Hedging longevity risk with longevity swaps

Master Thesis Quantitative Finance

Heleen Westland (307367)

Erasmus University supervisor: dr.H.J.W.G. Kole
Nationale-Nederlanden supervisor: drs.R.H.J. Rondagh

March 2009

Abstract

Longevity risk is the risk that people on average grow older than anticipated on; it is the negative effect of changes in mortality rates on the risk measures of companies involved in pensions. As a pension provider, the Group Life department of Nationale-Nederlanden is heavily confronted with this risk. A new possibility to cover (part of) the longevity risk arises with the emergence of a market in longevity derivatives: over-the-counter products based on indices of mortality rates, analogous to products in the established financial markets. One of these products is the longevity swap. This thesis examines the possibilities longevity swaps might offer in hedging the longevity risk in a standard group pension contract of Nationale-Nederlanden. A portfolio of longevity swaps is attached to the liability stemming from the contract, after which the hedge is evaluated in terms of its effects on a risk margin called the Market Value Margin. Hedging the standard group pension contract results in a reduction of the Market Value Margin of 29%, while the risk capital required to cover the longevity risk decreases substantially as well (20%). However, the hedge is not cost-effective: the decrease in Market Value Margin is smaller than its expected costs.

Contents

1 Introduction

Pension providers grant periodical payments to policyholders from their retirement onwards until the time of their death. The average timespan between these two events is steadily increasing. In 1951, the expected remaining lifetime of a 65 year old Dutch male was 13 years; a female could expect to live for 14 more years (JP Morgan, 2007). These numbers increased to 16 and 20, respectively, in 2006.

The increase in life expectancy implies that the pension payments have to be paid over a longer period. Part of the increase can be forecasted using scientific models, and hence anticipated on in the determination of the pension premiums.

According to Loeys et al (2007), inclusion of mortality forecasts in the past still resulted in an underestimation of the actual present value of the payments due: the improvements in life expectancy appeared to be stronger than forecasted. The reasons for the extra improvements are diverse, and their pattern has been different at different age groups over time (Cairns et al, 2005). Obviously, the extra decrease in mortality rates can not be anticipated on in the premium determination. The risk that the accumulated pension premiums turn out to be too low for the actual present value of the payments is called longevity risk.

Pension providers can respond to longevity risk in three different ways (Blake et al, 2006):

- do nothing;
- hedge the risk by reinsurance;
- hedge the risk by longevity derivatives (longevity-linked securities).

Decreases in interest rates increased the relevance of hedging for pension providers in the 1990s. In a high interest rate environment, fixed income returns had been strong enough to offset underestimates of longevity increases. The changing interest rates enlarged the consequences of misjudgements. A striking example is the history of the British life insurer Equitable Life. Embedded options in its pension annuity contracts became very valuable in the late 1990s as a result of improvements in longevity while interest rates were falling. Eventually, Equitable Life was forced to close for new business in 2000. The problems would probably have been avoided if Equitable Life had chosen to hedge its longevity and interest risks.

In the past, reinsurance was the only way to cover longevity risk. Because the costs for reinsurance are quite high, pension and annuity providers only applied it to a limited extent. In recent years, a strong increase in the demand for longevity reinsurance has been observed, as a result of the build-up to Solvency II (the future regulatory framework for insurers). Under Solvency II, reinsurance is recognized as a risk decreasing factor: it is deducted from the regulatory risk capital the insurer is obliged to hold. Shareholders require a certain return over this regulatory risk capital. The new reduction in this cost of capital could outweigh the cost of reinsurance: reinsurance has become attractive because of its opportunity costs (CEA, 2008; Risk Magazine, 2008). However, the capacity of reinsurers is limited because longevity risk cannot be diversified away, and as a consequence, most reinsurers do not accept new longevity risk anymore (Blake, 2008; Neville and Ho, 2006). The limited availability of longevity reinsurance highlights the relevance of longevity derivatives.

Academics have written about longevity derivatives since the early 2000s. In contrast to reinsurance contracts, these derivatives are not tailor-made to the specific needs of the insurer. Just like derivatives on financial assets, longevity contracts might be traded on an exchange (bonds, futures, options) or over-the-counter (swaps, forwards).

Despite the fact that these products all can be constructed in theory, their practical implementation turns out to be problematic. Until now, only one longevity derivative has been introduced in the financial markets: the EIB/BNP longevity bond (2004). This bond was marketed with a maturity of 25 years, and its coupon payments were attached to the number of survivors in a certain age group of British and Welsh males. The EIB/BNP longevity bond did not generate sufficient demand and was redrawn from the market for redesign.

In the over-the-counter category, there actually is one traded derivative: the longevity swap (Pengelly, 2008; Blake et al, 2006). In a longevity swap, the cash flows are tied to the difference between the actual number of surviving people in a certain age group according to a certain index, and the expected number set in advance.

In March 2007, JPMorgan made an attempt to boost the development of the longevity derivatives market by launching LifeMetrics, 'a toolkit for measuring and managing longevity and mortality risk' (Coughlan et al, 2007a). This platform consists of mortality indices for the United Kingdom, the United States, Germany and the Netherlands, as well as supporting technical documents and software. It is intended to help in the standardization and facilitation of longevity derivatives trading by removing part of the subjectivity in the determination of rates and prices, as these can be based on the indices now. JP Morgan also presents itself as counterparty in these products, which has resulted in a few transactions already (Longevity Risk Conference, September 2008).

This thesis examines whether longevity swaps offer opportunities for the longevity risk of the Group Life department of Nationale-Nederlanden. Within Nationale-Nederlanden (ING), risk is quantified as Economic Capital (EC). EC is the amount of risk capital a product adds to the company: the loss that will not be exceeded with 99.95% certainty. Nationale-Nederlanden has been assigned a certain maximum for its total EC as a subsidiary of ING. The capital charges over all present and future non-market ECs are combined in the Market Value Margin. The MVM for life risk is the most relevant price measure for longevity risk. Longevity swaps do take away part of the longevity risk: both the Market Value Margin and the Economic Capital associated to it decrease. This enlarges the room for other risky activities within Nationale-Nederlanden and ING, or it is an opportunity to decrease the investment risk for ING's shareholders. The expected costs of the swap are larger than the decrease of the MVM though. Therefore, Nationale-Nederlanden should not invest in longevity swaps.

1.1 Nationale-Nederlanden

Nationale-Nederlanden is one of the largest insurance companies in the Netherlands. The company was established in 1962 after a merge between De Nederlanden van 1845 and De Nationale Levensverzekering Bank. After a few years of merely official cooperation, in 1970 all life insurance activities were combined into Nationale-Nederlanden Levensverzekering Maatschappij NV; the property and casualty insurance activities were bundled in Nationale-Nederlanden Schadeverzekering Maatschappij NV. In 1991, Nationale-Nederlanden BV merged with NMB Postbank Group, into ING

Figure 1.1: Legal structure of Nationale-Nederlanden

The figure shows the legal structure of Nationale-Nederlanden Levensverzekering Maatschappij NV. One of the departments of this NV is Group Life.

Group. The current legal structure of the Nationale-Nederlanden Levensverzekering Maatschappij can be seen in figure 1.1.

In 2007, the balance sheet total of the Levensverzekering Maatschappij equaled €61.5bln, with a net income of €3.7bln (stand-alone insurance business (technical net income): €2.7bln). The book value of all pension liabilities of the Levensverzekering Maatschappij sums to €50bln (ultimo 2007). The Group Life part of this amount is €32.6bln. (Annual Report Nationale-Nederlanden Levensverzekering Maatschappij, 2007)

The Group Life department insures group pension contracts of both corporations and pension funds. In sum, around 17.000 contracts are outstanding (April 2008). The main part (13.000) of the contracts has premiums that are linked to the u-rate of interest, where the profit sharing takes place in advance. For larger companies, profit sharing takes place after their realization, and is based on the return on a specific portfolio of investments. The largest corporates effectively deposit their pension funds at Nationale-Nederlanden (this considers around 200 contracts). For them, Nationale-Nederlanden acts as a kind of reinsurer, while the influence over the investment policy is not completely lost.

1.2 Outline

In the remainder of this thesis, the theory of longevity swaps is worked out and applied to the longevity risk of the Group Life department of Nationale-Nederlanden. The main objective is to find out whether longevity swaps could reduce this risk.

We examine the background of the subject in chapter 2. First, an overview is provided of the scientific literature on the main theoretical concepts that are applied in later chapters: longevity, longevity risk and longevity derivatives. Second, the chapter sketches the setting of the longevity risk within the Group Life department of Nationale-Nederlanden. Internal documents and spreadsheets are used to determine how this risk is measured and handled.

