20.11	24.47
PLT **(2058)**	86.45			22.86		
RIVM **(2050)**	83.8			21.1		

Table 3.1: *Period life expectancies for men in 2060 (2058 for PLT, 2050 for RIVM) at birth and remaining at age 65, with bounds of the 95% confidence interval for the Lee-Carter model. Numbers for CBS and RIVM are taken from their report, numbers for AG and PLT are found via own calculations.*

Model	e_0	**Lower**	**Upper**	e_{65}	**Lower**	**Upper**
AG	87.57			23.86		
CBS	87.4			24.2		
Lee-Carter	87.25	84.43	89.98	24.15	22.01	26.39
PLT **(2058)**	88.17			24.97		
RIVM **(2050)**	88.1			24.6		

Table 3.2: *Period life expectancies for women in 2060 (2058 for PLT, 2050 for RIVM) at birth and remaining at age 65, with bounds of the 95% confidence interval for the Lee-Carter model. Numbers for CBS and RIVM are taken from their report, numbers for AG and PLT are found via own calculations.*

Figure A.4 shows a plot over time of the logarithm of the central death rates for people aged 65. Here it can be seen that the trend is linearly decreasing on the logarithmic scale, which is due to the linear decreasing trend of κ_t. This decreasing pattern holds for all ages x where $b_x > 0$. For the ages where $b_x < 0$, one would see a linearly increasing pattern in the logarithm of central death rates. When I compare the graph for men to the graph for women, I see that in 2060 it is expected that men aged 65 have about the same probability of dying within one year as women of the same age. This is due to the steeper decrease in mortality rates for men in the past, a trend which is projected onto the future. Figure A.3 showed that the variance in the process κ_t is larger for women than for men, but this cannot be observed in figure A.4. It should be noted that for

some future years, the upper 95% confidence bound is lower for the variant where no model risk is considered. This comes from the trend being slightly different when estimates from all 5 different start periods are considered. The confidence intervals when model risk is considered are however always wider than when model risk is neglected.

Chapter 4

Portfolio specific mortality

In this chapter I will discuss the modeling of mortality for a specific subpopulation, in this case an annuity portfolio of an insurer (such as ASR) or pension fund. In section 4.1 I will discuss the models developed in the academic literature. Section 4.2 presents an approach to modeling portfolio specific mortality as it is done in practice in the Netherlands. Section 4.3 presents the approach I have used for modeling the portfolio specific mortality in the pensions portfolio of ASR. Furthermore it presents the approach used for combining the models for population mortality and portfolio mortality, to come up with a stochastic model for future mortality in the ASR pensions portfolio. Results of this are presented in section 4.4.

4.1 Models from academic literature

In the literature, a lot of research concerning mortality patterns of entire populations (as discussed in section 3.1) is performed. These models can be used to predict future mortality rates and the uncertainty (longevity risk) surrounding these estimates. For insurers (and pension funds) however, it is almost equally important to know how mortality in their portfolio relates to general mortality in the country in which they are active. Therefore, also models have been developed to quantify these specific relations. It should however be noted that the number of papers written in this field is far less than the number of papers written on general mortality patterns. This would not be a problem if the data set of the insurer is of such a size that it can be seen as a specific population itself (not only in number of clients, but also in number of years for which data is available). In that case, the insurer can just apply one of the models discussed in the previous section to its own data set. Typically however the number of clients is considerably smaller and reliable data is only available for a small number of years. This is why an insurer will often need to revert to a model for population mortality to predict future development in mortality rates. The specific relation between mortality in the portfolio and mortality in the population can then be applied to this model.

In the same paper in which Brouhns et al. (2002) present their Poisson model for population mortality, they also come up with a Cox Proportional Hazard model to study mortality of Belgian annuitants relative to the general Belgian population. This model is based on the linear relationship on the logarithmic scale, which was already observed by Brouhns and Denuit (2001). The specification of the model is as follows:

$$\log m_{x,t}^{PORT} = \theta_1 + \theta_2 \log m_{x,t}^{POP} + \epsilon_{x,t}, \tag{4.1}$$

where $m_{x,t}^{PORT}$ denotes central death rates for the portfolio considered, $m_{x,t}^{POP}$ denotes central death rates for the general population and $\epsilon_{x,t}$ is the iid error term with mean 0 and variance σ_ϵ^2. If this relation remains valid over time (I will discuss this issue later in this section), I can relate

the future mortality rates in the portfolio to the future mortality rates of the population via:

$$\widehat{m}_{x,t}^{PORT} = \exp(\theta_1) \left(\widehat{m}_{x,t}^{POP}\right)^{\theta_2}. \tag{4.2}$$

Denuit (2007) uses this same relation for Belgian group annuitants and Belgian individual annuitants. For the group annuitants he finds the coefficient of determination R^2 to equal 88.63%, for the individual annuitiants he finds an R^2 of 81.72%.

Assuming that the error terms in equation (4.1) follow a Normal distribution, the left hand side in that equation follows a Normal distribution as well, with mean $\theta_1 + \theta_2 \log m_{x,t}^{POP}$ and variance σ_ϵ^2. The parameter risk can be quantified as well. In this way, it is possible to come up with a probability distribution for mortality in a portfolio, given the corresponding numbers for mortality in the population.

Using the Cox Proportional Hazard model, one would typically find $\theta_2 <> 1$. This means that if there is a decrease in future mortality, improvements in mortality rates go at a different rate for the general population than for the people in the portfolio of the insurer according to equation (4.2). It is questionable whether this is a desirable property. I would however like to point out that there is no consensus on whether the mortality rates of subsets of the population do or do not improve at different rates. Bruggink (2009) for instance argues that differences in life expectancies for different socio-economic groups in the Netherlands have not changed over time. Pitacco et al. (2009) on the other hand suggest that differences in mortality between different socio-economic groups have widened over time. Furthermore, if one considers a different scale, for instance when one looks at the development of q_x instead of m_x, the relation will change anyhow.

In the same book, Pitacco et al. (2009) propose a new model for portfolio mortality which assumes the rate of decline in mortality to be the same for both the general and the insured population. Also, they assume the relation between population mortality and portfolio mortality to be constant over time. The model is defined by the following equation:

$$\log m_{x,t}^{PORT} = f(x) + \log m_{x,t}^{POP} + \psi_{x,t}, \tag{4.3}$$

where $\psi_{x,t}$ is the iid[1] error term with mean 0 and variance σ_ψ^2. So for each age x, one can approximate the portfolio mortality by the population mortality plus a constant depending on x. Given that the portfolios that Pitacco et al. (2009) consider consist of people who typically have lower mortality than the general population, they find that for most x, $f(x) < 0$. The fit of this model is even higher than that of the Cox Proportional Hazard model, at least for the data used by Pitacco et al. (2009). They consider data of mortality in Belgium, distinguished by gender and by type of annuity (individual or group). They find coefficients of determination R^2 between 97.2% and 99.8% for males and females respectively. Note that in all these papers the analysis is limited to the ages 65-98. Confidence intervals for the regression are constructed as for the Cox Proportional Hazard model.

Another model which I would like to mention is the one developed by Plat (2009). The model he considers is stochastic of nature and makes no assumption about the relation between population and portfolio mortality being constant over time or not. Plat (2009) wants to estimate portfolio mortality factors $P_{x,t}$ for ages $x = x_1, \ldots, x_m$ and years $t = t_1, \ldots, t_n$ according to:

$$P_{x,t} = \frac{q_{x,t}^{port}}{q_{x,t}^{pop}}, \tag{4.4}$$

where $q_{x,t}^{port}$ denotes the one-year death probabilities as experienced in the portfolio and $q_{x,t}^{pop}$ represents the one-year death probabilities for the general population.

[1] Independent and identically distributed.

For every year t, Plat (2009) employs a regression model to approximate the vector of k portfolio mortality factors in year t:

$$P_t = 1 + X_t\beta_t + \varepsilon_t, \tag{4.5}$$

where X_t is the matrix of explanatory factors during year t (size $m \times k$), β_t is the corresponding vector of coefficients and ε_t is the error term. Plat (2009) also imposes a restriction to ensure that $P_{x,t} = 1$ for $x = x_m$:

$$\sum_{i=1}^{k} X^i_{x_m,t}\beta^i_t = 0, \tag{4.6}$$

where i denotes the explanatory factor (column of X). This restriction resembles the assumption that the effects in mortality of adverse selection disappear over age.

This regression model could be fit using Ordinary Least Squares. Plat (2009) argues however that there is significant heteroskedasticity present in the model, as the factors $P_{x,t}$ are a result of different numbers of deaths. Therefore he proposes to use Generalized (or Weighted) Least Squares (see for instance Wilmoth, 1993), based on a quantity like the observed number of deaths. For each year t, the coefficients β_t can then be estimated according to

$$\widehat{\beta}_t = (X'W'WX)^{-1}X'W'W(P_t - 1), \tag{4.7}$$

where W denotes the weighting matrix. Doing this for every year t, Plat (2009) ends up with a time series of β. Using standard Box-Jenkins methods, he can then find a suitable stochastic process. From the argument of biological reasonableness, Plat (2009) argues that the process should be stationary. If one does not impose this restriction, it can be that $P_{x,t}$ becomes equal to 0 for some pair (x, t), meaning that people aged x cannot die in year t.

Plat (2009) applies this model to data for a large portfolio of about 100,000 insured males above the age of 65, and to a medium-sized portfolio of about 45,000 insured males aged above 65. For these portfolios he finds an $AR(1)$ time series to be most appropriate, meaning that β is modeled according to

$$\beta_t = c + \phi\beta_{t-1} + \eta_t, \tag{4.8}$$

where c and ϕ are constants and η_t is the error term. Figure 4.1 shows the resulting portfolio mortality experience factors $P_{x,t}$ for the year 2016, together with their 99% confidence interval. Note that since the process is stationary, other years would present a similar graph. As the figure shows, uncertainty in the value of $P_{x,t}$ increases as the size of the portfolio decreases. This is a demonstration of the weak side of the model: one needs a large data set. This concerns not only the amount of insureds in the portfolio, but also the number of years during which mortality in the portfolio is registered. Furthermore, uncertainty is larger for lower ages (at least until age 65). Probably this uncertainty will continue to increase as age gets lower. This makes it a less desirable model to adopt when for instance a portfolio of deferred annuities is considered, where insureds can well be younger than 65.

Unlike in the previously mentioned models, Plat (2009) does find ways to combine the stochastic characteristics of both the population and portfolio mortality model. In order to simulate mortality rates for both the population and the portfolio, he needs to know more about the correlations. He uses the technique of Seemingly Unrelated Regression (SUR), which imposes the need to use the same historical observation periods for both population and portfolio mortality. First of all, Plat (2009) assumes that mortality as observed in the population can be written according to

$$\alpha^k_t = X^\alpha_k\eta^\alpha_k + \varepsilon^\alpha_k \quad \text{for } k = 1, \ldots, m. \tag{4.9}$$

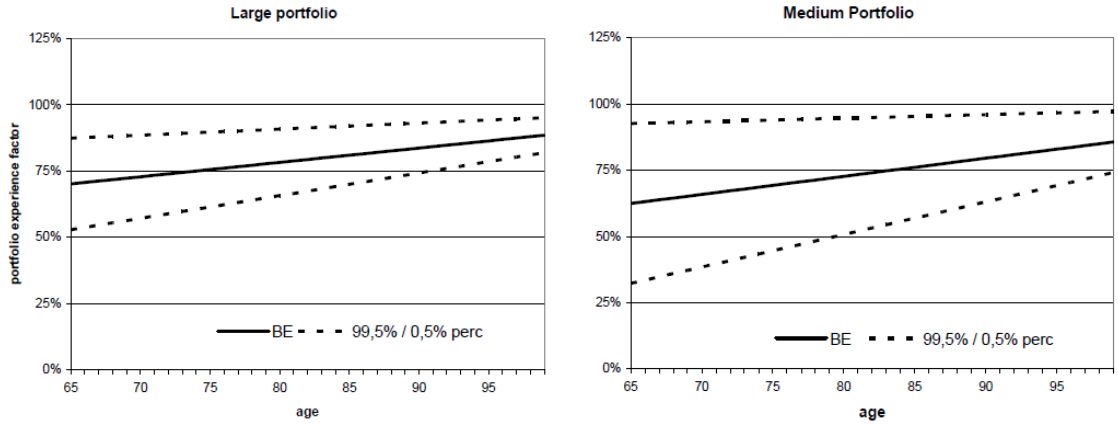

Figure 4.1: *Plot of $P_{x,2016}$ for a large and a medium-sized portfolio, together with 99% confidence interval. Source: Plat (2009).*

The technique of SUR does not require the processes to be similar, hence the name. Therefore, if (4.8) is for each element i rewritten in the more general form

$$\beta_t^i = X_i^\beta \eta_i^\beta + \varepsilon_i^\beta, \tag{4.10}$$

the combined process can be written as

$$Y = X^{\alpha,\beta}\eta^{\alpha,\beta} + \varepsilon. \tag{4.11}$$

The combined processes can then be fitted by first estimating the parameters equation by equation, by means of Ordinary Least Squares. The residuals of the process can then be used to estimate the covariance matrix $\widehat{\Sigma}$. The last step is then to estimate the factors $\widehat{\eta}^{\alpha,\beta}$ using Generalized Least Squares, meaning that it is estimated by

$$\widehat{\eta}^{\alpha,\beta} = (X^{\alpha,\beta'}\widehat{\Sigma}^{-1}X^{\alpha,\beta})^{-1}(X^{\alpha,\beta'}\widehat{\Sigma}^{-1}Y). \tag{4.12}$$

In this way, both the uncertainty in the population mortality model and the uncertainty in the portfolio mortality model are represented in the combined model.

4.2 Portfolio mortality in practice

So far, the only Dutch organization in practice that aims to model portfolio mortality is the Dutch Association of Insurers[2]. In the same report in which Workgroup PLT (2010) presents their model for (future) population mortality, they also demonstrate their model for portfolio specific mortality. This model is based on the mortality as experienced by the pension insurers that have provided data. These companies account for approximately eighty percent of the pension insurance market. It is important to note that ASR is not among the companies that have provided data. Therefore this data need not be representative for the mortality ASR experiences. Workgroup PLT (2010) imposes a simple model, characterized by the relation

$$q_{x,t}^{PORT} = e_x q_{x,t}^{POP}, \tag{4.13}$$

where $q_{x,t}^{PORT}$ reflects the one-year death probabilities of the insureds, the $q_{x,t}^{POP}$'s are the death probabilities of the general population and e_x is a portfolio mortality factor dependent of age. PLT

[2]It is expected that the new projection by AG, scheduled to be published in 2012, will also contain a model for portfolio mortality.

assumes that the relation between portfolio mortality and population mortality does not change over time. So future portfolio mortality can be projected based on a forecast of the population mortality and the relation described above. The actual estimation is only done for the ages 29.5 up to 94.5. For lower ages, they assume the factors to be constant at the level of age 29.5. Also, they assume that the effects of adverse selection have disappeared at age 104.5 ($e_x = 1$ for this age and higher). Between ages 94.5 and 104.5, PLT assumes a linear relation, which they extrapolate until value 1 is reached. For men, PLT finds the relation (4.13) for ages 29.5-94.5 to be linear. For women however they find a quadratic fit. The fits are displayed in figure 4.2.

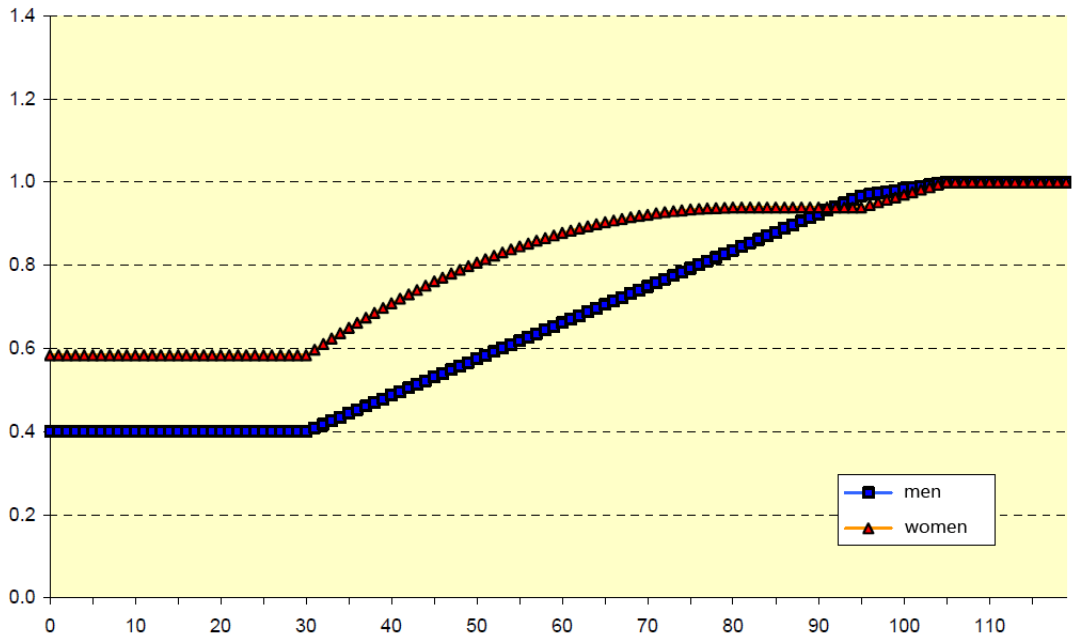

Figure 4.2: *Portfolio mortality factors for pension insurers. Source: Workgroup PLT (2010). Translated legend.*

PLT is able to provide life expectancies in 2058 (the final year of their projection) for people with a pension insurance at one of the companies that have provided data, as described above. They find a (period) life expectancy at birth of 87.93 years for men and of 88.90 years for women. Remaining life expectancies at age 65 in 2058 are 23.77 years for men and 25.39 years for women.