Chapter 3 further zooms in on the longevity swap. It discusses this product in detail, where particular attention is paid to the pricing of the swap. The pricing techniques are based on the scientific literature of chapter 2. The developed methodology is used to work out an example of a swap, which is used in a sensitivity analysis to determine the reactions of its price to changes in financial and actuarial variables.

In chapter 4, a series of longevity swaps is applied to an existing group pension contract of the Group Life department of Nationale-Nederlanden. The liability stemming from this contract is extrapolated into the future. The methods used to determine the liability can be employed to calculate the different components of Economic Capital. The Market Value Margin follows from the Economic Capital figures at each future time.

The impact of a specific portfolio of longevity swaps on the longevity risk associated to the contract is examined by correcting the liability, EC and MVM of the cashflows arising from the contract for the cashflows arising from the swaps. The hedge is evaluated on basis of the difference in liability and MVM; hence its cost-effectiveness is determined. A sensitivity analysis gives an indication of the extent to which the assumptions used in the model influence the outcomes of the analysis.

With this information, we assess that longevity swaps do reduce the longevity risk for the Group Life department of Nationale-Nederlanden: the Market Value Margin is reduced by 29% and the Economic Capital in the first year is 20% below its original value. However, the price of the hedge makes it not cost-effective.

2 Background

This chapter introduces the concepts that are applied in later chapters. Section 2.1 discusses longevity. It begins with a short overview of longevity in the Netherlands, along with the risk it brings to pension providers in general and the way that is accounted for at Nationale-Nederlanden. The other main building block of this thesis is introduced in section 2.2: longevity derivatives. Both the existing products in the financial markets and the products that theoretically could exist in the future are discussed.

2.1 Longevity

2.1.1 Longevity in the Netherlands

In the Netherlands, mortality rates have been declining steadily ever since they were recorded. In 1850, when the first registration took place, the average lifetime of men was 38 years. Women lived two years longer. Currently, the average age at death for men is 76 years, women live around 81 years. A substantial part of this dramatic improvement can be ascribed to lower child mortality, but people also live longer at older ages. For example, in 1951 the expected remaining lifetime of a 65 year old Dutch male was 13 years, against 14 years for a female. These rates increased to 16 and 20, respectively, in 2006. (De Beer, 2006; JPMorgan, 2007)

An illustration of the process is given in figure 2.1, which depicts two aspects of the improvements in longevity. The data are obtained from the LifeMetrics Index on the Netherlands (JP Morgan, 2007). The left figure shows the number of years a 65-year-old person on average has left (life expectancy, $e(65)$) over time. The life expectancy has clearly increased over the period, where the increase was stronger for females. The right figure depicts the level of the mortality rates in time: the probability that a person aged 65 will die during the upcoming year ($m(65)$). The improvement in longevity (decrease in mortality rates) is evident, particularly for the females. More information on measures and symbols used can be found in Appendix A.

Figure 2.1: Longevity improvements in the Netherlands, 1951-2006

The left figure shows the life expectancy of men and women in the Netherlands who are 65 years old (the official retirement age).The right figure shows the development of their mortality rates.
Source: JP Morgan LifeMetrics Index Netherlands

2.1.2 Forecasting longevity

Coughlan et al (2007a) find that each additional year of life adds 3-4% to the value of pension liabilities. In order to avoid losses, it is therefore important to incorporate future longevity improvements in the premium policy. Several methods are available for the extrapolation of past trends into the future. According to Gregorkiewicz and Plat (2004), the most relevant and most frequently used methods in the Netherlands are the Lee-Carter model, the CRC model and the CBS model.

- Lee and Carter (1992) model the probability of a person aged x in year t to die during the upcoming year in a simple log linear model. Their methodology is focused on replicating past mortality rates as good as possible and subsequently extrapolating this model into the future.

  The basis equation is as follows:

  $$\ln m(x,t) = a(x) + b(x)k(t) \tag{2.1}$$

  Here, ln $m(x,t)$ is the natural logarithm of the central mortality rate as defined in appendix A; $a(x)$ is the observed ln$m(x,t)$ per age in the past; and $b(x)$ can be seen as the extent to which ln$m(x,t)$ depends on a general mortality index $k(t)$. The mortality rate hence is modeled as a constant that depends on age, plus a log linear trend that indicates how the rate changes over time. The model is estimated in four steps:

  1. $a(x)$, $b(x)$ and $k(t)$ are estimated on historical data.
  2. $k(t)$ is estimated again to match the historical observations exactly.
  3. The reestimated $k(t)$ parametervalues are replaced by a time series model for $k(t)$:

     $$k(t) = k(t-1) + \mu \tag{2.2}$$

  4. The $m(x,t)$ are projected to the future. Initial mortality rates $q(x,t)$ can be deducted as in appendix A:

     $$q(x,t) = 1 - e^{-m(x,t)} \tag{2.3}$$

  The model can be adjusted by assuming $b(x)$ to be time related, resulting in a model with multiple $k(t)$s. This makes it possible to incorporate different consequences of a certain factor for the mortality rates of different age groups. (Booth et al, 2002)

- The CRC model is built by the Commissie Referentietarief Collectief of the Actuarieel Genootschap (the Dutch Association of Actuaries) (Internal memo, 2007). The objective of the model was to take into account that mortality rates of the part of the population that is in a pension fund are lower than those of the population at large. Experts included the effect of their own experience as well as distinct trends for separate age groups in the estimated mortality tables.
  The CRC models the initial mortality rate per age over time ($q(t,x)$)as follows:

  $$q(t,x) = c(x) \cdot q(t-1,x)e^{E(t,x)} \tag{2.4}$$

  where $c(x)$ is the mortality reduction factor for age x and $E(t,x)$ is an error term (with an expectation of 0, a variance of $\sigma^2(x)$, and no correlation over ages nor over time). The model assumes that the mortality rates fall each year by a fixed percentage $c(x)$. Writing the model in logarithms results in

  $$\Delta y(t,x) = m(x) + E(t,x) \tag{2.5}$$

where $y(t,x)$ = ln$(q(t,x))$, $m(x)$ = ln$(c(x))$ and $\Delta y(t,x)$ = $y(t,x)$ - $y(t-1,x)$. It follows that the variance of $\Delta y(t,x)$ is $\sigma^2(x)$. The best estimate after k years ($\hat{y}(T,k)$) is subsequently calculated as

$$\hat{y}(T,k) = y(T) + km(x) \tag{2.6}$$

- The third alternative is the use of CBS data (De Jong and Van der Meulen, 2005). These data are publicly available data[1] on population mortality rates for the Netherlands, and are updated every two years. Besides current mortality rates, also estimates until 2050 are provided. The CBS estimates these projections separately per cause of death (cancer, hart diseases, diabetes, etcetera). The projections are a result of a combination of a statistical analysis (historical mortality rates per sex, cause of death and age group) and the opinions of medical specialists. The model cannot be translated into a quantitative model, as the medical assumptions are often subjective.

The estimates of the Lee-Carter, CRC and CBS methods are different. As an example, the estimated mortality rates for a Dutch male aged 65 in 2008 are depicted in figure 2.2. The CBS estimates appear to give the lowest mortality rate estimates. CRC roughly coincides with Lee-Carter until 2035, and is above it afterwards.

Figure 2.2: Mortality rates for males aged 65 in 2008

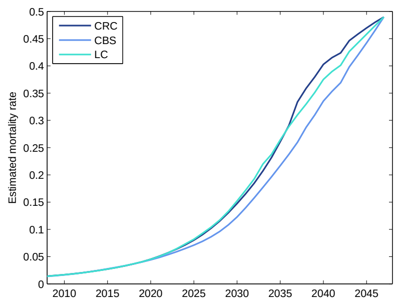

The figure depicts the initial mortality rates in time of males who are 65 years old in 2008, using the CBS, CRC and Lee-Carter methods.

2.1.3 Longevity risk

In the past, these (and similar) models have not been able to capture the full extent of the improvements in longevity (Loeys et al, 2007). An illustration of this is given in figure 2.3, which shows the increase in life expectancy projections for British males aged 60, during the past decades. In 1955, it was estimated that in 2030, a 60-years-old British male could expect to live for 17.9 more years. This forecast was re-estimated at 19.9 years in 1968, and increased further in the course of time, to 27.8 years eventually in 1999.

[1]The data can be obtained at http://statline.cbs.nl/statweb/.

Figure 2.3: Mortality rates for males aged 65 in 2008

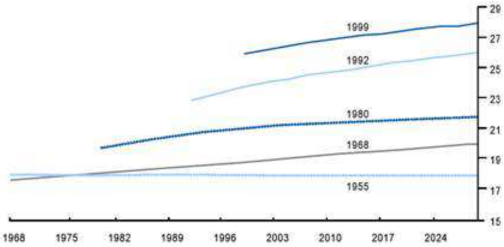

Life expectancy predictions for UK males aged 60 by year of projection.
Source: Coughlan et al, 2007a

Because a change in the realized mortality rates implies that the model forecasts have to be updated, losses are realized on the value of the pension liabilities at each estimate update. The risk of such an unanticipated increase in pension liabilities is called longevity risk. Idiosyncratic longevity risk is the risk that any particular policy holder lives longer than anticipated. This can be hedged by diversification (selling more policies). Systematic longevity risk is not diversifiable; it is the risk that policyholders on average live longer than expected (Cairns et al, 2006). In this thesis, the word longevity risk is used in the meaning of systematic longevity risk.

2.1.4 Longevity risk in Nationale-Nederlanden

In order to measure and manage risks in a consistent way, all ING Group business units use the concept of Economic Capital (EC) to determine the risks they are exposed to. Economic Capital quantifies the impact of adverse events on the market value surplus. It is an application of Surplus-at-Risk: the EC determines the maximum risk that can be engaged in, while the market value of the assets is sufficient to cover for the market value of the liabilities with a 99.95% probability (CIRM, 2007). 99.95% is the probability that an AA-rated company like ING will not default within one year.

The actuarial department of the Group Life department of Nationale-Nederlanden carried out several internal analyses in order to determine how to come to the Best Estimate Mortality and Economic Capital figures.

For the determination of the Best Estimate Mortality estimate, the CRC methodology is applied. This decision was made after a close examination of the three models presented in section 2.1.2. The drawback of the CBS method is that the underlying distribution is unknown, making it impossible to determine the lower bounds of the mortality rates that are required for the EC calculations. The Lee-Carter model results in risk margins for the estimates of the future rates that are considered unrealistically low. The Best Estimate of the future mortality rates found by applying the model; the Best Estimate for the insureds can be determined by multiplying these rates by a certain factor for experience mortality.

At Nationale-Nederlanden, the longevity risk is subdivided into four subrisks, resulting in four separate EC-calculations (c.f. Van Broekhoven, 2002).

1. Volatility risk is the risk that the actual number of deaths deviates from the BEM as a result of the randomness of the process, assuming that the forecasted estimates are true.

   The EC required for it is determined by the 99,95% upper bound of the Capital at Risk resulting from the Best Estimate Mortality rates, where the death or survival of each insured is assumed to be a binomial process.

2. Calamity risk is the risk of a one-time claim event of extreme proportions that can be attributed to a certain event. After the calamity, mortality rates will return to their previous levels. Examples of calamities are epidemics, natural catastrophes (earthquakes, meteorite strikes) and terrorist attacks.

   The EC for this risk is determined by the difference in the liability under the assumption of double mortality rates in the upcoming year, and the liability under the Best Estimate. This outcome will be negative most of the time, as less pensions have to be paid out when many people die. A negative EC for calamity risk is set to zero. The number can be positive when the widow pensions are high compared to the premiums collected.

3. Trend uncertainty is the uncertainty in the estimation of the trend in future mortality rates of the entire population. The estimation of the trend is based on past trends, but medical developments could result in a structural break. An example could be the discovery of a cure for a currently fatal and frequent disease like cancer.

   The effect of a change in the general mortality trend is quantified by supplementing the historical population mortality data with a worst case scenario before the model is calibrated. The worst case scenario that is added to the data is obtained by the 83% lower bound of the initial estimates. The 83% confidence interval is used because this is the probability of a currently AA-rated company to survive the upcoming 17 years (the duration of the pension liabilities, where this risk relates to).The estimates are scaled back to match the insurer's liability by multiplying them with the factor for experienced mortality. The resulting liability is lowered by the liability under the Best Estimate to come to the EC for trend uncertainty.

4. Level uncertainty is the uncertainty around the experienced mortality factor, hence the effect of a change in the experienced mortality while population mortality rates are unchanged.

   The EC for level uncertainty is determined by the difference between the BE liability and the liability resulting from the population Best Estimate Mortality multiplied with the lower bound of the experienced mortality factor. This lower bound is also determined with a 83% confidence level.

In the last quarter of 2007, the Economic Capital for life risk amounted to €819mln, against €657mln in the third quarter. An overview of the longevity sub risks can be found in table 2.1. Clearly, the largest source of risk is trend uncertainty; 88% of the Economic Capital stems from this source. The increase in the EC for trend uncertainty as compared to the third quarter of 2007 is due to updated life risk scenarios. The change in calamity risk is a result of a different way of measuring it.

Table 2.1: Longevity subrisks in Nationale-Nederlanden Group Life

	EC Q4 2007	EC Q3 2007
Volatility	25	20
Calamity	0	20
Trend uncertainty	721	544
Level uncertainty	73	73
Total	819	657

The table shows the development of the separate components of longevity risk. The numbers are in €mln.

2.2 Longevity derivatives

In order to cover (part of) the longevity risk, pension providers can engage in longevity derivatives trading. A longevity derivative is a financial derivative of which the payoff is linked to a mortality rate or mortality index. This section discusses these derivatives. After a description of the current longevity derivatives market, the products that can be constructed in theory are examined. These products will be set in the light of hedging the longevity risk of a pension provider.

2.2.1 The market for longevity derivatives

According to Coughlan (2007a), the worldwide exposure to longevity is over €15 trillion. The possibilities for a market in longevity derivatives are therefore numerous. Blake et al (2006) list the possible stakeholders for this market as follows:

- providers of pensions and life insurance, who have a particular exposure to longevity risk and want to hedge it;
- investors like investment banks and hedge funds, as the correlation with other investment risk is very low.
- speculators and arbitrageurs, they are essential for the success of traded futures and options markets because they create liquidity;
- governments, on the one hand because they are exposed to longevity risk via the pay-as-you-go state pension system, and on the other hand because the existence of a longevity derivatives market can reduce the probability that large companies are bankrupted by their pension funds, resulting in more stability in the economy as a whole.

The market is net short in longevity: there are more economic agents who want to sell their longevity risk than agents willing to buy it. Investors will therefore require compensation: the mortality rates in the derivatives should be settled below the expected mortality rate. (Hua, 2007)

2.2.2 Longevity derivatives in the financial markets

Starting from 2003, several mortality catastrophe bonds have been issued on the capital markets. These bonds are linked to the opposite of longevity (seen from the point of view of the counterparty who covers its risk, which in this case is the issuer): the payout to the bondholders decreases when the number of people deceased in a certain group turns out to be higher than expected. The reinsurers, who issue them, use these bonds to cover part of their short term calamity risk. The mortality catastrophe bonds issued vary in their maturity (3 to 5 years), credit spreads, rating and underlying population. An accurate summary of these bonds can be found in Coughlan et al (2007a). As these bonds do not pay out when mortality is lower than expected, they are not useful for a pension provider like the Group Life department of Nationale-Nederlanden. In 2004, a longevity bond was presented that actually could be useful for this purpose: the EIB/BNP longevity bond. This derivative was issued by the European Investment Bank (EIB), with BNP Paribas as structurer and marketer and Partner Re as longevity risk reinsurer. The face value of the bond was 540mln, its maturity 25 years. The bond was an annuity bond with floating coupons that were linked to the survivor index of British and Welsh males aged 65 in 2002. The coupons would be higher when the realized mortality rates turned out to be lower than expected. The initial coupon was set at 50mln; at each coupon payment date, the bond pays a coupon of 50mln times $S(t)$. $S(t)$ stands for the survivor index, which equals 1 in 2002 (the basis of the index). Each year after that, the new index could be obtained by multiplying last year's index by 1 minus the mortality rate for the cohort in that year. When $m(t,\tau)$ is the central mortality rate in year t for the group of people aged τ, the values for the survivor index are obtained as follows:

$$S(t,\tau) = S(t,\tau-1)(1-m(t,\tau)) \tag{2.7}$$

The issue price of the contract was specified using a series of $S(t)$s as forecasted by the Government Actuary's Department in 2002. These cashflows were discounted at LIBOR minus 35bps. The bond was constructed such that bond holders would only run the credit risk of EIB, and because the EIB curve typically is about 15 basis points below the LIBOR curve, the investors in the bond actually would pay 20 basis points to hedge their longevity risk.