In terms of complexity and the amount of data needed, this model is comparable to the Cox Proportional Hazard model (Brouhns et al., 2002 and Denuit, 2007). It is slightly more flexible, as it does not impose a linear relation over age between population and portfolio mortality. It shows therefore resemblance to the method by Pitacco et al. (2009). The main difference is that the analysis is performed on another quantity ($\log m_{x,t}$ for Pitacco et al., 2009, versus $q_{x,t}$ for Workgroup PLT, 2010). The model by Plat (2009) is more advanced and statistically sound, but also requires the availability of more data. However, one might suspect that the data set that PLT has available is at least as large as the data set used by Plat (2009). In the next section I will present the approach I have used for modeling portfolio specific mortality.

4.3 Approach

In this section I will discuss the way I model the mortality as experienced in the pensions portfolio of ASR. The first subsection is about the data I have used, the next about the model I have

adopted for my analysis.

The data I was provided with comes from the administration system of ASR and spans the period from 2002 up to and including 2011[3]. Unfortunately the data from 2010 and 2011 is not of much use, since I have no access to population mortality data for these years and I can therefore not compare the mortality in population and portfolio for those years. This is why my analysis is only based on the years 2002 up to and including 2009. I have chosen to limit myself to started old age pensions, deferred old age pensions and started surviving dependants' pensions. The reason for this comes from the incorrectness in the administration system which seems to be present for the deferred surviving dependants' pensions. Especially for the very old ages a lot of policies are still in the system, i.e., the clients are still alive to ASR, while in reality they have most likely already died. The policies which I now consider should be of such nature that ASR has good knowledge of whether the clients are still alive. In the past however sometimes a clean-up of the administrative system has been performed, which showed that there were still some deaths not processed in the system. It might therefore be the case that the numbers provided are a slight underestimation of the actual death rates. This should be taken into account when drawing conclusions, but a further investigation of the data quality will not be a focus point of this thesis.

For every year mentioned I was provided with the number of insureds per age and gender, together with the total insured amount for that group. Also, for every year I received the number of deaths per age and gender, again together with total insured amount for that group. Using these numbers I could find the observed death probabilities for each age. I have chosen to consider death probabilities not based on number of deaths, but on insured amounts 'passed away'. Typically insured amounts are higher for higher educated people, who have higher life expectancies (Bruggink, 2009). Therefore this measure is convenient when investigating the effects of adverse selection in the portfolio, which is why this is the method typically used in practice. This means that in my approach, in equation (4.1) $\log m_{x,t}^{PORT}$ should be replaced by $\log \widetilde{m}_{x,t}^{PORT}$, with

$$\log \widetilde{m}_{x,t}^{PORT} = \frac{\widetilde{D}_{x,t}}{\widetilde{E}_{x,t}}, \tag{4.14}$$

where $\widetilde{D}_{x,t}$ is the total insured amount belonging to people aged x that have died in year t and $\widetilde{E}_{x,t}$ is the exposure-to-risk for people aged x in year t, measured in insured amounts. Like for the normal exposure-to-risk, the latter approximately equals the insured amount at the beginning of year t minus half the insured amount belonging to the people who die during year t.[4]

As the sample is not that big (number of deaths for men aggregated over all age groups and the total period: 7,958; number of deaths for women aggregated over all age groups and the total period: 4,923), it was not possible to study the specific portfolio mortality for each year separately. Instead, I studied the total horizon and the horizon split in two parts (2002-2005 and 2006-2009).

4.3.1 Model for portfolio mortality

After having inspected my data, I decided to go for the Cox Proportional Hazard model, presented by Brouhns et al. (2002) and Denuit (2007). For this model I obtained a coefficient of determination R^2 of 98.90% for men in the age groups 42-96, over the period 2002-2009. For women aged 44-98 I obtained an R^2 of 97.74% over the same period. The fits for those groups in the model by Pitacco et al. (2009) were only 74.66% and 25.40% respectively. For the approach by Workgroup PLT (2010), the fits are 97.81% and 94.48%. The fit of the Cox Proportional Hazard model is thus the highest, which is why I choose for this approach. The objection raised by Pitacco et al. (2009) concerning the different pace of mortality improvements is in my opinion not a strong one,

[3]Thanks go to Tonny Knake for extracting the data

[4]See Chapter 2.

since experts do not even agree about the relation that there has been in the past. Unfortunately I do not have enough data to test myself how the relation in the ASR pensions portfolio evolves over time. I therefore need to make an assumption regarding this development. I adopt the most frequently used one, namely that it does not change over time.

I limit my analysis to the age ranges for which the exposure is such that I deem the estimated death probabilities in the portfolio to be reliable. The criterion I have used concerns the insured amount of people of certain age that have died during the observation period. The first age group I consider is the first one for which this amount lies above 25,000 euro, the last age group is the last one for which this amount is above 25,000 euro. In section 4.4 I will go into more detail on this and present the results of this part.

Ideally, I would have used the model proposed by Plat (2009). Unlike the models mentioned above, this method leaves room for the stochastic nature of portfolio mortality relative to population mortality and allows the relation to vary over time. However, this model also has the need for larger amounts of data than the other models. For small portfolios which are studied during a relatively short period, such as the ASR pensions portfolio, the method by Plat (2009) would yield high levels of uncertainty. For me, this makes it unattractive to use this time series model. When the amount of data available is sufficiently large however, I would prefer the model by Plat (2009) over the Cox Proportional Hazard model, as the former does not require the assumption that the difference in portfolio and population mortality is constant over time.

As mentioned before I have based the regression on a subset of all age groups. Brouhns et al. (2002) and Denuit (2007) have done this as well, but due to the scope of those papers they had no need to extend the relation to other age groups. I do have to consider ages that fall outside the subset of ages on which the regression is based, which means that I do need to extend the relation. I postulate that the relation found holds not only for the considered subset, but also for the other ages. Again, due to the relatively small data set (especially at very young and very high ages), I cannot test whether this is a valid assumption.

4.3.2 Combining the models

In this subsection I will discuss how I have combined the models for population longevity and mortality of people in the ASR pensions portfolio. In the literature there are only few articles available on this subject. To my knowledge, Plat (2009) is the only one who has succeeded in combining the models whilst not neglecting the stochastic characteristics of one of the parts. His approach is discussed in section 4.1. Gschlössl et al. (2011) have developed an approach based on Poisson regression analysis, but this requires the use of more covariates (such as smoking habits and the occurrence of hereditary diseases in the family). This method can be used in the business of individual life insurance, but not for pension insurance where there are typically no personal characteristics of the insureds known.

The Cox Proportional Hazard model, which is the one I have used, provides point estimates and confidence intervals for the coefficients θ_1 and θ_2, together with correlation between the coefficients. In this way I can derive the full distribution of the regression line, under the assumption that the errors in the regression equation follow a Normal distribution. In the setting of the combined model, future death rates $m_{x,t}$ depend on a_x, b_x and κ_t from the Lee-Carter model and on θ_1 and θ_2 from the Cox Proportional Hazard model. These parameters are all stochastic with a known distribution, but there is no information available on the correlation between the Lee-Carter parameters and the Cox Proportional Hazard parameters. I can thus unfortunately not come to a fully equipped stochastic model for future mortality in the portfolio. Brouhns et al. (2002) therefore propose to only implement the point estimates of the parameters for the portfolio mortality model into the stochastic population mortality model. I will adopt this method. It is important to note that in this way the only randomness considered in the total mortality model

comes from the Lee-Carter model. The relation between population and portfolio mortality is now seen as deterministic, which leads to believe that in this way total uncertainty is (slightly) underestimated. Results for this part are to be found in section 4.4.2.

4.4 Results

Section 4.3 presents the model which I have used for the mortality as experienced by ASR in its pensions portfolio and how to combine this with the model for population mortality. In this section I will present the results of this approach.

4.4.1 Portfolio mortality model

I have run the Cox Proportional Hazard model for three different periods. One time for the period 2002-2005, once for the period 2006-2009 and I have done a run for the whole period 2002-2009. In Appendix B, figure B.1 shows scatter plots for the logarithm of population central death rates versus the logarithm of portfolio central death rates are shown. These plots clearly suggest the existence of a log-linear relationship as in the model by Brouhns et al. (2002) and Denuit (2007).

For each gender and period I have chosen a different set of age groups. The starting age is the first age for which the insured amount belonging to the deceased in the portfolio is at least 25,000 euro for the gender and period at hand. The closing age is the minimum of the last age for which the insured amount belonging to the deceased in the portfolio is at least 25,000 euro and 98, as this is the age after which smoothing of the population mortality by NIDI starts. Note that the regression is thus based on a subset of all ages, but I postulate that the obtained relation holds throughout the whole age range 0-98. For the later ages 99-120 I will prolong the table by means of the procedure by Denuit and Goderniaux (2005), which is explained in section 3.3. In this way I ensure that death probabilities are increasing until age 120 and then reach value 1, so that no one will survive after age 120. Tables 4.1 and 4.2 present the results of the OLS regression. Per period, the first row indicates the estimate for θ_1 in equation (4.1). The second row indicates the estimate for θ_2. The column 'Mean' gives the point estimate, the two columns following present the 95% confidence interval which are based on the standard error of the regression. In the last column, the coefficient of determination R^2 is shown.

Period	**Age groups**	**Mean**	**Lower Bound**	**Upper Bound**	R^2
2002-2009	42-96	-0.2702	-0.3878	-0.1526	98.90%
		0.9904	0.9616	1.0191	
2002-2005	42-96	-0.1674	-0.3483	0.0135	97.20%
		0.9621	0.9171	1.0071	
2006-2009	42-95	-0.1670	-0.3262	-0.0076	98.39%
		1.0633	1.0253	1.1012	

Table 4.1: *Results regression Cox Proportional Hazard model: men.*

Figure B.2 shows graphs of the different fits. As one would expect, the fit for the total period 2002-2009 lies between the others. It should be noted that for some regressions, the intercept is not significant at the 95% confidence level. It can be seen that the confidence intervals for the periods 2002-2005 and 2006-2009 are considerably wider than the confidence intervals for the period 2002-2009, just as the coefficients of determination R^2 are lower. This is due to the sample being smaller, which causes more variation. Furthermore, the estimates vary quite much between the different periods, some are even significantly different at the 95% confidence level. This seems to be due to the lower ages, as figure B.1 suggests that the relations are closely together at the higher ages (for which more data is available). This significant difference is in contradiction with

Period	Age groups	Mean	Lower Bound	Upper Bound	R^2
2002-2009	44-98	-0.2828	-0.4528	-0.1129	97.74%
		0.9537	0.9138	0.9937	
2002-2005	46-97	-0.1974	-0.4450	-0.0503	95.06%
		0.9243	0.8645	0.9842	
2006-2009	55-98	-0.2888	-0.6226	0.0045	92.43%
		1.0062	0.9166	1.0959	

Table 4.2: *Results regression Cox Proportional Hazard model: women.*

the assumption that the relation between population and portfolio mortality is constant over time. However, as the data only allowed me to consider two different periods, I can still not come up with a time-dependent process for the portfolio mortality. Therefore I will keep the assumption that the relation between population and portfolio mortality does not change over time.

The resulting fits can also be translated into factors as modeled by Workgroup PLT (2010), equation (4.13). I have done this for the implied factors from the fits of the Cox Proportional Hazard model for the period 2002-2009. The graphs then obtained are shown in figure B.3, where for comparison also the factors found by PLT are displayed, just as the factors directly observed from the data. The graph for men is slightly decreasing for the lower ages of the interval and slightly increasing for the higher ages. This is different from what Workgroup PLT (2010) found for the market, as they discovered a strictly increasing relation. For women I find over the whole age range a decreasing pattern. This means that the differences in mortality between women in the population and in the portfolio get in fact more apparent as women get older. This contradicts with the findings of Workgroup PLT (2010), who found a more or less increasing pattern. This is why I have also performed regressions based on the approach by Workgroup PLT (2010) on the data of the ASR portfolio. This means that I have modeled the relation

$$e_x^m = \varphi_1^m + \varphi_2^m x + \gamma_x^m \tag{4.15}$$

for men and the relation

$$e_x^f = \varphi_1^f + \varphi_2^f x + \varphi_3^f x^2 + \gamma_x^f \tag{4.16}$$

for women, with e_x^m and e_x^f defined as in equation (4.13) and with γ_x^m and γ_x^f iid normally distributed error terms with mean 0 and respective variances $\sigma_{x,m}^2$ and $\sigma_{x,f}^2$. I have run these regressions for the ages 42-94 (men) and 44-94 (women). These are the overlaps of the intervals I have determined using the criterion based on insured amounts, and the intervals Workgroup PLT (2010) considered.

Parameter	Mean	Lower Bound	Upper Bound	R^2
$\widehat{\varphi}_1^m$	0.9031	0.7180	1.0882	96.88%
$\widehat{\varphi}_2^m$	-0.0014	-0.0041	0.0012	

Table 4.3: *Results regression following method by PLT: men.*

Parameter	Mean	Lower Bound	Upper Bound	R^2
$\widehat{\varphi}_1^f$	2.0342	0.5212	3.5472	95.04%
$\widehat{\varphi}_2^f$	-0.0279	-0.0731	0.0172	
$\widehat{\varphi}_3^f$	0.0002	-0.0002	0.0005	

Table 4.4: *Results regression following method by PLT: women.*

Table 4.3 shows the results obtained for men, table 4.4 shows the results for women. The bounds

indicated for the parameter values are the bounds of the 95% confidence intervals. Workgroup PLT (2010) found for their dataset the values $\widetilde{\varphi}_1^m = 0.1392$ and $\widetilde{\varphi}_2^m = 0.0087$. For women they found $\widetilde{\varphi}_1^f = 0.1022$, $\widetilde{\varphi}_2^f = 0.0201$ and $\widetilde{\varphi}_3^f = -0.0001$. Clearly these values are very different and they do not fall in the confidence intervals displayed in tables 4.3 and 4.4. I must therefore, based on my data, on the 95% confidence level strongly reject the hypothesis that the factors found by PLT hold for the ASR portfolio considered here.

That there are differences between mortality patterns at ASR and mortality patterns at other insurers is thus clear, but I do not have a sound explanation for why these differences are present. Several options are possible. It could for instance be due to errors which might still be present in the data on mortality at ASR (see the subsection in this Chapter on the data used). The decreasing pattern of the portfolio mortality reduction factors leaves this possibility open, as potential errors would especially lead to lower factors at higher ages, where exposure is lower. Another explanation could be that the composition of ASR's portfolio is fundamentally different from that of other insurers. At lower ages, people in the ASR portfolio would then be less healthy than people with a pension insurance at other companies, whereas at higher ages the people in the ASR portfolio would be healthier. If that is indeed the case, this suggests that in the past ASR had many companies with 'white collar' workers in the portfolio, whereas now they would have many companies with 'blue collar' workers. The terms 'blue collar' and 'white collar' refer to lower educated and higher educated people respectively. Empirical investigation of the contract history does however not suggest a significant change in terms of the types of companies in the portfolio. Of course there is also the possibility that the discrepancy is due to the relatively small sample considered, meaning that the differences are merely coincidental. It should furthermore be noted that the observed differences are stronger for women than for men.

Even though the portfolio mortality factors derived differ from those found by Workgroup PLT (2010), the empirical life table derived from the mortality data for ASR does not seem to deviate that much from the population life table over the same time period, multiplied by the portfolio mortality factors from Workgroup PLT (2010). Figure B.4 shows this in a graph.

4.4.2 Combined model

Now that I have developed a model for the future population mortality in the Netherlands and a model for the portfolio mortality relative to the mortality of the Dutch population, I can combine those to predict the future mortality in ASR's pensions portfolio. As already stated in subsection 4.3.2, I will do this by implementing a point estimate of the model for portfolio mortality into the stochastic model for population mortality. As I have three point estimates, based on observation periods 2002-2005, 2006-2009 and 2002-2009, I first need to decide which point estimate to use. Also I should decide what to do with the ages that are not included in the sample on which the estimates are based.

Concerning the first problem, I have decided to use the point estimates that result from evaluating the entire sample period. As sample size is an important issue in this research, I do not want to 'throw away' parts of my data that can be very useful. Furthermore, taking the whole sample period allows me to consider the effects of adverse selection for a broader age range (see previous section). By using the whole sample period I implicitly assume that each year in the past is equally representative for the future. This is in line with assuming that the effects of adverse selection are constant over time. Of course this assumption might prove to be not valid, which is suggested by the results of the regression. The size of my data set is however not such that I can model how the relation varies over time. Therefore it is very important for ASR to keep adding new years of data and to improve the quality of administration. In this way ASR can in the future decide to either neglect early years of the sample, stating that effects of adverse selection have changed over time, or possibly to adopt a model which leaves more room for stochastic features. Think for instance of the model by Plat (2009).

By implementing the coefficients of the Cox Proportional Hazard model into the Lee-Carter model, I can find period life expectancies for people born in 2060. For men that would be in ASR's pensions portfolio I find a period life expectancy at birth in 2060 of 87.21 years. This is 1.70 years higher than the corresponding number for a new-born male in the general Dutch population according to my population mortality model. For new-born females in the ASR pensions portfolio, the period life expectancy in 2060 is 88.07 years. This is 0.82 years higher than the same number for a new-born female in the general Dutch population as obtained via my model for population mortality. The 95% confidence interval for males is $[85.00; 89.99]$, for females it is $[85.28; 90.66]$. The estimate by Workgroup PLT (2010), including their factors for portfolio mortality, is for both men (87.93) and women (88.90)[5] slightly higher than the values I obtain. Both however fall well into the confidence intervals. It should be noted that the implementation of their factors has a slightly lower absolute impact on life expectancies for their table than the implementation of the Cox Proportional Hazard model has for my mortality model.