The EIB/BNP longevity bond was withdrawn from the market after one year because it was not fully subscribed. This lack of interest can be explained in several ways (Blake et al, 2006; Neville and Ho, 2006; Hua, 2007):

- The coupon payments were based on males in England and Wales aged 65 in 2002, while the portfolio of pension and annuity providers generally contains both males and females, people of various ages and inhabitants of geographic areas other than Wales. This basis risk requires a compensation, and the compensation offered might be considered insufficient.
- The horizon of 25 years might be too short for the risk insurers and pension providers want to hedge.
- The liquidity of the bond might be assessed to be too low, i.e. investors might perceive that the probability that they will be unable to sell the bond at the desired moment is too large.
- Because of the long duration of the bond, the degree of model and parameter risk is quite high, which might make potential investors and issuers uncomfortable.
- Insurance firms might not have mandates yet that allow them to hedge longevity risk.

Besides the EIB/BNP longevity bond, no derivatives have been publicly introduced that could cover for the longevity risk a pension provider faces.

Loeys et al (2007) however conclude that the possibilities for a longevity market in the Netherlands are numerous. Reasons are the size of the pension funds and the market's awareness of longevity risk.

A large obstacle in the development of a liquid longevity derivatives market is the difference in needs on both sides of the market: Investors want a liquid product, and hence standardization, while pension and annuity providers want a customized hedge that closely follows their longevity risk.

In an attempt to support the standardization of the longevity market, JPMorgan launched its LifeMetrics platform in March 2007. LifeMetrics consists of an index with historical and current mortality rates for the US, the UK, Germany and the Netherlands (differentiated by age and gender), along with supporting technical documents and software. The index could be used as an objective index on which mortality estimates can be based (e.g., using one of the models presented in section 2.1.2.2) to come to a generally accepted price of a specific longevity derivative.

In March 2008, the first publicly announced large deal on the LifeMetrics Index was set up. In this deal, JP Morgan is the counterparty to Lucida, which bought the entire longevity risk of the Bank of Ireland. The longevity risk involved is €40mln each year, the maturity is 10 years. If the LifeMetrics Index indicates that people live longer than forecasted by the UK Statistics Office, Lucida will receive payments from JPMorgan and vice versa. It is hence an over-the-counter longevity swap. (Life & Pensions, 2008)

2.2.3 Theoretically feasible longevity derivatives

In theory, the variety in longevity derivatives is practically endless. Many features from financial bonds, futures, options, forwards and swaps could be translated to the longevity market. Several prospective longevity derivatives are described by Blake et al (2006), Blake (2008), Coughlan et al (2007b) and Dowd et al (2005). We discuss them in turn.

Longevity bonds
Longevity bonds could vary in structure and maturity.

- Structure
  - Principal-at-risk longevity bonds: the part of principal that is repaid at the maturity depends on a mortality event (e.g., the mortality catastrophe bonds discussed above. A specific kind of this bond could be a zero-coupon bond, which could be a building block for tailor-made positions.
  - Coupon based longevity bonds: annuity bonds whose coupon payments depend on mortality rates (e.g., the EIB/BNP longevity bond discussed in section 2.2.2).
- Maturity
  - Fixed maturity: predetermined number of years.
  - Stochastic maturity: the bond matures when the longest lived member of the reference cohort dies.

Longevity futures
The largest barrier for the emergence of a longevity futures market is that it needs a suitable underlying instrument (Blake et al, 2006). The spot market of this underlying has to be large, liquid, volatile, and not heavily concentrated on either buy or sell side, in order to make it attractive for hedging and speculation activities. For comparison, consider the US CPI futures contract. This future was listed in 1985

and delisted after two years because the number of trades was extremely low. According to Srinivasan (2004), this failure was a result of the combination of two factors: the underlying index was not published frequently and there were no stable pricing relationships with other instruments. The analogy with longevity futures contracts is evident. However, weather futures contracts did survive, despite the fact that the two factors supposed to cause the failure of the US CPI futures contract were also present.

Blake et al (2006) suggest the use of annuities or longevity bonds as underlying instruments. The UK market for annuities is fairly active, but it is also illiquid and inefficient. Prices are not updated frequently, but the pricing relationship with the yield curve of government bonds makes their moves predictable. Another problem is that insurers are not obliged to reveal their prices, and it may not be in their best interest to do so. An alternative would be the use of longevity bonds. There is no liquid market for longevity bonds yet, but when it will exist, it will probably fluctuate with the interest rate, making it a suitable underlying.

Longevity swaps
A longevity swap involves the exchange of a series of fixed (preset) payments for a series of random longevity-dependent payments. For each payments at t = 1, 2, , T, the fixed amount $K(t)$ is swapped for a random amount $S(t)$ (which is linked to the realized mortality rates: the fraction of people currently aged x, who are alive at some time t): firm A pays firm B an amount $K(t) - S(t)$ when $K(t) > S(t)$ and B pays A an amount $S(t) - K(t)$ if $S(t) > K(t)$. Hence, the amounts firm A has to pay are higher when more people live longer.

Blake et al (2006) report that several insurance companies already have entered into longevity swaps on an over-the-counter basis. Supposedly, they would have declared that the fixed leg payments are linked to a published mortality projection, and the floating leg is linked to the realized mortality of the counterparty. The parties consist of life insurance companies and investment banks.

Coughlan et al (2007b) propose a kind of longevity forward that could function as a building block in the construction of longevity swaps. In these q-forwards, a fixed cashflow is exchanged for a floating one at the maturity of the contract. JP Morgan presents itself as a counterparty in these trades, and claims that a few over-the-counter transactions have already been realized.

Besides ‘vanilla’ longevity swaps, many varieties are possible. Examples are swaps with two floating rate sides, cross-currency longevity swaps and mortality swaps in which one or more floating payments depend on a non-longevity random variable (e.g. interest rate, stock index).

Longevity options
Many financial options could be translated into longevity counterparts.

Longevity caps and floors can be based on a mortality index $S(t, x)$ (the fraction of people currently aged x, who are alive at some time t). When $s_c(t)$ is the cap rate for exercise date t, a caplet pays a certain amount times $max[S(t, x) - s_c(t), 0]$ at time t. Similarly, a floorlet would pay a certain amount times $max[s_f(t) - S(t, x), 0]$. Survivor caplets and floorlets can be packaged in longevity caps and floors.

Longevity swaptions could be written on longevity swaps of a specific type and maturity. The swaption might be American, European, or Bermudan. If the underlying swap is a vanilla longevity swap, the swaption can be a payer longevity swaption (the holder has the right to become the fixed rate payer) or a receiver longevity swaption (the holder has the right to become the fixed rate receiver).

2.3 Hedging longevity risk with longevity derivatives

In hedging the longevity risk of the Group Life department of Nationale-Nederlanden, the amount of current and future Economic Capital required for the longevity risk is central (section 4.1 explains how they can be combined in one measure, called the Market Value Margin). The main interest is whether a difference exists in the employed ECs with and without the hedge. This difference then can be weighted against the cost of the hedge. The selection of the most appropriate hedging instrument hence actually is a trade-off between the costs and the fit of the hedge.

A customized longevity hedge would be tailored to reflect the longevity risk actually faced in the organization, resulting in a 100% effective hedge. It requires over-the-counter products like swaps or forwards, preferably with a stochastic maturity to cover for the full longevity risk in the liabilities. Disadvantages of these products are their high costs, large counterparty credit exposure and high costs. In exchange for that, there is no residual basis risk and the hedge requires minimal monitoring.

A standardized hedge considers standardized products based on the longevity experience of the national population, with a fixed maturity of 5, 10 or 20 years. Its costs are lower, and the products potentially are more liquid. Also, the shorter maturity lowers counterparty credit exposure. But the hedge is not perfect; there is basis risk and roll-over risk, as periodical rebalancing is necessary. (Coughlan, 2007; Coughlan et al, 2007a)

Besides this trade-off, also the practical possibilities should be taken into account. The emergence of a market in longevity futures is highly unlikely because of the lack of a suitable liquid underlying product. Also, longevity options are a long way ahead, as the options have to be written on a market that does not exist as of yet.

This leaves longevity bonds and longevity swaps (the q-forwards mentioned in section 2.2.3 essentially are one-period longevity swaps). Longevity swaps have certain advantages over longevity bonds (Dowd et al, 2006). Their transaction costs are lower and it is easier to equal out a position. They are more flexible and could be customized to suit individual circumstances. Moreover, they do not require the existence of a liquid market, just the willingness of counterparties to exploit their comparative advantages or trade views on the development of mortality rates over time. Longevity bonds do need a liquid market. This market does not exist yet.

As such, longevity swaps appear to be the most relevant derivative to consider for hedging longevity risk. The degree of customization will be the outcome of the trade-off between risk and costs described above.

3 Longevity swaps

After the short introduction in section 2.2.3, longevity swaps are discussed more extensively in this chapter. The chapter starts off with a detailed explanation of the mechanics of the swap, which are also compared to those of interest rate and credit default swaps. In section 3.2, two pricing frameworks are developed, consistent with two pricing methods that are found in the literature and deemed relevant. Both frameworks are translated into a worked example, enabling a sensitivity analysis of the two prices in section 3.3.

3.1 Longevity swap mechanics

The basic building block of a vanilla longevity swap is a time t random payment that depends on a mortality index. This index might consist of the mortality rates of the entire population, as well as the mortality rates among a certain group of the population. At the initiation of the contract, two parties agree to swap at time t a fixed amount $K(t)$ (linked to the currently expected future level of the mortality index) for a random amount $S(t)$ (linked to the realized level of the mortality index in the future). The payment at time t is the difference between $K(t)$ and $S(t)$, which is paid by the party whose side is higher at that moment. The total swap consists of a series of these building blocks.

As an illustration, consider a pension fund with 10,000 policyholders from a homogeneous group of 65-year-old individuals. The pension consists of a payment of €1 each year to each policyholder alive. $S(0)$ hence is 10,000 and $S(t)$ gradually falls over time. The pension fund could enter the fixed side of the longevity swap contract. At each payment date, $K(t)$ is paid, which is an amount of euros equal to the number of people that were expected (at the beginning of the contract) to be alive at that time, no matter whether the actual number is higher or lower. The floating leg counterparty makes up for the difference when the mortality rate is lower than expected, and the difference has to be paid to this counterparty when the estimates appear too high. Hence, the pension fund does not run the risk that its policy holders live longer than expected, but it will also not profit when mortality rates appear to be higher.

There are many similarities between the mechanics of longevity swaps and those of interest rate swaps and credit default swaps.[2] All three provide their buyers with insurance against movements in the market, by fixing the effective (net) cashflows. Besides, they all involve periodical payments and all add a premium to the payments of the fixed leg counterparty, in order to make the present value of the contract zero at its initiation.

A notable difference between longevity swaps and interest rate swaps is found in the fixed payments: while they are constant for the interest rate swap, they decline with the anticipated number of people alive in the reference group for the longevity swap. Furthermore, the floating leg of the interest rate swap is tied to a market interest rate, while the floating leg of the longevity swap depends on the realized value of a survivor index. The most important difference between interest rate and longevity

[2] In an interest rate swap, two counterparties pay each other the difference between a specific interest rate that was expected at the beginning of the contract (a fixed interest rate) and its realized value. Interest rate swaps are traded on public exchanges. A credit default swap is linked to a bond. Its buyer receives the difference between the market and the face value of the bond when the issuer of the bond defaults before the maturity of the contract. In return for this, he pays a periodical fee, either until the contracts maturity or until the default. Credit default swaps are only traded over-the-counter. More information on interest rate swaps and credit default swaps can be found in Hull (2006).

swaps is that the market of the first is complete, enabling valuation based on the yield curve. Because the yield curve incorporates both the expected future interest rates and the amount of risk embedded in these estimates, forward rates can be used to determine the floating payments. The fixed payments of the interest rate swap are subsequently set such that the present values of the two legs are equal. The market for longevity swaps is incomplete and hence requires a different pricing procedure.

A comparison of the longevity swap with credit default swaps shows that both are directly linked to survival probabilities: of a company or of a reference group of people. Unlike the fixed payments in the longevity swap, the fixed leg of the credit default swap is constant over time. Another difference is that these fixed payments are not paid over the entire maturity of the contract, as the payments stop in case the bond issuer defaults. Regarding the floating side of the contract, the longevity swap makes periodical payments, while the floating counterparty in the credit default swap only makes one payment in case of default and none otherwise. A similarity is that both swaps have an incomplete market. In the valuation of credit default swaps, pricing equations are used along with market information on quotes of new swaps to bootstrap the survival probabilities of companies. With this information, the price of the credit default swap in time can be determined. Market information is not available for longevity swaps, hence this method can also not be applied to the valuation of longevity swaps. The procedures that are available are discussed in the next section.

3.2 Longevity swap pricing

The (prospective) market for longevity derivatives contains more short than long investors: the amount of longevity risk supplied is larger than the amount of longevity risk demanded (Loeys et al, 2007). Investors willing to accept longevity risk hence require compensation for it. In the determination of the terms of the swap, a premium is therefore added to the fixed side. This premium, the 'price' of the swap, is tied proportionally to the value of the fixed payments. At the initiation of the swap, the premium (π) is set such that both sides of the swap are equal (the initial value of the contract is zero). The $K(t)$ employed in the introduction of this chapter is equal to $(1+\pi)H(t)$.

$$\begin{aligned} PV(S(t)) &= PV[(1+\pi)H(t)] \\ &= (1+\pi)PV[H(t)] \\ \Leftrightarrow \pi &= \frac{PV(S(t)}{PV(H(t)} - 1 \end{aligned} \tag{3.1}$$

The value of the premium π can be determined by two different procedures:

- The Wang transform method uses the insurance market as point of departure. It converts the currently expected floating payments into their risk neutral equivalents using a specific market price of risk (from the insurance market), and determines the premium by solving equation 3.1. This procedure was presented by Dowd et al (2005) as a pricing method used in the pricing of several over-the-counter longevity swaps in practice.

- The Sharpe ratio method is based on parallels with the capital market. This method assigns a Sharpe ratio and determines the forward longevity premium by assuming that the total of $(1+\pi)PV[H(t)]$ equals the present value of the floating leg. This method is explained by Loeys et al (2007) as the procedure JP Morgan intends to use in the pricing of its q-forwards.

Both procedures are explained below. They are illustrated by a 'typical' longevity swap with a maturity of 20 years, on Dutch males aged 65 in 2008. The notional of the swap is set at €100,000. In

Figure 3.1: Risk-free and market interest rates and survival rates used in the swap example

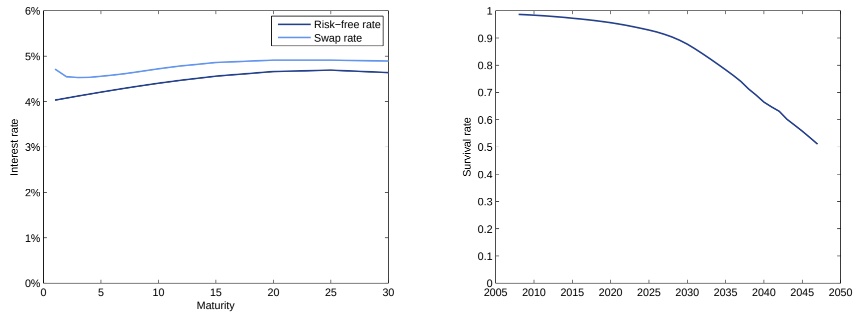

This figure illustrates the data used in the swap example. The left figure shows the risk-free interest rate structure (Euro AAA government bonds) and the market interest rate structure (Euro vs. Euribor swaps) as of January 1, 2008. The right figure shows the CBS survival probabilities for males aged 65 in 2008.

discounting the cashflows back to the beginning of 2008, we use the term structures of January 1, 2008 for Euro AAA government bonds (risk-free interest rate) and Euro versus Euribor interest rate swap rates (market interest rate). The interest rates are obtained from Thomson Datastream. Following Hull (2006), linear interpolation is used to determine interest rates of maturities that are not quoted, and the curves are assumed to be horizontal before the quote with the shortest and after the quote with the longest maturity.

The contract is built on the CBS survival rates for Dutch males aged 65 ($p(65)$) in 2008 (the probability of survival for this group of people from 2008 until 2047 when alive at the beginning of each year). All payments take place at the end of the year. The relevant interest rates and survival rates are shown in figure 3.1.

3.2.1 Pricing longevity swaps using the Wang transform

The first procedure to determine the price of a longevity swap determines the present value of the fixed and floating payments in order to solve for the premium. The present value of the fixed payments equals the discounted expected payments plus this premium. The present value of the floating payments is determined by the discounted payments that are expected under the risk-neutral measure. This procedure was proposed by Dowd et al (2006).

The current best estimates of the mortality rates are converted to their risk-neutral counterparts by the Wang transform. The Wang transform (Wang, 2002) essentially is an application of the Capital Asset Pricing Model to not normally distributed variables. The approach determines the standardized form of a series of probabilities (z-scores) as if they were normally distributed. These z-scores are uniformly shifted by an amount λ, which stands for the market price of risk. The shifted z-scores are then transformed back, again using the standard normal distribution.