I can also calculate period life expectancies for people in ASR's pensions portfolio aged 65 in 2060. For men, this number is 23.42 years. This is 1.49 years higher than for people in the general Dutch population following from my population mortality model. For women the number is 25.06 years, which is 0.90 years higher than for people in the general Dutch population according to my population mortality model. The 95% confidence interval for males is $[21.69; 25.84]$, for females it is $[22.95; 27.15]$. As for the life expectancies at birth, the values obtained by Workgroup PLT (2010) when they include their portfolio mortality factors are slightly higher for men (23.77) and for women (25.39)[6], but both still fall well into the confidence intervals.

To get a grasp of the influence of neglecting the uncertainty in the Cox Proportional Hazard model, I have also calculated life expectancies where not the point estimates but the bounds of the 95% confidence intervals of the regression are implemented into the Lee-Carter model. Testing for the underlying Normal distribution of error terms, using D'Agostino's K-squared test, shows that I cannot reject the null hypothesis of normality on the 95% confidence level. Implementing the resulting lines into the best estimate of the Lee-Carter model gives bounds on the life expectancies at birth which are about a year apart for men and about one and a half year apart for women. For remaining life expectancies at age 65 the differences are about 20% smaller. Uncertainty in the relation between portfolio mortality and population mortality thus seems to have a considerable effect on the life expectancies. One should however not draw conclusions here, since no information on correlations is known.

Concluding one can see that the implementation of the model for portfolio mortality leads to a similar decrease in mortality for men and women. My model predicts slightly lower life expectancies (higher mortality) for men than the model by Workgroup PLT (2010). For women it holds the other way around. The difference is, at the level of life expectancies in 2060, however not statistically significant. In other words, using my model I can on the 95% confidence level not reject the null hypothesis that the (remaining) life expectancies in 2060 found by Workgroup PLT (2010) are the right ones for the pensions portfolio of ASR that I have considered here.

[5] Source: own calculations.
[6] Source: own calculations.

Chapter 5

Longevity risk in Solvency II

In the upcoming financial regulation framework Solvency II for insurers active in member states of the European Union, regulation is put forward for the amount of capital insurers should reserve to cover for their risks. For each branch of the insurance business (life, non-life, health etc.), different risks are distinguished for which companies should hold extra reserves. The total amount of money an insurer should reserve is called the basic Solvency Capital Requirement (SCR). This basic SCR is comprised of the Solvency Capital Requirements for individual lines of business, where diversifications are taken into account by means of a correlation matrix. Note that I will not discuss the basic SCR further. For an insurer issuing lifelong (pension) annuities, longevity risk is an important component of the totals risk to be assessed. In this Chapter I will go deeper into the SCR for longevity risk. First of all I will present the analytics on which the Solvency II standards are based. In section 5.2 I discuss two approaches for calculating the Solvency Capital Requirement. The first is the standard approach as proposed in QIS5, a Quantitative Impact Study regarding the calibration of Solvency II. The second approach concerns a proposal for an internal model, which should give a more accurate approximation of the SCR for longevity risk than the standard approach. I will apply both methods to two model portfolios resembling the pensions portfolio of ASR and to four modified portfolios. In section 5.3 the results obtained for both approaches are presented and compared. Note that throughout this section I will use the time subscript t as starting time for the calculations. Future time periods are denoted by time $t+s, s = 1, 2, \cdots$.

5.1 Analytics

In this section I will present the analytics behind the determination the Solvency Capital Requirement for longevity risk. The reasoning behind Solvency II is that insurers should with 99.5% certainty be able to cover for their liabilities over a one-year time horizon. In other words, the insurer should calculate the 99.5% Value-at-Risk over this one-year horizon. In yet other words, the probability at time t that the funding ratio FR_{t+1} falls below one should be lower than 0.5%, where FR_{t+1} is determined according to

$$FR_{t+1} = \frac{A_{t+1}}{L_{t+1}} \tag{5.1}$$

and where A_{t+1} denotes the value of the assets at time $t+1$ and L_{t+1} denotes the value of the liabilities also at time $t+1$ (to be defined later in this section). This means that the minimum level of assets A_t^* that the insurer should hold at time t is defined by

$$A_t^* \equiv \min\left\{ A_t |\, \mathrm{P}_t\left(\frac{A_{t+1}}{L_{t+1}} < 1\right) \leq 0.005 \right\}, \tag{5.2}$$

where P_t denotes the objective probability measure at time t.

The capital the insurer should hold at time t, in excess of the value of the liabilities, is called the *Solvency Capital Requirement* (SCR) and is defined by

$$SCR_t \equiv A_t^* - L_t. \tag{5.3}$$

If I denote the liabilities to be paid at the end of year t by $\check{L}_t$ and assume a return r_t on the assets during year t, A_{t+1} can be written as

$$A_{t+1} = A_t^* \cdot (1 + r_t) - \check{L}_t. \tag{5.4}$$

Combining this equation with (5.2) and (5.3) gives

$$SCR_t = \mathrm{Q}_{0.995,\mathrm{t}} \left(\frac{\check{L}_t + L_{t+1}}{1 + r_t} \right) - L_t, \tag{5.5}$$

where $\mathrm{Q}_{0.995,t}$ denotes the function returning the 99.5% quantile of the distribution of $\frac{\check{L}_t + L_{t+1}}{1+r_t}$, conditional on the information available at time t.

The Solvency Capital Requirement is thus the extra amount on top of the expected present value of the liabilities, that the insurer should hold to be approximately 99.5% sure that he can in fact cover for his liabilities. The SCR depends on how the portfolio develops, in terms of inflow (new participants) and outflow (mortality and surrender), and thus changes every year. If one wants to determine the SCR at time t, one needs to have an expression for the value of the liabilities L_t. In Solvency II it is proposed to decompose the liabilities into a best estimate (BEL_t) plus a risk margin (RM_t):

$$L_t = BEL_t + RM_t, \tag{5.6}$$

with BEL_t as the expected present value of all future payments to be made

$$BEL_t = \sum_{s \geq 0} \mathrm{E}_t \left[\check{L}_{t+s} \right] \cdot P_t^{(s+1)}, \tag{5.7}$$

where $P_t^{(s+1)}$ is the discount factor (or: price at time t of a zero coupon bond maturing at time $t + s + 1$). Note that I need to take $P_t^{(s+1)}$ to discount $\check{L}_{t+s}$ and not $P_t^{(s)}$, as the payments are due ultimo year t.

The risk margin is calculated according to the so-called cost of capital approach, which is based on the idea of selling the liabilities to another insurance undertaking with an empty portfolio (European Commission, 2010). This (fictional) counterparty would also need assets to cover for these liabilities, as well as a risk premium. This risk premium is based on all values of $SCR_{t+s}, s \geq 0$. Furthermore, this number should be multiplied with a Cost of Capital rate (CoC), which should reflect the interest rate on top of the risk-free rate that the counterparty should pay to the market in order to attract the capital needed. In QIS5 this number is set at 6%. Many companies in the business have argued that this number should be lowered, but hitherto no justifying evidence has been delivered (CEIOPS, 2010). It should be noted that the insurer does not have to pay the risk-free rate to the counterparty, as the latter can invest the needed assets for a risk-free return itself (CEIOPS, 2009). This means that the risk margin at time t is determined by

$$RM_t = CoC \cdot \sum_{s \geq 0} SCR_{t+s} \cdot P_t^{(s+1)}. \tag{5.8}$$

It is quickly observed that when determining the SCR and the risk margin, one arrives in a loop. The SCR depends on risk margins to be held now and in the future, whereas in return the risk

margin is a function of all future SCRs. Furthermore, at time t one cannot determine the value of SCR_{t+s}, as this capital requirement depends on the run-off behavior of the portfolio at hand between time t and time $t+s$, which is subject to randomness. Assumptions need to be made in order to solve these problems. CEIOPS (2010) proposes to neglect the risk margin as part of the liabilities when the future SCRs are calculated. The underlying idea is that the risk margin only accounts for a small part of total liabilities, making this a valid approximation and solving the first problem. Using this assumption it is in principle possible to calculate the future SCRs and the risk margin, but it is still a cumbersome process. Since the portfolios considered typically span a long time horizon, one would either need to perform many simulations or to make closed form approximations of the distribution of the underlying variables. The latter approach is for instance used by Stevens et al. (2010). QIS5 (CEIOPS, 2010), which is the fifth Quantitative Impact Study on the calibration of Solvency II, proposes some simplifications in calculating the SCR, which should make it easier to calculate the SCRs. In subsection 5.2.2 I will discuss these simplifications.

5.2 Determining the SCR for longevity risk

Section 5.1 has provided a short introduction about the Solvency Capital Requirement for longevity risk in Solvency II. In this section I will discuss two approaches regarding the determination of this capital requirement in ASR's pensions portfolio. In the first subsection I will shortly discuss the model portfolios and the term structure for determining the discount factors that I have used. Subsection 5.2.2 presents the standard, simple approach as proposed in QIS5. This approach is based on deterministic mortality probabilities. I will also discuss two other simplifications put forward in QIS5. In the subsection thereafter I will present an alternative, more sophisticated model based on simulation techniques and which makes use of the stochastic characteristics of my mortality model.

5.2.1 Portfolios considered

The aim of this part of my thesis is not to come up with a fully equipped internal model for the SCR for longevity risk, but more to investigate whether it could be beneficial for ASR to adopt an internal model over the standard approach. In the literature, typically simple and closed portfolios consisting only of people aged 65 (only holding an old age pension or holding an old age and a partner pension) are considered. This is for instance the case in Stevens et al. (2010) and in Jansen (2012). These investigations suggest that for those portfolios it might indeed be beneficial to adopt an internal model. This conclusion might however be invalid when a different, for instance younger, portfolio is considered. As the ASR pensions portfolio does not consist of only people aged 65, it might be worthwhile to do the analysis based on a portfolio which shows more resemblance to ASR's. I have constructed two of those portfolios, one only containing old age pensions and one containing old age pensions and partner pensions. I will use my mortality model discussed in Chapters 3 and 4 to predict the future mortality rates of people in these portfolios. These mortality rates can then be used to evaluate the run-off of the portfolios. Section 5.3 presents the results of this approach. Later I will also shortly present the results for four modified portfolios. At that point I will also present how these portfolios are modified.

For this part I was provided with an extract of the total ASR pensions portfolio, containing number of insureds and insured amounts per age group and gender[1]. In order to decrease the computational load, I have decided not to evaluate the development of each policy separately, but to look at the totals per age group. For each age group and gender I have the insured amount that the people in the group in total are (expected to be) entitled to receive ultimo each year they are 65 or older and still alive. These insured amounts are based on what an insured is, based on his current salary, expected to receive during retirement. I assume career patterns to be flat for the

[1]Thanks go to Hans den Uijl for providing this extract.

participants with age below 65, meaning that the insured amounts will not change in the future. Furthermore I will not consider the possibility of indexation. The entitlements for each cohort will thus only change (decrease) due to mortality in the portfolio (note that I assume the portfolios to be closed to new participants). The first portfolio consists of males and females between the ages 19 and 109. The second portfolio consists of the same people, but with an additional entitlement. In case the insured dies but his or her partner is still alive, the partner receives each year a payment equal to 70% of the insured amount built up, until he or she is passed away. So these people do not only hold an old age pension, but also a partner pension. Note that for men I assume a female partner that is three years younger. Likewise I assume for women a male partner that is three years older. Note furthermore that the construction of these portfolios is such that not all participants will have died by the time the year 2060 is reached. This means that I will also need to use the Lee-Carter model to estimate death probabilities in the years thereafter. I will do this using the parameters coming from the Lee-Carter model that is based on all five starting years of estimation (1970, 1975, 1980, 1985 and 1990). I run this model for 105 years into the future. By this year, all participants and partners of participants will for sure have died, since in my mortality model people cannot become older than 120. The first year of the calculations will be 2012, as the portfolio considered resembles the situation in 2012 as well. Note that since the input data of my Lee-Carter model does not reach beyond 2009, I need to assume that development of mortality rates between 2009 and 2012 is as expected from the Lee-Carter model, i.e., following the best estimate.

For the discounting I use the nominal term structure of zero coupon bonds provided by the Dutch National Bank (DNB) as it was on June 30, 2012. This term structure gives discount rates $P_t^{(s+1)}$ based on zero coupon bonds with a duration smaller than 60 years. For the longer terms, I adopt the rate that holds for the 60-year term. Furthermore I assume that the rates remain constant over time, i.e., I neglect interest rate risk (interest rates are deterministic). So the price of a τ-year zero coupon bond in year t is the same as the price of a τ-year zero coupon bond in year $t+s, s>0$. I have printed the list of resulting discount factors in Appendix D.

5.2.2 Standard approach

Using the formulas presented in Chapter 5 it is possible to determine the SCR for a given portfolio of liabilities. The process is however still far from trivial, which is why some simplifications are proposed in QIS5 (European Commission, 2010). The first simplification regards the use of stochastic death probabilities. In the original formulation of Solvency II, the insurer is asked to calculate a 99,5% Value-at-Risk (European Union, 2009). Many insurers however use deterministic death probabilities, possibly with a future trend incorporated. In that case it is not possible to calculate quantiles, since no distribution is known. Therefore, QIS5 proposes a 20% immediate shock in all death probabilities (thus for each age x and for each year t) as scenario corresponding to the 99.5% quantile, whereas the best estimate is calculated using the deterministic probabilities. The SCR for year t is then in straightforward manner determined, by looking at the discounted values of the liabilities in year t for the best estimate scenario and the shock scenario. The SCR for year t is, according to the standard approach, equal to the difference of these two discounted values:

$$SCR_t = \sum_{\tau \geq 0} \left(\check{L}_{t+\tau}^{shock} - \check{L}_{t+\tau}^{BE} \right) \cdot P_t^{(\tau+1)}. \tag{5.9}$$

This approach can in principle also be used for the later years $t+s$, where the expectation of SCR_{t+s} can be calculated by

$$\mathrm{E}_t \left[SCR_{t+s} \right] \approx \sum_{\tau \geq 0} \mathrm{E}_t \left[\left(\check{L}_{t+s+\tau}^{shock} - \check{L}_{t+s+\tau}^{BE} \right) \cdot P_{t+s}^{(\tau+1)} \right], \tag{5.10}$$

where the superscript *shock* refers to the shock scenario and the superscript BE refers to the best estimate. The risk margin at time t can then be calculated following

$$RM_t \approx CoC \cdot \sum_{s \geq 0} \mathrm{E}_t \left[SCR_{t+s} \right] \cdot P_t^{(s+1)}. \tag{5.11}$$

Even though the basic structure for calculating SCR_{t+s} is now specified in equation (5.10), it still depends on the development of the portfolio between time t and time $t+s$ and the expectation to be calculated can thus not yet be determined. Therefore QIS5 allows for other simplifications to calculate future SCRs. The second proposal for simplification in QIS5 is to determine future values of the SCR according to a constant fraction of the best estimate. This can be written as

$$\frac{\mathrm{E}_t \left[SCR_{t+s} \right]}{SCR_t} \approx \frac{\mathrm{E}_t \left[BEL_{t+s} \right]}{BEL_t}, \tag{5.12}$$

with BEL_t as defined in the previous section. The underlying assumption here is that the SCR is proportional to the best estimate. This might however prove not to hold in practice, for instance when the underlying risks change over time. This is explained in QIS5 (CEIOPS, 2010). Börger (2010) argues that this method of approximating the future SCRs is highly inaccurate, as he finds the SCR as fraction of the best estimate value of the liabilities to be generally increasing over time for portfolios that are closed to new participants.

The third and last simplification proposal also regards the calculation of the SCR for future years. As explained in section 5.1, at time t it is not possible to determine the exact value of SCR_{t+s}. Therefore QIS5 allows to assume that between time t and time $t+s$ the run-off pattern of the portfolio is according to the best estimate. The SCR at time $t+s$ thus only depends on deviations from the best estimate after time $t+s$. If I denote the scenario where death probabilities between time t and time $t+s$ develop according to the best estimate by ξ_t^s, I can write this as

$$\mathrm{E}_t \left[SCR_{t+s} \right] \approx SCR_{t+s} \left(\xi_t^s \right), \tag{5.13}$$

where the right hand side denotes the realization of the random variable SCR_{t+s} in case development of the death probabilities between time t and time $t+s$ is as in the scenario ξ_t^s.

According to Börger (2010), this method of calculating the SCRs and consequently the risk margin is in practice generally seen as yielding the "exact" risk margin. Stevens et al. (2010) however find, using an analytical approximation, that this simplification might not always be accurate. It however still seems to be the more accurate than the previous proposal. I will hereafter refer to the standard approach as the method that combines the 20% longevity shock with the latter simplification for calculating future SCRs.

It is not that difficult to calculate the amount of money ASR should hold for the longevity risk in the pensions portfolio according to the standard approach. For this approach I use the best estimate death probabilities from my stochastic mortality model for the best estimate value of the liabilities. For the shock scenario I use the same death probabilities, but multiplied with a factor 0.8. In the way this method is built up, individual mortality risk is neglected. Individual mortality risk is the risk coming from the fact that, given a certain probability distribution of mortality, the remaining lifetime of a person is still unknown. He or she can thus still live longer or shorter than expected. This individual mortality risk can be pooled away by attracting a large amount of insureds. In that case, the observed number of deaths will converge to the expectation. In section 5.3 the SCR for the treated portfolio as calculated using the standard approach is given. But first, the next subsection deals with the internal model that I propose for calculating the SCR for longevity risk, which allows for considering the stochastic characteristics of my mortality model.