In mathematical form, the Wang distortion operator (g_λ) is defined as [3]

$$g_\lambda(u) = \Phi(\Phi^{-1}(u) + \lambda) \tag{3.2}$$

where $\Phi(x)$ is the cumulative standard normal distribution, u is a probability between 0 and 1, and λ is the market price of risk. This implies that the probabilities (u) can be transformed in the following way:

$$u^\star = g_\lambda(u) \tag{3.3}$$

In the setting of the longevity swap, the distribution of the best estimates for the survival rates (p) is filled in as u:

$$F^\star(p) = \Phi(\Phi^{-1}(F(p)) + \lambda) \tag{3.4}$$

The Wang transform adds the risk of longevity to the survival rates, which makes these rates risk neutral as no additional premium is required for the longevity risk when these rates are used. The risk neutral rates are the basis of the present value determination of the floating payments ($PV[S(t)]$): they are multiplied with the notional of the contract, and discounted with the risk-free interest rate structure.

The present value of the fixed payments ($PV[H(t)]$) is determined by discounting the best estimate mortality rates at the initiation of the swap against the prevailing market interest rate curve. The premium (the price of the swap) can then be obtained by filling in equation 3.1.

The main drawback of this pricing procedure is that the Wang transformation requires a value for the market price of risk parameter (λ). Lin and Cox (2005) derive λ from the market annuity price (the price currently paid for an immediately starting annuity on a 65-year-old male), assuming that the uncertainty in the mortality table is taken into account in the determination of this price. The market price of risk would then be determined by the difference between the market annuity price and the purchase price of an annuity based on the 'true' (actually expected) survival rates. This procedure is complicated by two factors:

- market prices do not only contain the best estimate of future mortality and the market price of longevity risk, but also the market price of interest rate risk and a margin for costs;
- substantial differences exist in prices for annuities between insurers in the Netherlands[4] ; this calls for a decision on which insurer knows the 'right' risk premium.

In addition to the problem of the determination of the market price of risk, there also is criticism on the concept of the Wang transform as a means to determine risk-neutral insurance liabilities. Bauer and Ruß (2006) state that the Wang transform is just an arbitrary possibility to distort the distribution of survival probabilities. Its results are consistent with prospects on longevity: the difference between the estimated and the transformed survival rates are higher at old ages than at younger ages (as the estimates for nearby dates are probably more accurate and the uncertainty about mortality is higher for higher ages). Further rationale behind the procedure is not presented. Wang's (2002) main argument for the usefulness of his method is that its results are consistent with the CAPM results for normally

[3]Because the investors on the long side of the market demand risk compensation, the market price of risk is added to the floating survival rates. This implies that the payments that are paid to the fixed counterparty are lower in case the survival rates appear to be higher than expected.

[4]See for example a comparative analysis of annuity prices ('direct ingaande lijfrente') for a 65-year old male at http://www.independer.nl.

distributed variables and with the Black-Scholes prices for Geometric Brownian Motions with constant coefficients. Pelsser (2008) however determines that the price arising from the Wang transform is only consistent with arbitrage-free pricing when the coefficients of the process are functions of time exclusively. This would imply that the procedure should not be used for the determination of the market price of mortality risk, as the process is also linked to age.

Example
The valuation of the fixed leg payments is straightforward: the CBS Best Estimates of the survival rates are discounted to 2008 by the Euribor swap curve of January 1, 2008.

The floating leg payments are determined using the Wang transform, hence a value for the market price of longevity risk is required. The procedure of Lin and Cox (2005) is modified in order to exclude cost margins and interest rate risk. The modification encompasses that both the market price and the price implied by the currently estimated survival rates (the 'true' price) are determined by the mortality tables insurers generally use to determine annuity prices. The Collectief 2003 mortality table was used to determine the net market price of an annuity for a 65-year-old male. Collectief 2003 is the standard mortality table used by Dutch insurers to determine the purchase prices of group contracts (issued by the Verbond van Verzekeraars). The obtained survival rates were used to find the net purchase price (present value) of a payment of €1 at the end of each future year in which the policyholder is alive:

$$a_{65} = \sum_{t=66}^{120-66} \frac{(t)p(x)}{(1+r)^t} \tag{3.5}$$

where

$$(t)p(x) = \prod_{s=0}^{t-1}(1 - q(x+s))$$

The 'true' net annuity price is determined by the mortality rates of the generation table that was recently developed by the actuaries of the Group Life department of Nationale-Nederlanden (based on CRC estimates, the own mortality experience, and cost margins). The survival rates of a male who currently is 65 of age are obtained and transformed with the Wang transform, where an arbitrary temporary λ is used. This λ is added to the survival probabilities, as the risk is at the side of the seller of the contract (the insurer runs the risk that the insured person lives longer than expected). The resulting rates are discounted against the risk-free curve to come to the amounts that sum to the present value of the expected annuity payments at 1-1-2008. The market price of risk parameter is solved by equaling the two purchase prices.

For a male aged 65, the market price of risk parameter calculated in this way appears to be -0.18295. The negativity of the number is not surprising, because the mortality rates of the Collectief 2003 table underestimate the longevity as now determined in the internal mortality table: the insurer knows he will lose on longevity. So whereas it is generally known that mortality has improved, the annuity prices are not updated because of competition reasons. Figure 3.2 shows the evolution of the obtained market price of risk for people of different ages. Because a longer maturity (lower age) implies a higher λ as the uncertainty of future mortality rates is larger, the market price of risk parameter should increase continuously with maturity. The figure shows that this is indeed true, but also that the premium is negative over the entire range of ages. It is clear that the method of Lin and Cox (2005) does not work in determining the market price of the longevity risk.

Figure 3.2: Market price of risk deducted from theoretical annuity prices

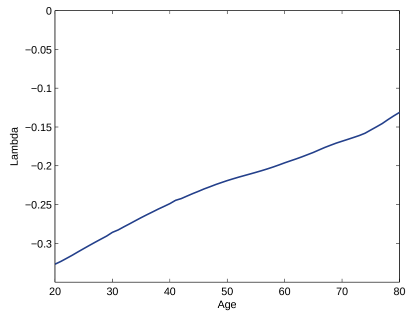

The evolution of the market price of risk parameter for Dutch males of different ages, as deducted from the theoretical market and theoretical 'true' annuity prices (modified Lin and Cox (2005) procedure).

As an alternative method to determine the market price of risk, we use the internal price of risk Nationale-Nederlanden assigns to its longevity products in the Economic Capital. The amount of capital required to cover the present values of the Economic Capital (the Market Value Margin, MVM) as a percentage of the provision held for the product is assumed to equal the market price of risk. The concepts of Economic Capital and Market Value Margin will be elaborated for a specific contract in the chapter 4, but as an indication, the overall λ is assumed to equal the ratio of the overall MVM for longevity risk to the total provision held for longevity liabilities. According to internal ING spreadsheets, the MVM for longevity risk in the last quarter of 2007 was €199.32mln. The book value of the provision for the total longevity liabilities outstanding was €22,865mln, implying a λ of 0.9%.

0.9% appears to be very low for a market price of longevity risk. But because of the lack of other acceptable indications, we use this 0.9% to transform the Best Estimate Mortality probabilities into their risk-neutral counterparts. The risk-free interest rate curve is used in the discounting.

As all cashflows and their present values are determined now, the premium λ can be found. It is obtained using equation 3.1.

$$\pi = \frac{PV[S(t)]}{PV[H(t)]} - 1 = \frac{1,620,523}{1,226,942} - 1 = 0.02737$$

The price of the swap is hence 2.737%: the mortality table the fixed payments are based on are multiplied by 1-0.02737 to come to the fixed leg of the contract.

3.2.2 Pricing longevity swaps using the Sharpe ratio

An alternative method, that does not use a transformation of survival probabilities nor a value for the market price of risk, is proposed by Loeys et al (2007). The procedure determines the payments of the fixed counterparty ($(1+\pi)H(t)$) based on analogies with other investment classes, and assumes the present value of these payments equals the present value of the floating leg at the initiation of the contract. The basis of the procedure is the Sharpe ratio, which in theory is equivalent to the market price of risk parameter λ (Wang, 2002). The procedure described in the previous section requires assessing the market price of risk, while the method of Loeys et al (2007) assumes a certain value for the Sharpe ratio.

In liquid markets, the Sharpe ratio provides a benchmark for the determination of risk premia. In general, the required risk premium will be higher when the underlying risk (volatility) is higher, when the correlation with other asset classes is higher, or when the liquidity of the asset is lower. Longevity risk is uncorrelated with other risks, therefore the required Sharpe ratio should be below the averages in riskier markets such as equities, but the liquidity of the instruments is fairly low. Sharpe ratios of different liquid markets are listed in table 3.1. Loeys et al (2005) state that an annualized Sharpe ratio of around 0.25 can be expected.

The (absolute) annualized expected return of a forward ($E(R)$) is equal to the difference between the expected mortality rate (q_e) and the forward mortality rate (q_f), divided by the number of years the return is realized on (t):

$$E(R) = (q_e - q_f)/t \tag{3.6}$$

The (absolute) annualized risk of a forward equals the standard deviation of year-on-year percentage changes in the mortality rate (σ): the historical rate of changes in mortality rates (σ_q; the relative rate) times the expected mortality rate (to result in an absolute rate):

$$\sigma = \sigma_q \cdot q_e \tag{3.7}$$

As the annualized Sharpe ratio (SR) is assumed to be 0.25, q_f can be obtained straightforwardly.

Table 3.1: Sharpe ratios for various asset classes

Historical Sharpe Ratios (1976-2006)	
Stock market US	0.38
Stock market Pacific	0.28
Stock market Europe	0.39
Commodities	0.18
LT US Treasury	0.25
5y US Treasury	0.33
Bond market index	0.37

The table depicts historical annualized Sharpe ratios for several asset classes.
Source: http://www.sharpeinvesting.com/2007/06/sharpe-ratio.html

$$SR = \frac{E(R)}{\sigma} = \frac{(q_e - q_f)/t}{\sigma_q \cdot q_e}$$
$$\Leftrightarrow q_f = (1 - SR \cdot t \cdot \sigma_q) \cdot q_e$$

The forward discount (negative premium) implied equals the difference between the forward mortality and the expected future mortality as percentage of the expected future mortality. This premium equals the premium π in equation 3.1; which equalizes the floating and the fixed side of the contract. Because this premium only related to the cashflow at the maturity date, the π is different for each time. A subscript t is therefore attached to the symbol.

$$\begin{aligned} q_f - q_e &= (1 - 1 - SR \cdot t \cdot \sigma_q) \cdot q_e \\ &= -(SR \cdot t \cdot \sigma_q) \cdot q_e \\ \Leftrightarrow \pi_t &= (q_f - q_e)/q_e \\ &= -(SR \cdot t \cdot \sigma_q) \end{aligned} \tag{3.8}$$

Our pricing framework uses the survival rates p instead of the mortality rates q. As the risk of the market is equivalent (because p = 1 - q), the Sharpe ratio will be the same. The risk is at the upside of the survival rates, hence q_f is above q_e. The resulting premium (again, π in equation 3.1, but only for the cashflow at maturity t) is:

$$\begin{aligned} (1 - p_f) - (1 - p_e) = p_f - p_e &= (SR \cdot t \cdot \sigma_q) \cdot (1 - p_e) \\ \Leftrightarrow \pi_t = \frac{p_f - p_e}{1 - p_e} &= SR \cdot t \cdot \sigma_q \end{aligned} \tag{3.9}$$

The premium is set separately for each maturity; the swap is built from a series of longevity forwards. The fixed payments of the longevity forwards are determined using the survival rates as determined by the currently expected mortality rates times (1 - π_t). The present value of this leg is assumed to equal the present value of the future floating payments.

For clarification, we draw a parallel between these longevity forwards and forwards on a financial asset. The price of a forward on a financial asset is determined by its current price, reduced with the present value of benefits resulting from owing the asset until the maturity of the contracts (e.g. dividends) and raised by the present value of costs stemming from owning the asset until the maturity of the contract (e.g. storage costs). At the maturity date, the then prevailing spot price of the asset is compared to the price at the moment the contract was initiated, exclusive the benefits and inclusive the costs. When the prevailing price is higher, the seller of the contract suffers a loss equal to the difference. When the prevailing price is lower, the buyer of the forward contract has to come up with the initially agreed price.

For the longevity forward (as building block of the longevity swap), the forward price equals the currently expected future survival rate for the cohort and the mortality under concern, proportionally raised with the premium. The payment at the maturity is determined by the difference between the realized survival rate and (1 + π_t) times the expected survival rate. When the realized rate is higher than (1 + π_t) times the expectation, the short counterparty pays the difference times the notional amount; the long counterparty pays when the realized rate is lower.

Figure 3.3: Premiums for each future payment of the longevity swap

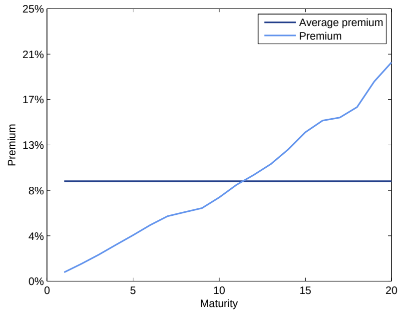

The premiums of a 20 year longevity swap determined using the Sharpe Ratio method and its average premium.

Example
We first simulate a longevity swap with only one future payment (a longevity forward), on the 65-year old cohort of Dutch males, with a maturity of 10 years. According to CBS projections, the survival rate for 75 year olds in 2017 is 0.9635. The volatility of the changes in mortality rates for 75 year olds has been 3.07% of the rate. The forward survival rate is determined as follows:

$$\pi_T = 0.25 \cdot 10 \cdot 0.0307 = 0.07675$$

This method thus results in a risk premium of 7.68%.

A 20 year swap consists of 20 one-period swaps with different maturities. The resulting premiums per period are shown in figure 3.3. The average premium is 9.17%. This average premium can be compared with the (single) premium for the 20-year swap of the Wang transform method.

Figure 3.4 shows that the Sharpe ratio procedure cannot be used to price swaps on 65-year-olds with a maturity longer than 25 years. This is because the historical data for the higher ages is too volatile, hence σ_q is high. Besides, there is no volatility data available for maturities above 34 years, as the JP Morgan LifeMetrics database only contains survival rates for people up to 98 years of age.

Of course, swaps with longer maturities can be applied to younger cohorts. As an example, figure 3.5 shows the path of the premiums of a 40 year swap on Dutch 45-year old males. The concept is clear: the premiums increase substantially over time, along with the uncertainty in mortality rates. We analyze a maturity of 20 years as the base case, because this enables swaps on the older ages as well. Besides, in light of the standardization of the market, the 20 year horizon will probably be more liquid than the 40 year horizon.

The main problems of the Sharpe ratio procedure hence are that it is not possible to price a longevity swap on people above the age of 90 because of reasons of high historical volatility, and that the premiums increase rapidly with maturity. Besides, several aspects of this procedure are disputable. The assumption of a constant Sharpe ratio of 0.25 and the length of the sample period used in determining the volatility are open to question. We come back to this in the sensitivity analysis in the next section.

Figure 3.4: Premium per year for a 40 year longevity swap

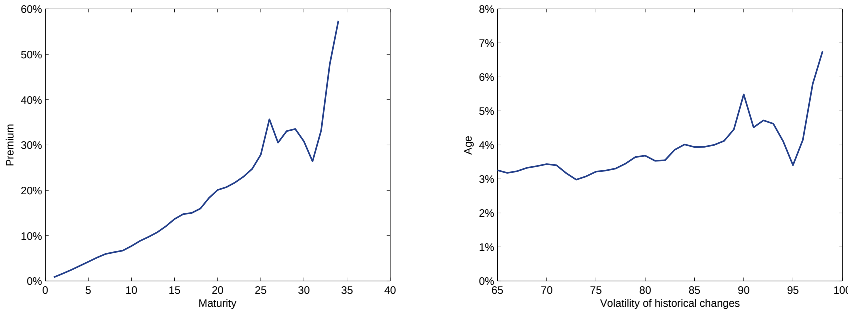

The left figure shows the premiums of the 40 year swap on males who are currently 65 years old, using the Sharpe ratio method. The right figure shows the volatility of the historical changes in mortality rates for the different ages.

Figure 3.5: Premiums for a 40-year swap on 45-years-old Dutch males

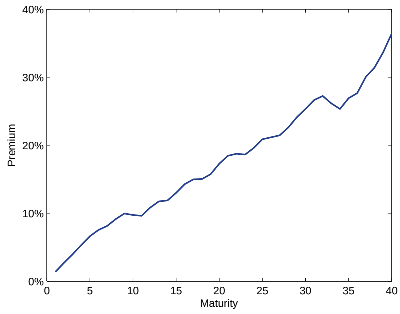

The premiums of a 40 year swap on males who are currently 45 years old, using the Sharpe ratio method.