5.2.3 Internal model

Many insurers in Europe deem the shock scenario proposed in QIS5 to be too conservative (CEIOPS, 2010). Also it neglects the influence of age of insureds on the uncertainty of the present value of the liabilities. Therefore the general feeling is that using an internal model instead of the standard approach might yield a more accurate SCR that the insurer needs to hold for longevity risk. This is why I would like to investigate whether a (simple) internal model indeed suggests a substantially different SCR to be held.

In my mortality model (presented in Chapters 3 and 4) the future death probabilities are not deterministic, but stochastic. This means that there is no need to adopt the 20% decrease in mortality shock to approximate the 99.5% Value-at-Risk, since it can be determined explicitly. I will base my proposal for an internal model on my ASR specific mortality model.

For the internal model I will follow an approach closely related to the approach followed by Jansen (2012), but try to expand that method where possible. An important property of her approach is that not only longevity risk, but also the individual mortality risk can be modeled. This is done by simulating the number of deaths for age group x in year t using a Binomial distribution, analogous to what Cossette et al. (2007) do. I will however not adopt this step. For the simple portfolios that are normally considered, with only people aged 65 receiving a unity payment each year they are still alive, this method is not that computationally intensive. For the portfolio I consider however I could only model the individual mortality risk by separately modeling each individual lifespan. This would enormously inflate the time needed for computations, which is the main reason why I will not consider individual mortality risk here. Furthermore, the size of the portfolios (total number of insureds is 248,839) is such that the individual mortality risk will likely not have that much influence on the results, as this risk disappears for large portfolios. Note that for the closed-portfolio approach followed here, there will be some individual mortality risk present in later years, when less people are left in the portfolio. Since this will only happen relatively far in the future however, this does not concern large amounts of money (as the portfolio size has already decreased substantially) and furthermore these amounts should be heavily discounted in the present risk margin.

In the internal model the calculation of the best estimate of the liabilities is exactly the same as in the standard approach, as the best estimate death probabilities are then used. Calculation of the 99.5% quantile is however more involved. Starting from the year 2012, I run 5,000 simulations where I for every year $t = 2012, \ldots, 2116$ and age groups $x = 0, \ldots, 120$ simulate a $q_{x,t}$ according to the Lee-Carter model. These simulations are drawn from the parameter estimates as described in section 3.3, based on all five starting periods. The simulations thus reflect the longevity risk that is present, as these future $q_{x,t}$'s are random. Note though that the number of simulations is small compared to the number of simulations that is used to reflect the parameter risk and starting period risk in the Lee-Carter model. The amount $\pi_{x,t}$ that should be paid to cohort x at the end of year t, for the old age pensions portfolio, is then obtained via

$$\pi_{x,t} = \pi_{x-1,t-1} \cdot (1 - q_{x,t}). \tag{5.14}$$

Note that these amounts only need to be paid out for $x > 65$. For every year, $\check{L}_t$ is then defined according to

$$\check{L}_t = \sum_{x=65}^{120} \pi_{x,t}. \tag{5.15}$$

For the portfolio consisting also of partner pensions, extra payments come into play. For each cohort x and year t, the amount $\pi'_{x,t}$ to be paid to the partners of the insureds is defined by

$$\pi'_{x,t} = \pi'_{x-1,t-1} \cdot (1 - q'_{x,t}) + 0.7 \cdot \pi_{x-1,t-1} \cdot q_{x,t}, \tag{5.16}$$

where $q'_{x,t}$ denotes the death probability of the partner. Note that the partner pension thus only pays out ultimo each year that the insured has already died and the partner is still alive. As said before, men are assumed to have a female partner who is three years younger, women are assumed to have a male partner who is three years older. For every year, $\check{L}_t$ is then defined according to

$$\check{L}_t = \sum_{x=65}^{120} \pi_{x,t} + \sum_{x=0}^{120} \pi'_{x,t}. \tag{5.17}$$

For every simulation the run-off pattern of the portfolios can be followed, which yields for each year a number to be paid out to policy holders aged above 65. By appropriately discounting those numbers to be paid, a present value of the total liabilities is obtained. By taking the 99.5% quantile of the empirical distribution obtained for this present value and then subtracting the best estimate for the first year (2012), the SCR for this year is obtained. In formula:

$$SCR_{2012} = \mathrm{Q}_{0.995,2012}\left(L_{2012}\right) - L^{BE}_{2012}, \tag{5.18}$$

where L_{2012} is defined by

$$L_{2012} \approx \sum_{s\geq 0} \check{L}_{2012+s} \cdot P^{(s+1)} \tag{5.19}$$

and L^{BE}_{2012} is defined by

$$L^{BE}_{2012} \approx \sum_{s\geq 0} \mathrm{E}_{2012}\left[\check{L}_{2012+s}\right] \cdot P^{(s+1)}. \tag{5.20}$$

Note that I drop the subindex for the discount factor, as I assume that interest rates are deterministic and constant over time.

For calculating the future SCRs, I will adopt the third simplification put forward in the previous subsection. Using this simplification yields a significant reduction in simulations needed, while the comment by Börger (2010) suggests that the loss of accuracy is only small (note that Stevens et al., 2010, are more cautious but still find it to be the most exact simplification proposal). For each year $2012+s$, $s = 1, \ldots, 104$ I thus first assume that run-off of the portfolio between year 2012 and year $2012+s$ has occurred following the best estimate, i.e., following the best estimate death probabilities. I denote this scenario by ξ^s_{2012}. Starting from year $2012+s$ I then run 5,000 simulations as described in the previous paragraph. This means that I run in total $105 \cdot 5,000$ simulations. Like before, the 99.5% quantile can be determined from the empirical distribution of the discounted value in year $2012+s$ of the liabilities to be paid. The SCR for year $2012+s$ is then obtained by subtracting the best estimate of the discounted value of liabilities for year $2012+s$ from this quantile. In formula:

$$\mathrm{E}_{2012}\left[SCR_{2012+s}\right] \approx \mathrm{Q}_{0.995,2012}\left(L_{2012+s}\right) - L^{BE}_{2012+s}, \tag{5.21}$$

where L_{2012+s} is defined by

$$L_{2012+s} \approx \sum_{\tau\geq 0} \check{L}_{2012+s+\tau}\left(\xi^s_{2012}\right) \cdot P^{(\tau+1)} \tag{5.22}$$

and L^{BE}_{2012+s} is defined by

$$L^{BE}_{2012+s} \approx \sum_{\tau\geq 0} \mathrm{E}_{2012}\left[\check{L}_{2012+s+\tau} | \xi^s_{2012}\right] \cdot P^{(\tau+1)}, \tag{5.23}$$

where $\check{L}_{2012+s+\tau}\left(\xi^s_{2012}\right)$ denotes the realization of $\check{L}_{2012+s+\tau}$ in case development between year 2012 and year $2012+s$ was according to scenario ξ^s_{2012}. Once these calculations are performed for all s, the risk margin for the first year can be determined following

$$RM_{2012} \approx CoC \cdot \sum_{s\geq 0} \mathrm{E}_{2012}\left[SCR_{t+s}\right] \cdot P^{(s+1)}. \tag{5.24}$$

This simulation-based approximation for calculating the Solvency Capital Requirement should be more accurate than the standard approach proposed in QIS5, as it specifically calculates the 99.5% quantile of the liabilities. The deterministic shock approach of the standard model imposes that the quantile of the present value of liabilities to be paid is independent of the age-wise composition of the portfolio considered. It is however clear that a 'green' portfolio, where the average age of participants is lower, suffers from substantially larger longevity risk than an old portfolio. A person of age 30 has more years left in which his death probabilities can vary from today's standards than a person of age 70. This is why, if one assumes that the shock of 20% is appropriate for a portfolio at some average age, the shock is probably too conservative for older portfolios and too optimistic for younger portfolios. For an old age pension portfolio with people aged 65, the determinstic shock seems to be too conservative (see for instance Stevens et al., 2010, Börger, 2010 and Jansen, 2012). In section 5.3 I will present the results of both approaches for the model portfolios that I have considered and compare these results.

5.3 Results

In this section I will present the results found for both the standard approach and the internal model. I will do this for the portfolio consisting of only old age pensions and for the portfolio where also partner pensions are insured, starting with the former. First of all I have obtained the results for the standard approach. In the first year (2012), the best estimate present value of the liabilities in this approach is slightly over 9 billion euro. The SCR for this year is little more than 616 million euro, which is 6.81% of the best estimate present value of the liabilities. In Appendix C, figure C.1 shows the development of the best estimate present value of the liabilities over the years. As the average age of people in the portfolio is well below 65, the present value of the portfolio is increasing in the first years, as the bulk of the payments is coming nearer and hence discounted less heavily. In figure C.2, the Solvency Capital Requirement as a percentage of the best estimate in this portfolio in the first seventy years is shown. Note that the other line in this graph is discussed in the next paragraph. The fraction considered basically follows a pattern which is reversed compared to the pattern of the best estimate, as the fraction decreases in the first years, but then rapidly increases. The rapid increase at the end comes from the high ages having more influence on the results as the portfolio runs off. For the very high ages, death probabilities in the best estimate scenario are close to 1. In the shock scenario however, these probabilities are close to 0.8. This means that in the shock scenario, those people have a probability of surviving which is almost 20 times as high as in the best estimate scenario. This is why the SCR as fraction of the best estimate is much higher for years where the very old dominate the portfolio. Note that this has virtually no influence in absolute terms, as in the later years there are only a few euros to be paid. Calculating the risk margin for the first year according to equation (5.11) yields a value of about 940 million euro. This is 10.38% of the best estimate of the present value of the liabilities for year 2012.

For the internal model, the best estimate present value of the liabilities is for each year the same as in the standard approach. In the first year the SCR is again approximately 616 million euro, which is 6.81% of the best estimate. Figure C.2 shows the development of the SCR as fraction of the best estimate over the first seventy years. This fraction follows roughly the same pattern as for the standard approach, but is on average slightly lower. The increasing pattern is not what I would expect, as for later years it seems more reasonable that uncertainty decreases. This because the people in the portfolio are older and have less years to live, meaning less years in which their experienced mortality rates can deviate from the best estimate. My suspicion is that this comes from the implementation of the parameter risk, which means that over every simulated path the drift in κ_t is different. The small number of simulations undertaken relative to the number of parameter estimates I have, makes the model more sensative towards the selection of a particular set of parameter estimates. If I neglect parameter uncertainty, i.e., I assume there exists only one set of parameter estimates, the SCR as fraction of the best estimate is decreasing over time. This

is what I would expect. The SCR being on average slightly smaller than the SCR in the standard approach, yields a slightly lower risk margin as well. This risk margin is for the first year in this approach about 920 million euro. This is 10.15% of the best estimate for that year, which is 0.2 percentage point smaller (2.27%) than in the standard approach.

For the portfolio where old age pensions are combined with partner pensions, results are slightly different. Of course the best estimate present value of the liabilities is higher than for the portfolio with only old age pensions considered, as the entitlement to a partner pension leads for every policy holder to either the same amount of payments as when only an old age pension is insured (in case the partner dies before the insured) or to more payments (in case the insured dies when the partner is still alive). Figure C.3 shows the development of the best estimate present value of the liabilities throughout the years. In the first year, the value of the best estimate is slightly under 11 billion euro. The value of the SCR for this year is almost 548 million euro, which is 5.01% of the best estimate. Figure C.4 shows the path of the SCR as fraction of the best estimate over the first seventy years (the other line in this figure is discussed in the next paragraph). This fraction is now rising over the whole period. For the first years, the fraction is lower than for the portfolio with only old age pensions. This is because in the shock scenario, the value of the partner pensions is for young people lower than in the best estimate scenario, since it is less likely that the insured person will die. The high levels of the SCR relative to the best estimate for later years occur for the same reason as before, namely because of the large differences in survival probabilities between the best estimate scenario and shock scenario. The risk margin in the first years is about 1.04 billion euro, which is 9.52% of the best estimate.

For the internal model, the best estimate does of course not change. The SCR does change and like for the pure old age pensions portfolio, it becomes lower than the SCR obtained via the standard approach. In the first year, the SCR is now approximately 514 million euro, which is 4.70% of the best estimate present value of the liabilities. This is 8.51% lower than in the standard approach. Figure C.4 shows the path of the SCR as fraction of the best estimate over the first seventy years. As can be seen, this path shows roughly the same pattern as in the standard approach, but on a lower level. The risk margin for the first year is now almost 850 million euro, which is 7.77% of the best estimate present value of the liabilities. Note that this is 22.52% smaller than the risk margin in the standard approach. This indicates that the results for the standard approach are less close to those for the internal model when also partner pensions are included.

Besides the model portfolio, I have (as mentioned before) also considered a few other portfolios. These other portfolios are equal to the model portfolio, but with shifted ages. So have I for instance calculated the SCRs for a portfolio where all insureds are 10 year older than they truly are. This makes the portfolio thus older and would thus in expectation lead to a lower SCR. Furthermore, the SCR coming from the internal model should be lower than that from the standard approach and the difference should be bigger than before. I have also performed the same calculations for a portfolio where all insureds are 10 year younger, which should lead to opposite results. This is a meant as a small analysis of sensitivity of the results towards the average age in the portfolio. The results for the modified portfolios are, along with the results for the actual model portfolio, displayed in table 5.1.

The first thing that comes to attention is that the SCR for the first year as fraction of the best estimate for that year is increasing with average age in the standard approach, whereas it is decreasing with age in the internal model. It seems that in the standard approach the uncertainty for older people is much larger than for younger people, due to the larger differences in survival probabilities discussed earlier. In the younger portfolios the SCR for the first year is thus lower than in the older portfolios. This contradicts with what I would expect. In the internal model however, differences in survival probabilities for the old are not that big, which makes that the uncertainty is mainly determined by the length of the horizon over which the people in the portfolio can find improvements in mortality rates. The amount of uncertainty is thus decreasing in the

Portfolio	$\frac{SCR^S_{2012}}{BEL_{2012}}$	$\frac{SCR^I_{2012}}{BEL_{2012}}$	$\frac{RM^S_{2012}}{BEL_{2012}}$	$\frac{RM^I_{2012}}{BEL_{2012}}$
OAP (-10)	6.28%	7.47%	12.35%	14.08%
OAP	6.81%	6.81%	10.38%	10.15%
OAP (+10)	7.42%	5.78%	8.62%	6.93%
OAP + PP (-10)	4.42%	5.31%	10.81%	10.27%
OAP + PP	5.01%	4.70%	9.52%	7.77%
OAP + PP (+10)	5.77%	4.55%	8.35%	5.88%

Table 5.1: *Calculated Solvency Capital Requirements and risk margins, as fraction of best estimate value of the liabilities. OAP denotes portfolio with only old age pensions, OAP + PP denotes portfolio with also partner pensions. Superscript S indicates use of the standard approach, superscript I indicates use of the internal model. Suffix (+10) indicates the portfolio where all insureds are 10 year older than in reality, (-10) points to the portfolio where all insureds are 10 year younger.*

average age of the portfolio. To me this seems more reasonable. It can be seen that results for the SCR in the first year are closest together for the portfolio resembling ASR's pensions portfolio. For the younger portfolio the SCR found via the internal model is higher, for the older portfolio the SCR found via the standard approach is higher. This implies, as expected, that the standard approach yields a value which is too high in case the portfolio is old, whereas the resulting SCR is not high enough for young portfolios. The risk margin in the first year as fraction of the best estimate for that year is for both approaches decreasing as average age increases. This is in line with what I would expect, since a younger portfolio should yield more uncertainty in future death probabilities and thus a higher reserve to be kept. The risk margin is basically a sum of discounted future SCRs. For younger portfolios there are more future SCRs than for older portfolios, as the run-off horizon is longer. Therefore it is logical that the risk margin is higher for younger portfolios. Compared to the (+10)-portfolios, there is 20 years of extra uncertainty present in the (-10)-portfolios. Furthermore it seems that the SCR and risk margin from the standard approach are higher compared to those from the internal model in case partner pensions are considered as well.

Concluding it seems like the standard approach for calculating the Solvency Capital Requirement and risk margin for longevity risk is quite accurate for portfolios that more or less resemble ASR's pensions portfolio, compared to the results found for my internal model. For a pure old age pension portfolio with insured amounts as built up in ASR's pensions portfolio, the SCR and risk margin are only a bit smaller when a more accurate approach (the internal model) is used. This is more or less in line with what Börger (2010) found for portfolios consisting of deferred annuities. The difference is bigger when a portfolio consisting of old age pensions and partner pensions is considered, but still not that large. Therefore, even though the results suggest that ASR can do better in determining their SCR for longevity risk, it is doubtful whether ASR should indeed adopt some form of an internal model for determining the SCR and risk margin for longevity risk. Implementing an internal model will most definitely yield different amounts of capital to be held, but I can at this point not say whether the benefits outweigh the costs. Assuming that the portfolios of other insurers are similar to ASR's, this might hold for those other insurers as well. It seems furthermore more attractive to adopt the internal model when a portfolio consisting of both old age pensions and partner pensions is considered, since differences in results between standard approach and internal model are bigger for these portfolios. It should however be noted that the results presented here are only based on a small number of simulations, compared to the number

of simulations that was used to construct the Lee-Carter model.

Even though I can based on this research not say that the standard approach yields an inaccurate approximation for the SCR of ASR's pensions portfolio, other papers (mentioned before) have shown that the standard approach for determining the SCR might still be too conservative for older portfolios. Also, Towers Perrin (2009) have performed research towards an age-dependent expected improvement in future mortality rates. When whole life annuities are considered, they found only for ages below 40 a 99.5% quantile of decrease in mortality which approximately matches the shock proposed in QIS5. For the higher ages, the improvement to be used for the 99.5% VaR suggested by their approach is far below the percentage set in QIS5.

In my opinion, another shortcoming of the standard approach is that the SCR seems to be increasing in the average age of the portfolio. This effect can be explained in the model, see before, but it does not seem reasonable to me. I would expect that a younger portfolio would require a higher SCR than an older portfolio, as there are simply more years left during which death rates can vary. This effect is only partially present within my internal model. For younger portfolios, the required SCR as percentage of the best estimate is higher than for older portfolios, which is in line with my expectations. During the run-off of a portfolio however, the SCR as fraction of the best estimate increases over time, thus when people in the portfolio get older. This still seems undesirable and unreasonable to me. My suspicion is that this is due to the relatively small number of simulations that is taken compared to the number of simulations for parameter uncertainty used to constitute the Lee-Carter model.

As a sidestep, I take this paragraph to question how the 20% shock is determinated. In my interpretation, it is in QIS5 neglected that in lifetables used in practice, most of the times a decreasing trend in mortality is already incorporated. I base this on a passage in CEIOPS (2010). There it is argued that the shock is set at 25% (note that this was later altered to 20%), because research among nine European countries has shown that death probabilities had in the prior fifteen years decreased by about 25%. If this is indeed how the 20% level was established, the apparent reasonable approximation of the 99.5% quantile for a representative pensions portfolio is merely a lucky shot.

Chapter 6

Conclusions and recommendations

The goal of this thesis was to develop a stochastic model for future mortality in the ASR pensions portfolio, and to investigate how this model could be used for calculating the Solvency Capital Requirement for longevity risk. As the data set of mortality in the portfolio was not large enough to come up with a stochastic model solely based on the data from ASR, I had to take a different approach. First of all I have developed a stochastic model for the future mortality in the Dutch population, using the Lee-Carter method. Then I have separately modeled the past mortality in the portfolio relative to mortality in the general Dutch population over the same period. The point estimates of this relation I have implemented into the Lee-Carter model, which provided a stochastic model for future mortality in the pensions portfolio of ASR.

Resulting life expectancies in 2060 for people in the ASR portfolio do not differ that much from what Workgroup PLT (2010) found based on data of other pension insurers in the Netherlands, accounting for about eighty percent of the market. In fact I can, based on my model, on the 95% confidence level not reject the null hypothesis that the life expectancies found by Workgroup PLT (2010) are the correct ones for the ASR pensions portfolio considered. However, the portfolio mortality factors found by PLT are on the 95% confidence level significantly different from the portfolio mortality factors for ASR. There are multiple possible reasons for why the mortality pattern in the ASR portfolio is different from that pattern for the insurers considered by Workgroup PLT (2010). These are discussed in section 4.4.

In the ideal situation I would have modeled population mortality and portfolio mortality simultaneously. To make this possible, a larger data set is needed. The approach I have used now slightly underestimates uncertainty, as I am forced to neglect uncertainty surrounding a linear regression. Furthermore I need to assume that the relation between population mortality and portfolio mortality remains constant over time. I therefore recommend ASR to continue the gathering of data on mortality in the portfolio. In that way it might in the future be possible to investigate how the relation between portfolio mortality and population mortality has changed over time and to incorporate this into the portfolio mortality model. At this point in time, no conclusion can be drawn in this respect, as the data set is simply too small. The results of the regressions on the relation between population and portfolio mortality do however suggest that it is not constant over time. Experts do not agree whether the general relation between mortality of a population and a subpopulation has changed or remained constant in the past. It should be noted that the administrative system from which the data on portfolio mortality was extracted has in the past proven to suffer from a lag in administrating deaths (and in some cases even failed to register them at all). This might still be the case today, which makes the data less reliable. A further investigation of the quality of the data was not in the scope of this thesis, as this meant that for thousands of policies one would need to find out whether the insured is still alive.

The second part of my thesis was devoted to the calculation of the Solvency Capital Require-

ment that ASR should hold for longevity risk according to the Solvency II regulation framework for insurers. I have calculated this reserve for six different model portfolios, of which one resembles the ASR pensions portfolio, and using two different approaches. The first method was the standard approach as proposed in QIS5, the second method was a proposal for an internal model based on not only the best estimate death probabilities from my mortality model, but also using the stochastic characteristics. Three of the portfolios considered consist of people entitled to an old age pension, the other three consist of people also entitled to a partner pension.

Results for this part showed that the standard approach proposed in the Quantitative Impact Study QIS5 seems to perform quite well in estimating the SCR for longevity risk. For the resembling portfolio consisting only of old age pensions, the SCR calculated using the internal model is about the same as the SCR obtained via the standard approach. When people in the portfolio are made 10 years younger, the shock percentage of 20% is too low. For a portfolio with people that are 10 years older, the shock is too conservative. For the portfolios containing also partner pensions, differences between the internal model and standard approach are a little bigger. The results thus suggest that ASR could benefit from implementing an internal model for the determination of the SCR for longevity risk, but I am not convinced that these benefits would outweigh the associated costs.

A shortcoming of the standard approach that I find, is that the way the standard approach is modeled yields that the SCR as fraction of the best estimate present value of the liabilities is increasing with the average age of the portfolio. This is counterintuitive to me, but the effect is also partially present in my proposal for an internal model. This might be due to the relatively small number of simulations used in the model. When parameter uncertainty would be neglected, this effect disappears. It should furthermore be noted that the needed SCR according to the internal model is decreasing with age. Therefore it might still be beneficial to implement an internal model when the average age of people in the portfolio is high. Think for instance of a portfolio of immediate annuities. This is in line with results from other papers.

Regarding future research, some improvements can be made to the research performed here. As discussed in an earlier paragraph, ASR needs to collect more data in order to come to more advanced models and more reliable conclusions regarding the mortality in the portfolio. The underlying model for population mortality can be improved as well. Throughout this thesis I have made some assumptions regarding this model. I have for instance based the prediction of future mortality rates in the population entirely on the Lee-Carter method, whereas there are also other accurate models known in the literature. This form of model risk is neglected in my analysis and could be incorporated in future research. Furthermore I have assumed that there is no cohort effect for mortality in the Netherlands since there is, to my knowledge, no clear evidence that one does exist. For future research this could however be investigated further. Another possible point of improvement regards the extrapolation of the time series κ_t. As commonly done in the literature, I have adopted a linear ARIMA(0,1,0) time series model. Even though empirical evidence from the past suggests that this model performs well for low mortality countries, more careful analysis could show that a slightly more advanced model provides a better fit and more reliable future estimates. It could also be the case that the fit of the model improves when the penalized least squares approach for smoothing b_x would successfully be implemented, instead of the polynomial regression used now.

The internal model for the calculation of the Solvency Capital Requirement can be improved as well. The results of this approach could for instance be compared to results obtained when analytical techniques, instead of simulations, are used. One could for instance do this by means of the approach by Stevens et al. (2010). This analysis becomes however also more complicated when parameter uncertainty is considered and no closed-form analytical solutions exist. Some other recommendation is to do this analysis again, but with more simulations. This however requires a significant amount of time and possibly the use of faster computer software. Another proposal

for improvement would be to consider more advanced input data. This regards the portfolio used, where for instance premiums to be paid could be included, the portfolio could be set open to new participants and a stochastic term structure for the interest rates could be implemented. As it is now, I have neglected interest rate risk, assuming that the rates will not change over time. Of course this is not really in line with practice, so either the use of this model should be combined with the appropriate hedging of interest rate risk in the market, or it should include a stochastic term structure of interest rates if it would actually be implemented in practice. Furthermore individual mortality risk is not included. Even though the portfolio is large and the effects of the individual mortality risk being implemented are likely to be not big, results will be (slightly) more accurate when this form of risk is considered as well.

Chapter 7

References

Actuarieel Genootschap (2010). AG-Prognosetafel 2010-2060. Available from AG website.

ASR (2011). The essence of trust. ASR 2011 Annual Report. Available from ASR website.

Börger, M. (2010). Analysis of the Solvency II Standard Model Approach to Longevity Risk. Working paper.

Brouhns, N. and Denuit, M. (2001). Risque de longévité et rentes viagères. III. Elaboration de tables de mortalité prospectives pour la population assurée belge, et évaluation du coût de l'antisélection. Belgian Actuarial Bulletin, 2, 64-72.

Brouhns, N., Denuit, M., Vermunt, J.K. (2002). A Poisson log-bilinear regression approach to the construction of projected life-tables. Insurance: Mathematics and Economics, 31, 373-393.

Bruggink, J.-W. (2009). Ontwikkelingen in (gezonde) levensverwachting naar opleidingsniveau. Bevolkingstrends, 57(4), 71-75.

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

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Epstein, D., Ong, A., Balevich, I. (2009). A quantitative comparison of stochastic mortality models using data from England & Wales and the United States. North American Actuarial Journal, 13, 1-35.

Carolina, N. and Van Duin, C. (2010). Onzekerheidsmarges voor de sterfteprognose van het CBS. Bevolkingstrends, 58(2), 32-37.

CEIOPS (2009). Final CEIOPS' Advice for Level 2 Implementing Measures on Solvency II: Technical Provisions - Article 86(d) Calculation of the Risk Margin. Available from EIOPA website.

CEIOPS (2010). QIS5 Calibration Paper. Available from EIOPA website.

Coelho, E., Magalhães, M.G. and Bravo, J.M. (2007). Mortality and Longevity Projections for the Oldest-Old in Portugal. Work session on demographic projections, Bucharest, 10-12 October 2007, 117-132.

Cossette, H., Delwarde, A., Denuit, M., Guillot, F. and Marceau, E. (2007). Pension Plan Valuation and Mortality Projection: A Case Study with Mortality Data. North American Actuarial Journal, 11, 1-34.

Delwarde, A., Denuit, M. and Eiler, P. (2007). Smoothing the Lee-Carter and Poisson log-bilinear

models for mortality forecasting : A penalized log-likelihood approach. Statistical Modelling, 7(1), 29-48.

Denuit, M. (2007). Comonotonic approximations to quantiles of life annuity conditional expected present value. Insurance: Mathematics and Economics, 42, 831-838.

Denuit, M. and Goderniaux, A.-C. (2005). Closing and projecting life tables using log-linear models. Bulletin of the Swiss Association of Actuaries, 29-49.

Dowd, K., Cairns, A.J.G., Blake, D., Coughlan, G.D., Epstein, D., Khalaf-Allah, M. (2010). Evaluating the goodness of fit of stochastic mortality models. Insurance: Mathematics and Economics, 47, 255-265.

European Union (2009). Directive of the European Parliament and of the Council on the taking-up and pursuit of the business of insurance and reinsurance (Solvency II) (Recast).

European Commission (2010). QIS5 Technical Specification. Annex to Call for Advice from CEIOPS on QIS5. Available from EIOPA website.

Girosi, F. and King, G. (2007). Understanding the Lee-Carter Mortality Forecasting Method. Working paper.

Gschlössl, S., Schoenmaekers, P. and Denuit, M. (2011). Risk classification in life insurance: methodology and case study. European Actuarial Journal, 1, 23-41.

Haberman, S. and Renshaw, A.E. (2011). A comparative study of parametric mortality projection models. Insurance: Mathematics and Economics, 48(1), 35-55.

Jansen, A. (2012). Risk Margin for Mortality Risk. Master Thesis Executive Master of Actuarial Science.

Jasilionis, D. (2011). About Mortality Data for the Netherlands. Documentation from Human Mortality Database. Available from HMD website.

Lee, Ronald D. and Carter, L. (1992). Modeling and Forecasting U.S. Mortality: Differentials in Life Expectancy by Sex. International Journal of Forecasting, 8(3), 393-412.

Lee, Ronald D. (2000). The Lee-Carter method of forecasting mortality, with various extensions and applications. North American Actuarial Journal, 4, 80-93.

Li, N. and Lee, R. (2005). Coherent Mortality Forecasts for a Group of Populations: An Extension of the Lee-Carter Method. Demography, 42, 575-594.

Luijben, A.H.P. and Kommer, G.J. (2010). Tijd en Toekomst. Deelrapport van de VTV 2010 Van gezond naar beter. Available from RIVM website.

Peters, F., Nusselder, W. and Mackenbach, J. (2012). Assessing the model risk of longevity. The projection model of the Actuarial Society in comparison with two other Dutch projections. Netspar NEA paper, yet to be published.

Pitacco, E., Denuit, M., Haberman, S., and Olivieri, A. (2009). Modelling Longevity Dynamics for Pension and Annuity Business. Oxford University Press.

Plat, R. (2009). Stochastic Portfolio Specific Mortality and the Quantification of Mortality Basis

Risk. Insurance: Mathematics and Economics, 45(1), 123-132.

Plat, R. (2009a). On stochastic mortality modeling. Insurance: Mathematics and Economics, 45(3), 393-404.

Renshaw, A.E. and Haberman, S. (2006). A cohort-based extension to the Lee-Carter model for mortality reduction factors. Insurance: Mathematics and Economics, 38(3), 556-570.

Stevens, R., De Waegenaere, A. and Melenberg, B. (2010). Calculating capital requirements for longevity risk in life insurance products. Using an internal model in line with Solvency II. Available from Netspar website.

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

Towers Perrin (2009). UNESPA Longevity Risk Investigation. Confidential report.

Van Duin, C. and Garssen, J. (2010). Bevolkingsprognose 2010-2060: sterkere vergrijzing, langere levensduur. Available from CBS website.

Van Duin, C., De Jong, G., Stoeldraaijer, L. and Garssen, J. (2010). Bevolkingsprognose 2010-2060: model en veronderstellingen betreffende de sterfte. Available from CBS website.

Van Broekhoven, H. (2002). Market Value of Liabilities Mortality Risk: A Practical Model. North American Actuarial Journal, 6(2), 95-106.

Van de Poel, J., Palm, F. and Nijman, T. (2010). De validiteit van het door het Actuarieel Genootschap ontwikkelde prognosemodel voor overlevingskansen. Rapport van de commissie van deskundigen. Available from AG website.

Wilmoth, J.R. (1993). Computational methods for fitting and extrapolating the Lee-Carter model of mortality change. Technical report. Department of Demography. University of California, Berkeley.

Wilmoth, J.R. (1995). Are Mortality Rates Falling at Extremely High Ages: An Investigation Based on a Model Proposed by Coale and Kisker. Population Studies, 49(2), 281-295.

Wilmoth, J.R. (2000). Demography of longevity: past, present and future trends. Experimental Gerontology, 35, 1111-1129.

Workgroup PLT (2010). Generatietafels Pensioenen 2010. Verbond van Verzekeraars. Material only intended for pension insurers.

7.1 Data

Statistics Netherlands. Annual death counts by sex and single year of age (1970-2009). Data obtained through the Human Mortality Database (www.mortality.org) on May 4, 2012.

ASR pension portfolio data was acquired within the company.

Nominal term structure on zero coupon rates as used by DNB as they were on June 30, 2012. Obtained on July 4, 2012.

Appendix A

Graphs Lee-Carter model

This appendix presents graphs of the parameter estimates of the Lee-Carter model. Comments on the graphs are given in section 3.4.

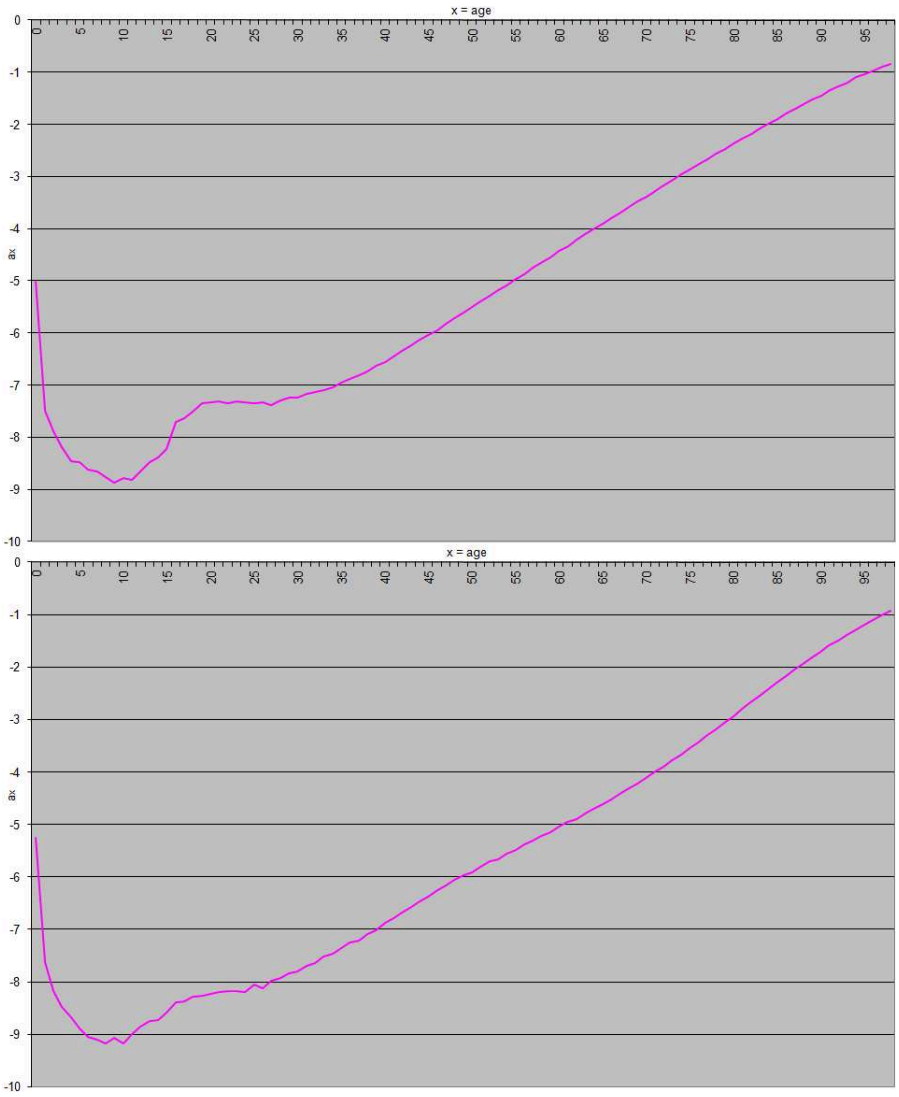

Figure A.1: *Estimates for a_x according to equation* (3.15)*. Men in upper panel, women in lower.*

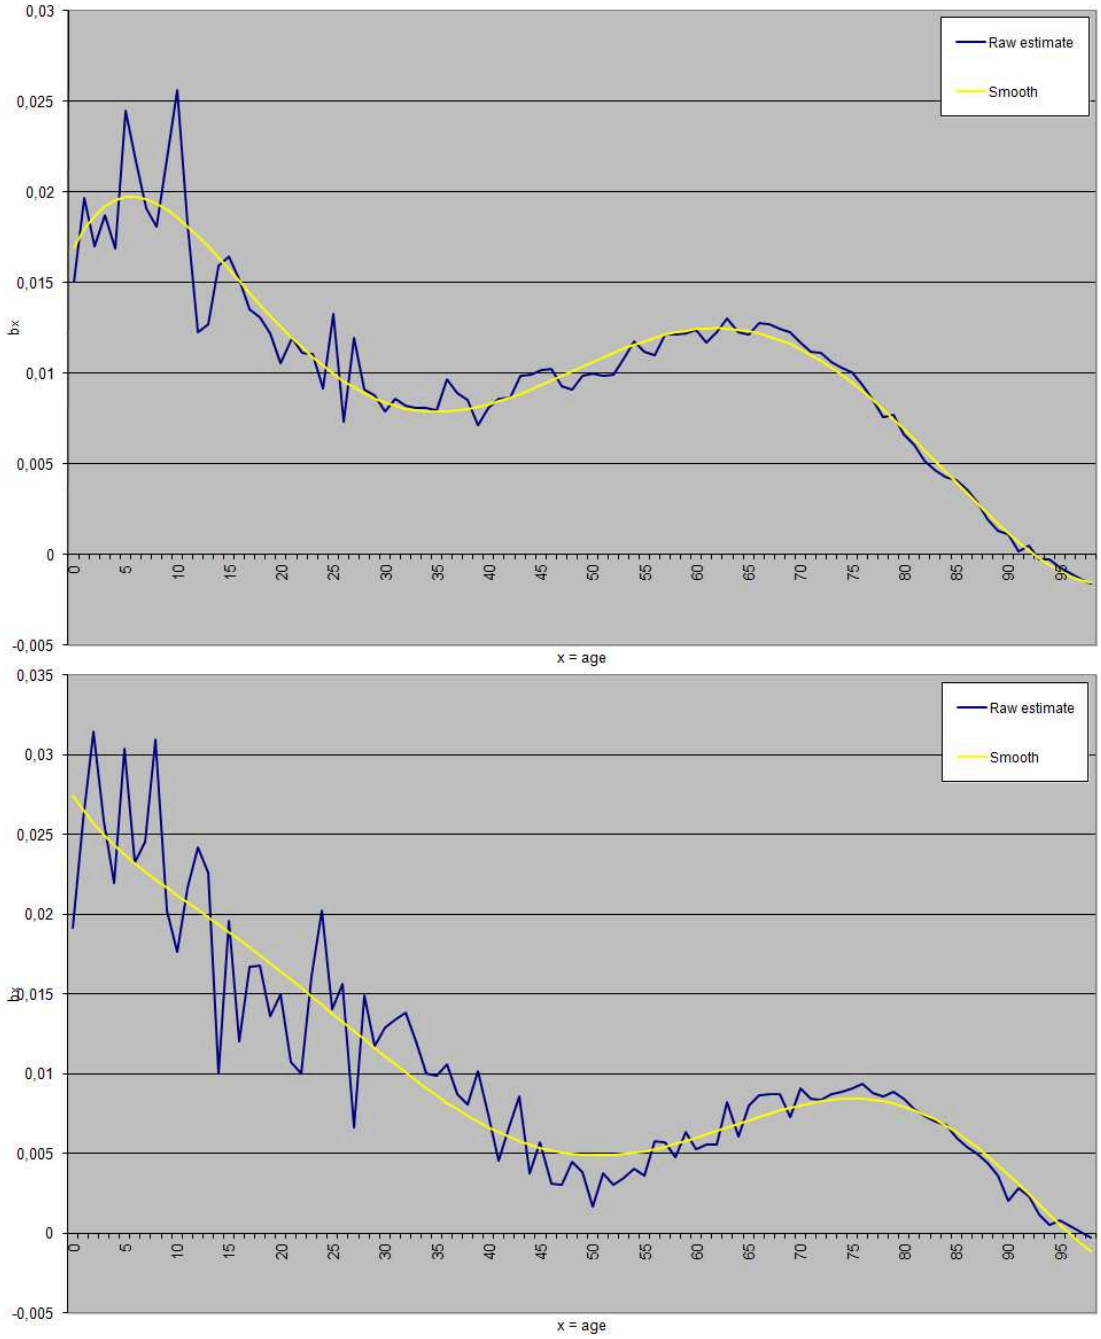

Figure A.2: *Estimates for b_x according to equation* (3.18) *and smoothed version (using polynomial regression). Men in upper panel, women in lower.*

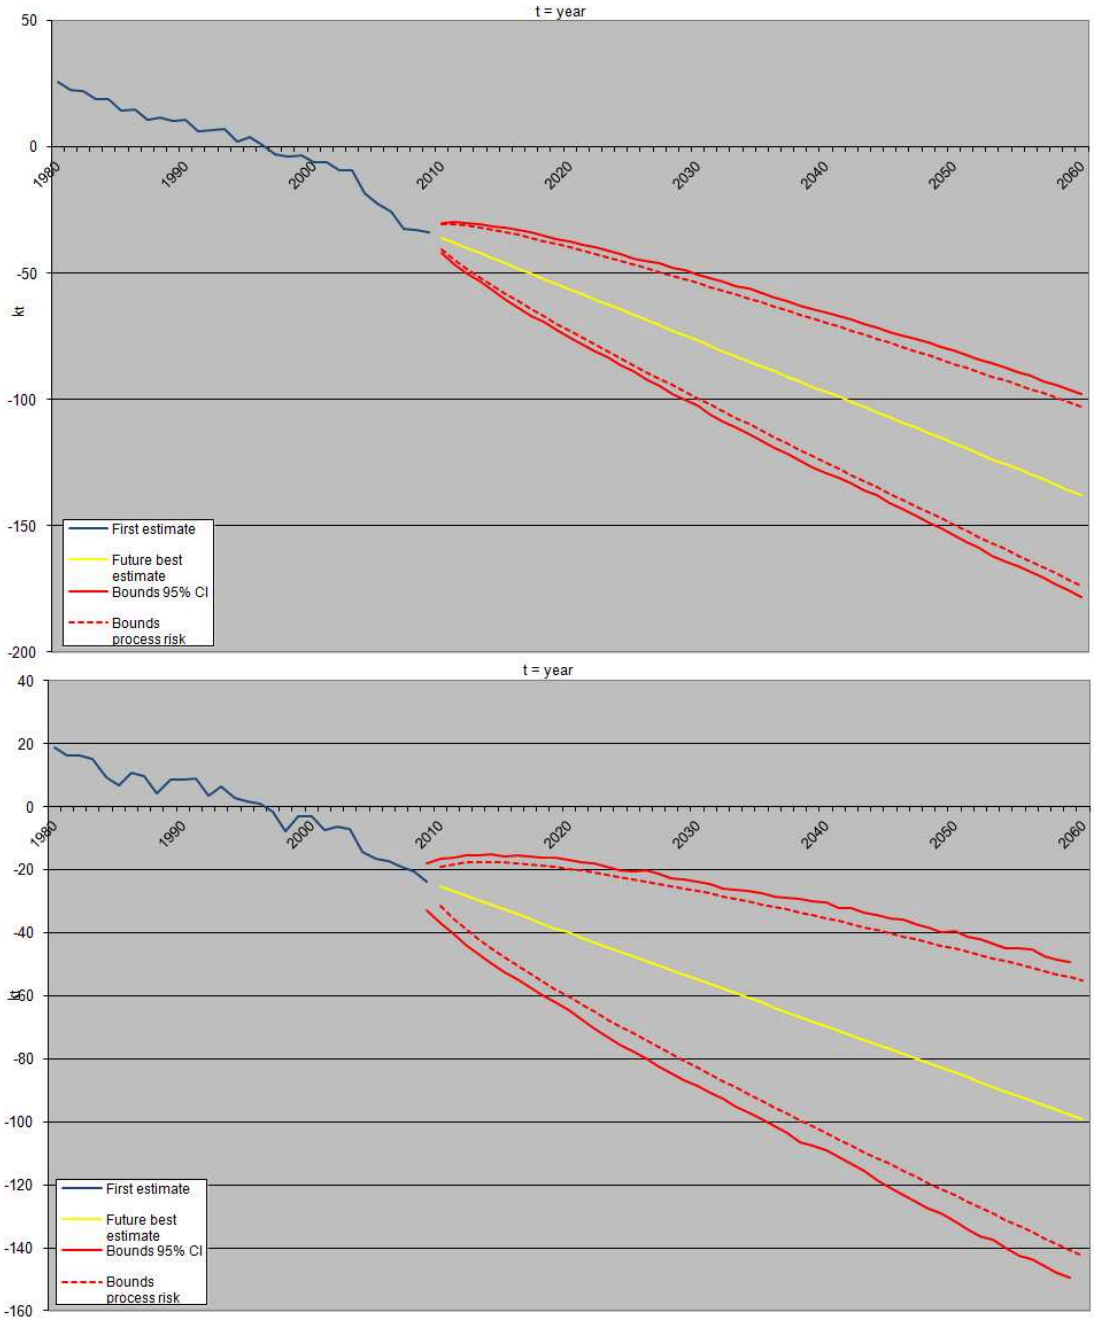

Figure A.3: *Estimates for κ_t according to equation* (3.19)*, future best estimates and 95% confidence intervals. Confidence intervals are based on only process risk (dashed lines) and on process risk and parameter risk (solid lines). Men in upper panel, women in lower.*

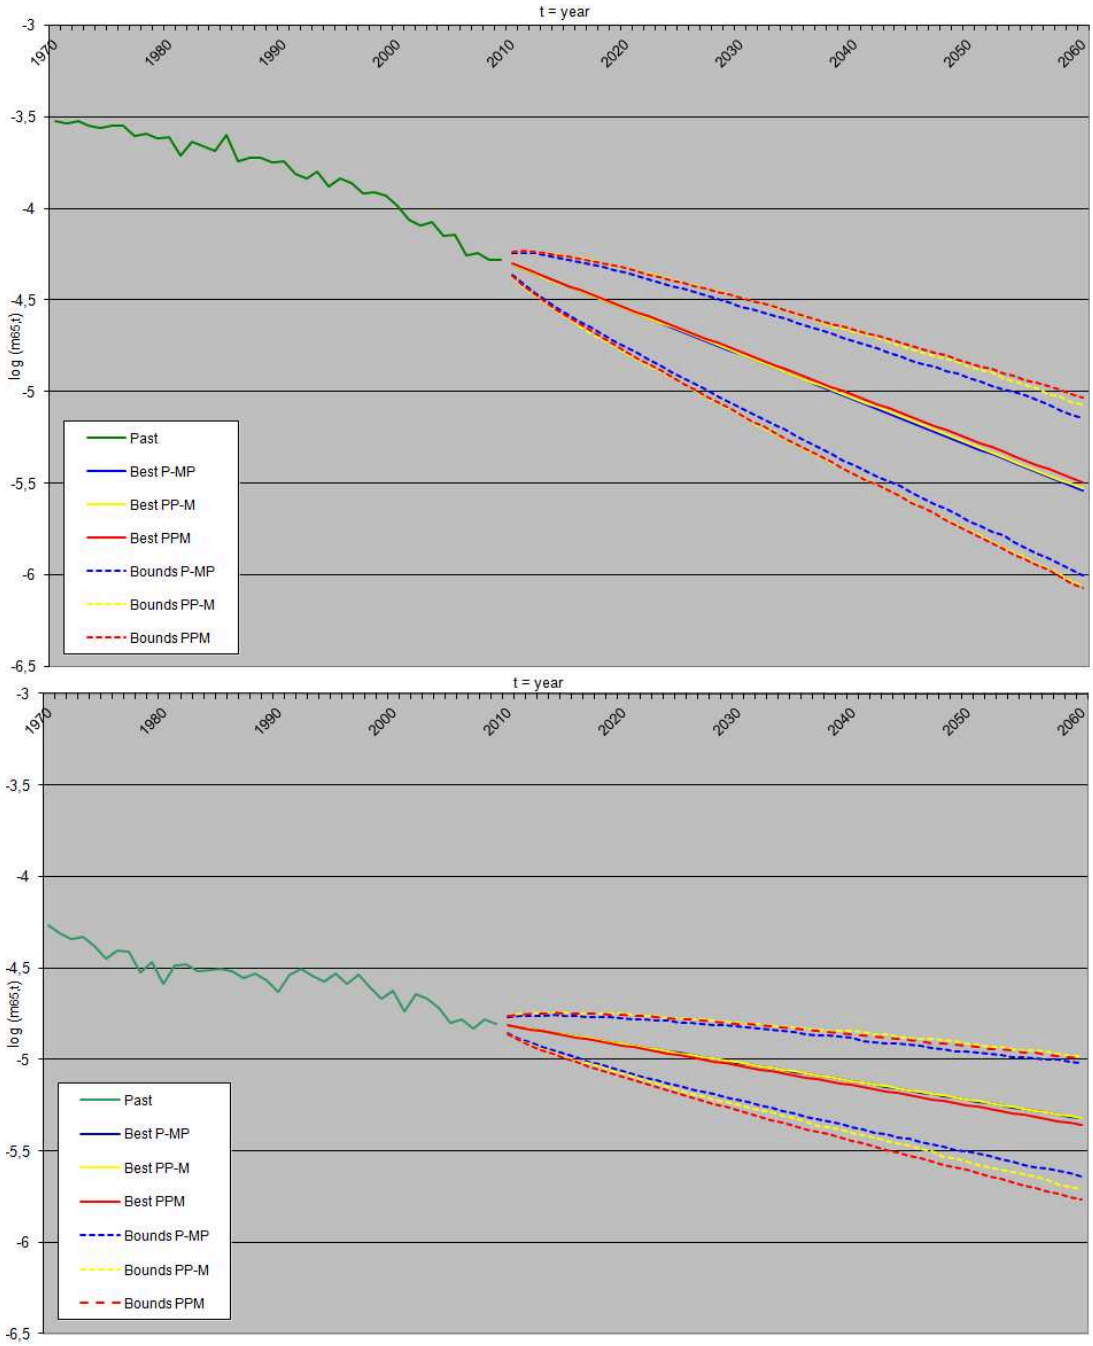

Figure A.4: *Best estimates (solid lines) of logarithm of central death rates for people aged 65, with 95% confidence intervals (dashed lines). P-MP means only process risk is considered, PP-M means process risk and parameter risk are considered and PPM means that also model risk is (partially) taken into account, by considering multiple starting periods. Men in upper panel, women in lower.*

Appendix B

Graphs portfolio mortality

In this appendix, graphs for the part on portfolio mortality (section 4.4) are displayed.

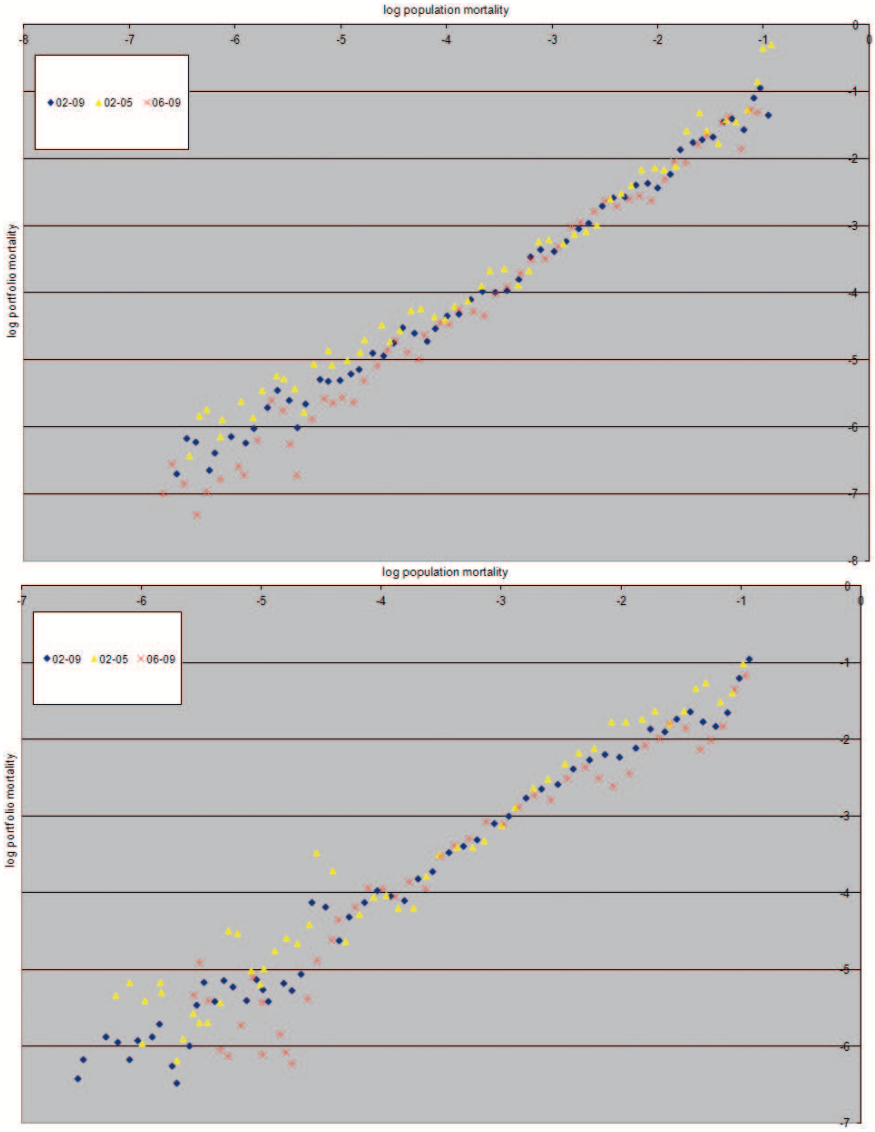

Figure B.1: *Scatter plots of* $\log \widetilde{m}_{x,t}^{PORT}$ *against* $\log m_{x,t}^{POP}$ *for different time periods. 02-09 refers to period 2002-2009, 02-05 refers to period 2002-2005 and 06-09 refers to period 2006-2009. Men in upper panel, women in lower.*

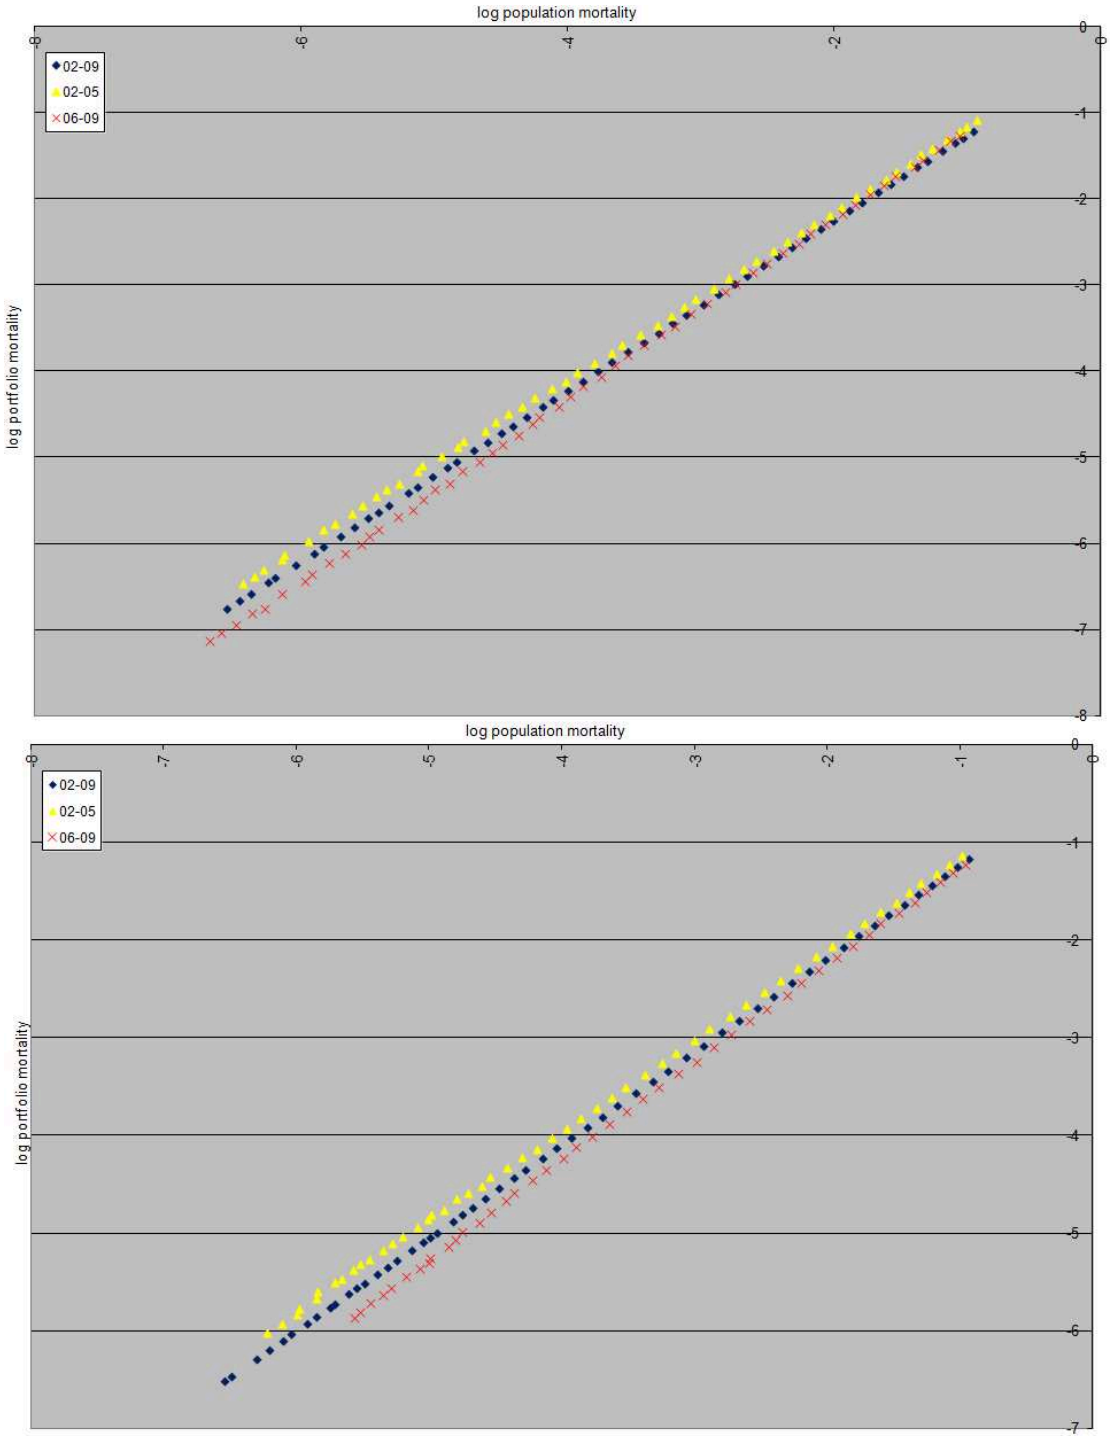

Figure B.2: *Fits of* $\log \widetilde{m}_{x,t}^{PORT}$ *against* $\log m_{x,t}^{POP}$ *for different time periods, using the Cox Proportional Hazard model. Men in upper panel, women in lower.*

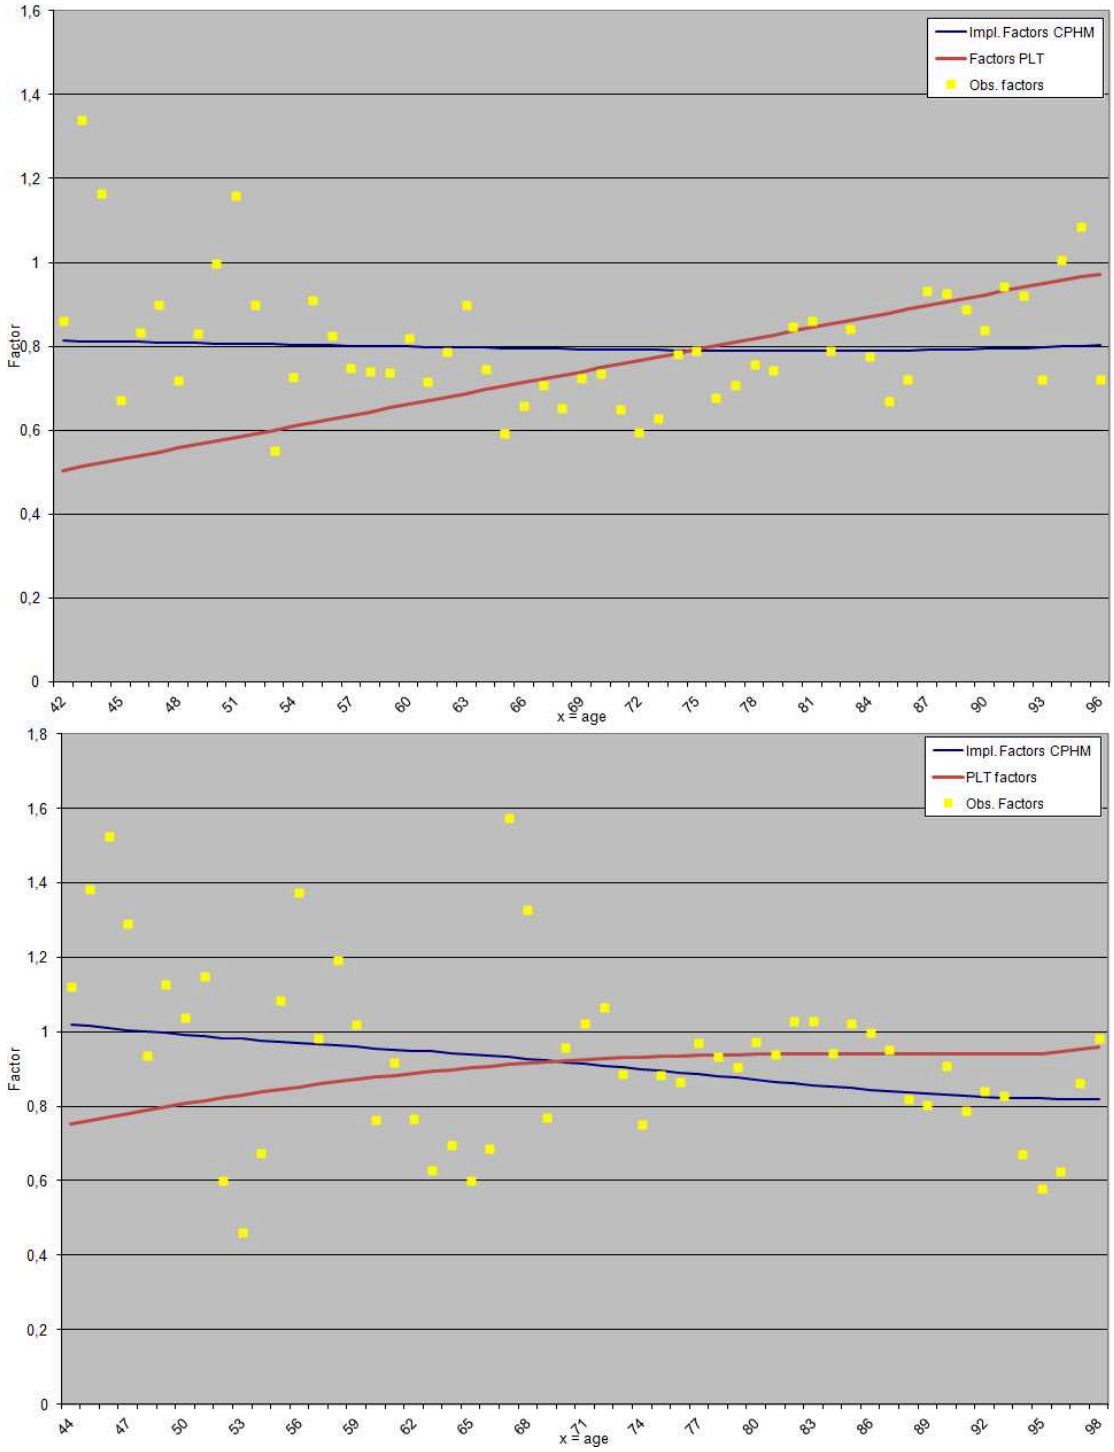

Figure B.3: *Plots of implied portfolio mortality factors (factors by which one should multiply population mortality on the q_x-scale to obtain portfolio mortality) for period 2002-2009, coming from the Cox Proportional Hazard model (CPHM), the portfolio mortality factors found by Workgroup PLT (2010) and the observed factors from the data. Men in upper panel, women in lower.*

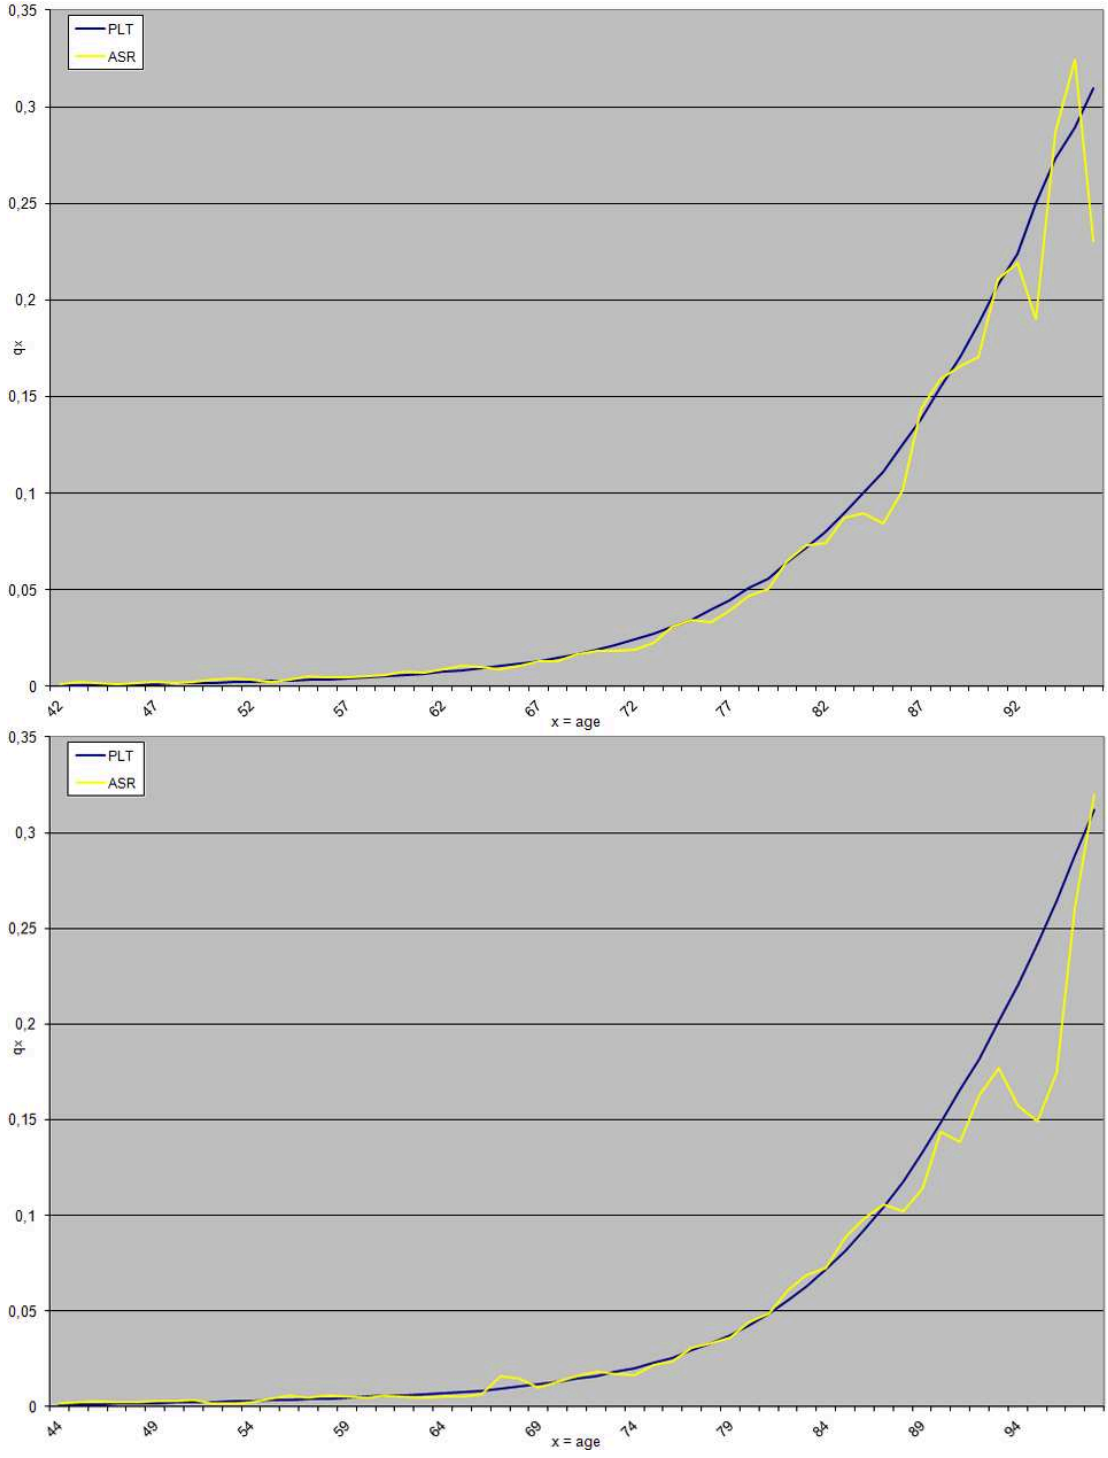

Figure B.4: *Plots of implied life tables for ASR in the period 2002-2009 (based on observed death rates), versus the implied life table from PLT (population mortality over 2002-2009, multiplied by population mortality factors from PLT). Men in upper panel, women in lower.*

Appendix C

Graphs SCR for longevity risk

In this Appendix some graphs for the Solvency Capital Requirement for longevity risk are displayed. Comments on the graphs can be found in section 5.3.

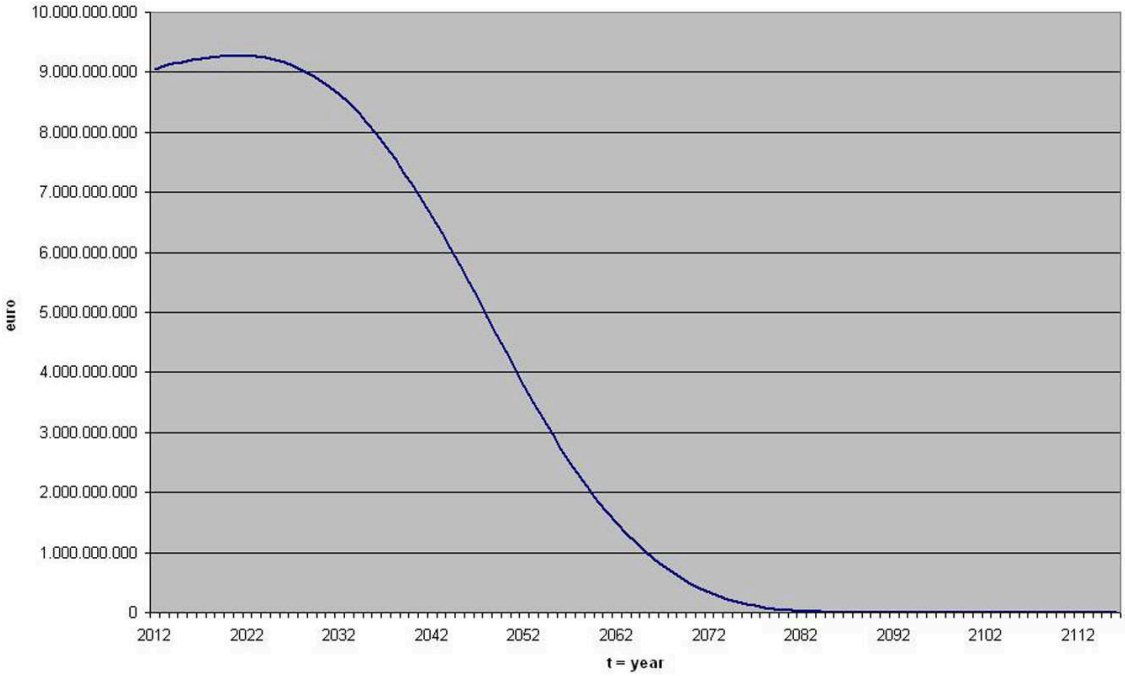

Figure C.1: *The best estimate present value of the liabilities over the years, for the portfolio consisting only of old age pensions. Portfolio is closed, discount factors are based on the term structure of a zero coupon bond issued by DNB, as they were on June 30, 2012.*

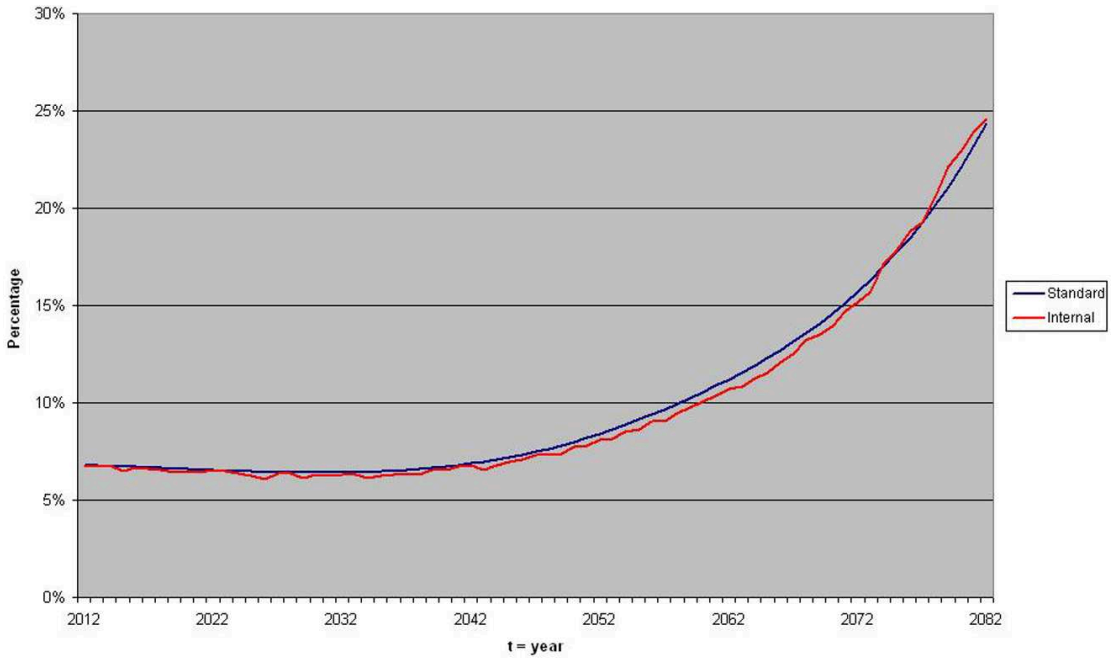

Figure C.2: *Pattern of the development of (the expected value of) the SCR as fraction of the best estimate present value of the liabilities over the years, for the portfolio consisting only of old age pensions. Lines are drawn for the standard approach as proposed in QIS5 and for the internal model as proposed in this thesis.*

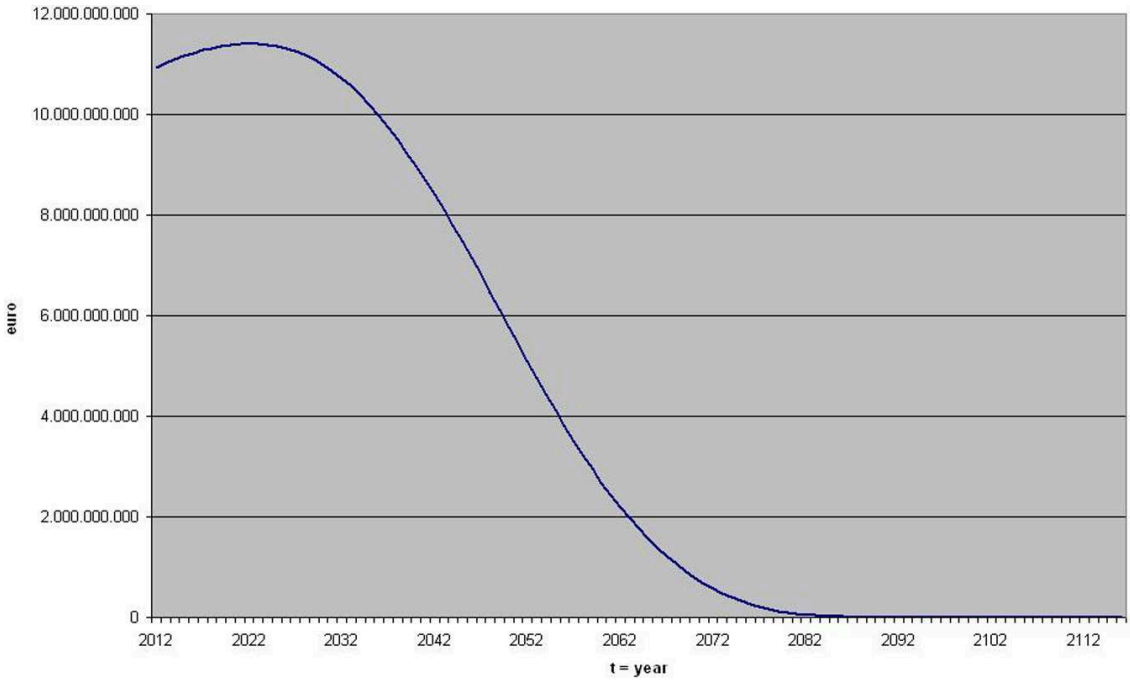

Figure C.3: *The best estimate present value of the liabilities over the years, for the portfolio consisting of old age pensions and partner pensions.*

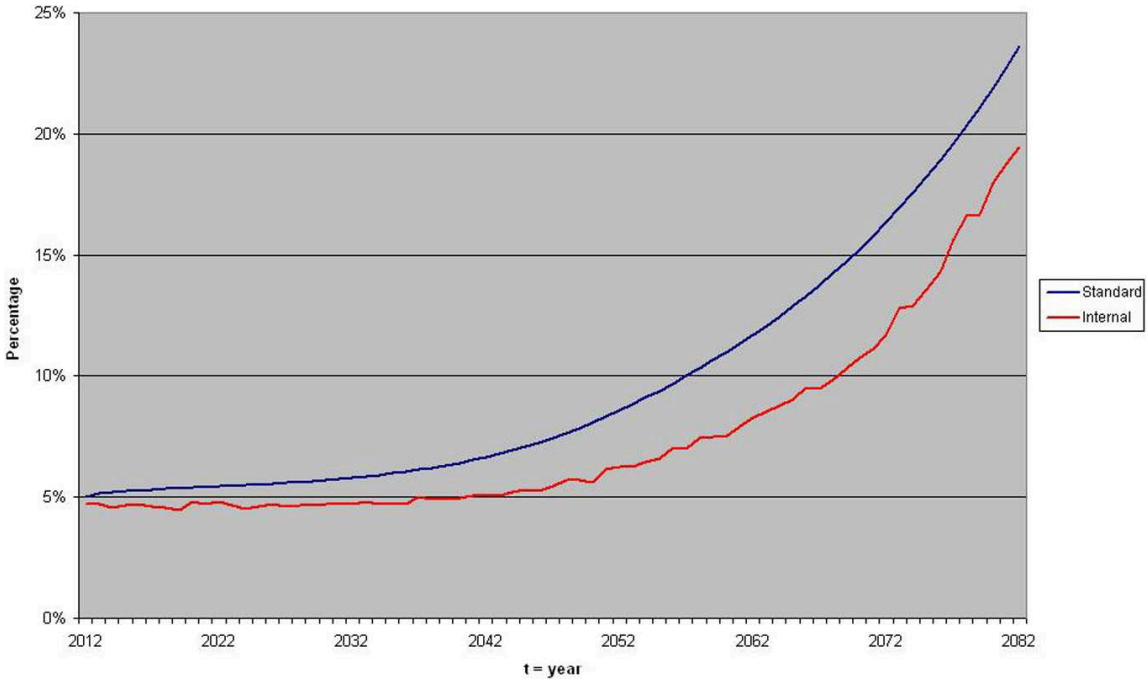

Figure C.4: *Pattern of the development of the SCR as fraction of the best estimate present value of the liabilities over the years, for the portfolio consisting of old age pensions and partner pensions. Lines are drawn for the standard approach and for the internal model.*

Appendix D

Discount factors used for SCR

In this appendix the discount factors used for the determination of the Solvency Capital Requirement are displayed. Discount factors are based on the nominal term structure for zero coupon bonds, as issued by DNB on June 30, 2012.

Term	Factor	Term	Factor	Term	Factor	Term	Factor
1	0.9908	28	0.5298	54	0.2875	80	0.1559
2	0.9816	29	0.5196	55	0.2805	81	0.1524
3	0.9699	30	0.5097	56	0.2738	82	0.1489
4	0.9539	31	0.4982	57	0.2671	83	0.1454
5	0.9346	32	0.4870	58	0.2607	84	0.1421
6	0.9124	33	0.4758	59	0.2544	85	0.1388
7	0.8890	34	0.4651	60	0.2482	86	0.1357
8	0.8650	35	0.4546	61	0.2425	87	0.1325
9	0.8409	36	0.4443	62	0.2369	88	0.1295
10	0.8166	37	0.4343	63	0.2315	89	0.1265
11	0.7922	38	0.4244	64	0.2261	90	0.1236
12	0.7685	39	0.4148	65	0.2209	91	0.1208
13	0.7463	40	0.4054	66	0.2159	92	0.1180
14	0.7246	41	0.3957	67	0.2109	93	0.1153
15	0.7035	42	0.3861	68	0.2061	94	0.1127
16	0.6877	43	0.3766	69	0.2013	95	0.1101
17	0.6703	44	0.3675	70	0.1967	96	0.1075
18	0.6542	45	0.3586	71	0.1922	97	0.1051
19	0.6387	46	0.3499	72	0.1878	98	0.1027
20	0.6234	47	0.3414	73	0.1835	99	0.1003
21	0.6105	48	0.3332	74	0.1793	100	0.0980
22	0.5978	49	0.3250	75	0.1752	101	0.0957
23	0.5853	50	0.3172	76	0.1711	102	0.0935
24	0.5732	51	0.3095	77	0.1672	103	0.0914
25	0.5613	52	0.3020	78	0.1634	104	0.0893
26	0.5506	53	0.2947	79	0.1596	105	0.0873
27	0.5401						

Table D.1: *Table with discount factors used for calculating the Solvency Capital Requirement for longevity risk in ASRs pensions portfolio.*

Appendix E

Penalized least squares smoothing

This appendix discusses the smoothing of parameter b_x during estimation. Besides smoothing b_x after the estimation process, it is also possible to use a method developed by Delwarde et al. (2007) where b_x is already smoothed during the process of estimation itself. The advantage of this approach is that it also accounts for possible changes of value in the other parameters a_x and κ_t.

Delwarde et al. (2007) employed a penalized least squares approach in which they minimize the objective function

$$\sum_{x_1}^{x_m}\sum_{t_1}^{t_n}(\log m_{x,t} - a_x - b_x\kappa_t)^2 + b'P_b b, \tag{E.1}$$

where (with π_b a smoothing parameter)

$$P_b = \pi_b \Delta'\Delta \quad \text{with } \Delta = \begin{bmatrix} 1 & -2 & 1 & 0 & \dots & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & -2 & 1 \end{bmatrix}. \tag{E.2}$$

The term $b'P_b b$ yields a penalty for irregular estimates of b, measured by the sum of the square of the second order differences $b_{x+2} - 2b_{x+1} + b_x$. Hence, the objective function represents a trade-off between the goodness of fit and the smoothness of b. One can easily see that the larger π_b, the more smooth b will be. Delwarde et al. (2007) argue that as $\pi_b \to \infty$, the fit becomes linear. The optimal choice of the smoothing parameter π_b is based on cross validation.

When there is an initial set of estimates $(\widehat{a}^{(0)}, \widehat{b}^{(0)}, \widehat{\kappa}^{(0)})$ available, one can enter the iterative procedure below. This algorithm is based on standard Newton-Raphson methods, iteratively minimizing the objective function (E). As initial set of estimates I have used the set obtained via the standard SVD method.

$$C_a^{(h)}\widehat{a}^{(h+1)} = C_a^{(h)}\widehat{a}^{(h)} + r_x^{(h)} \tag{E.3}$$

$$C_\kappa^{(h)}\widehat{\kappa}^{(h+1)} = C_\kappa^{(h)}\widehat{\kappa}^{(h)} + r_\kappa^{(h)} \tag{E.4}$$

$$(C_b^{(h)} + \mathrm{diag}P_b)\widehat{b}^{(h+1)} = (C_b^{(h)} + \mathrm{diag}P_b - P_b)\widehat{b}^{(h)} + r_b^{(h)}. \tag{E.5}$$

Here, diagPb is a diagonal matrix containing the main diagonal of P_b. $C_a^{(h)}, C_\kappa^{(h)}$ and $C_b^{(h)}$ are diagonal matrices of dimensions $(m \times m), (n \times n)$ and $(m \times m)$ respectively. They are defined by

$$(C_a^{(h)})_{x_i,x_i} = n \tag{E.6}$$

$$(C_\kappa^{(h)})_{t_j,t_j} = \sum_{x_1}^{x_m} (\widehat{b}_x^{(h)})^2 \tag{E.7}$$

$$(C_b^{(h)})_{x_i,x_i} = \sum_{t_1}^{t_n} (\widehat{\kappa}_t^{(h+1)})^2. \tag{E.8}$$

The column vectors $r_a^{(h)}, r_\kappa^{(h)}$ and $r_b^{(h)}$ have dimensions m, n and m respectively and are defined by

$$(r_a^{(h)})_{x_i} = \sum_{t_1}^{t_n} (\log m_{x_i,t} - \widehat{a}_{x_i}^{(h)} - \widehat{b}_{x_i}^{(h)} \widehat{\kappa}_t^{(h)}) \tag{E.9}$$

$$(r_\kappa^{(h)})_{t_j} = \sum_{x_1}^{x_m} \widehat{b}_x^{(h)} (\log m_{x,t_j} - \widehat{a}_x^{(h+1)} - \widehat{b}_x^{(h)} \widehat{\kappa}_{t_j}^{(h)}) \tag{E.10}$$

$$(r_b^{(h)})_{x_i} = \sum_{t_1}^{t_n} \widehat{\kappa}_t^{(h+1)} (\log m_{x_i,t} - \widehat{a}_{x_i}^{(h+1)} - \widehat{b}_{x_i}^{(h)} \widehat{\kappa}_t^{(h+1)}). \tag{E.11}$$

Delwarde et al. (2007) state that convergence only requires a few steps. No ending criterion is specified, but one could for instance use a criterion based on the total quadratic change in b_x.