3.2.3 Comparison

The two presented pricing methods result in different premiums for the longevity swap. The Wang transform method determines one market price of risk parameter, and hence results in a constant premium over the life of the swap. The Sharpe ratio prices a single-payment swap for each year, resulting in an increasing premium over the maturity.

For a longevity swap with a maturity of 20 years on the 65-years-old Dutch males cohort, the premium according to the Wang transform method is 2.737%, the Sharpe ratio method results in an average of 9.17%. The Wang transform method can never reach the average of the Sharpe ratio method; its maximum is 6.3%, but that is only reached when the market price of risk is set unrealistically high (above 100%). The Sharpe ratio method would result in the Wang transform price in case the assumed Sharpe ratio was 0.0746, which is quite far from the current 0.25 as well.[5] The two methods hence are not consistent with each other.

[5]These numbers were found by increasing the market price of risk parameter λ or the Sharpe ratio in the respective pricing equations.

Dowd et al (2005) present the Wang transform method as a method used in the financial world to price over-the-counter longevity swaps. Loeys et al (2007) present the method JP Morgan wants to use for its market in longevity derivatives that it is trying to launch. Both methods hence are relevant for practice, and both will be discussed in the remainder of this thesis.

3.3 Sensitivity analysis

This section will evaluate the change in premium resulting from a change in the parameters of the pricing model. The influence of changes in different parameters on the price of the swap is examined, where the other parameters are held constant. As the two methods employ different parameters, they will be discussed in turn. The base case is determined by the 20 year longevity swap on Dutch males aged 65 at the beginning of 2008.

3.3.1 The longevity swap under the Wang transform method

The premiums resulting from the Wang transform pricing method are evaluated on the influence of changes in the interest rates (both risk free and market), the market price of risk, the maturity of the swap, the mortality estimates (BEM), and the age group (cohort) the swap is based on.

Interest rates
First, the effect of parallel moves of both interest rate curves is determined. Also, the market and risk-free rates are shifted separately. The results are shown in figure 3.6, where the left panel shows the premiums and the right panel the changes in premiums for the movements of the interest rates.

When both curves move by the same amount and in the same direction, the premium remains practically unchanged. As the interest rates are only used for discounting and the differences between the risk free and the market interest rates are not substantial, this is not surprising. There is a slight decrease in the premium when both curves are moved upwards. The reason for this is found by rewriting equation 3.1:

$$
\begin{aligned}
\pi &= \frac{PV[S(t)]}{PV[H(t)]} - 1 \\
&= \frac{\sum_{t=1}^{T} \frac{\Phi(\Phi^{-1}(p(t))+\lambda)}{(1+r_f(t))^t}}{\sum_{t=1}^{T} \frac{p(t)}{(1+r_m(t))^t}} - 1 \\
&= \sum_{t=1}^{T} [(\frac{\Phi(\Phi^{-1}(p(t))+\lambda)}{(1+r_f(t))^t} \cdot \frac{(1+r_m(t))^t}{p(t)} - 1) \cdot \frac{\frac{p(t)}{(1+r_m(t))^t}}{\sum_{s=1}^{t} \frac{p(s)}{(1+r_m(s))^s}}] \\
&= \sum_{t=1}^{T} [(\frac{\Phi(\Phi^{-1}(p(t))+\lambda)}{p(t)} \cdot \frac{(1+r_m(t))^t}{(1+r_f(t))^t} - 1) \cdot \frac{\frac{p(t)}{(1+r_m(t))^t}}{\sum_{s=1}^{t} \frac{p(s)}{(1+r_m(s))^s}}] \qquad (3.10)
\end{aligned}
$$

The last factor in this equation is a scaling factor, which makes up for the bias caused by the necessity to link each probability to its own interest rate (this ratio makes the ratios proportional to the denominator of the initial premium equation, equation 3.1).

The premium is hence subdivided into three parts:

1. the ratio of the Wang transformed to the initial survival probabilities;
2. the ratio of the discount factors;
3. the scaling factor.

The market interest rate curve is above the risk free interest rate curve. Therefore, when both curves are shifted up by the same number of basispoints, ratio (2) decreases. Also, the scaling factor will decrease. As the rest of the equation is unchanged, the premium will decrease as well.

The effects of a change in only the market interest rate or only the risk free interest rate are opposite. This can also be explained by equation 3.1. When only the market interest rate rises, the ratio increases (as the present value of the fixed side payments decreases while the floating side present value stays the same) and vice versa for the risk free rate. One could also reason this: an increase in the market interest rate indicates a larger compensation per unit of risk, hence the untransformed survival probabilities are discounted by a higher rate, resulting in a lower denominator in equation 3.1 and therefore a higher premium. An increase in only the risk free interest rate has exactly the opposite result. It suggests that the trustworthiness of the government decreases, and that of the companies becomes relatively better.

These effects are not linear, because the effect of a shift in the interest rate increases with the shift. The reason for this is that the difference between $(1 + 0.04)^3$ and $(1 + 0.045)^3$ is larger than the difference between $(1 + 0.03)^3$ and $(1 + 0.035)^3$, and hence has a larger effect on the rewritten ratio described above (equation 3.10).

The curvature of the premium changes as result of risk free rate shifts is larger, because the floating side payments are higher than the fixed payments (as a result of the Wang transform). The ratio in equation 3.1 hence is more strongly affected by changes in the risk free rate. This can also be seen from equation 3.10: an increase in the market interest curve results in a higher discount rate ratio, but this is partially offset by a decrease in the scaling factor.

Figure 3.6: Results of a parallel move in interest rate curves

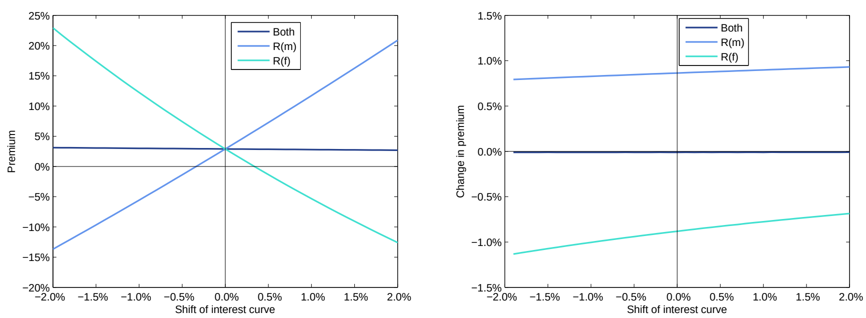

The figure shows the premiums for the longevity swap when interest rates change. Besides a parallel move in both curves, also the effect of only a change in the risk-free or only a change in the market interest rate are considered.

Figure 3.7: Premium level for different market prices of risk

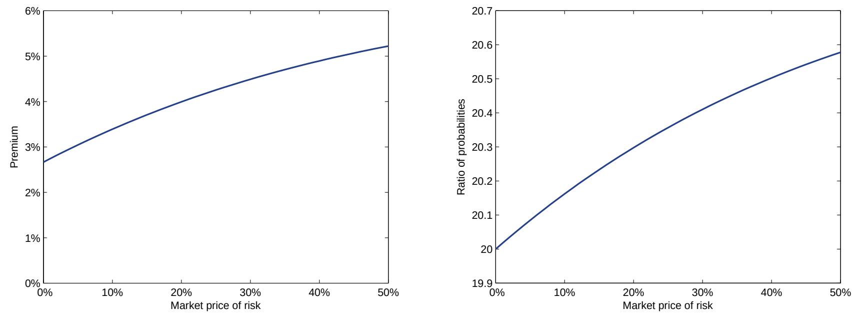

The left panel shows the premium of the longevity swap for different values of the market price of risk. The right panel shows the shape of the ratio of mortality rates (ratio 2 of equation 3.10 over different market prices of risk.

Market price of risk
Increasing the market price of risk parameter results in a higher premium for the longevity swap. The base case premium is 2.737%. Figure 3.7 shows the premium when the market price of risk is varied between 0% and 50%. When the market price of risk is zero, the premium is only determined by the difference in discounting, which is caused by the difference between market and risk-free interest rates. Increases in the market price of risk result in a decreasing increase of the premium. The reason for this can be found in the ratio of the probabilities from equation 3.10, which is shown in the right panel of figure 3.7: this ratio has the same shape as the premiums (because there is no λ in the other parts of the equation, it is not surprising that they are constant over the different market prices of risk).

The result can be explained by the definition of the Wang transform for the survival probability (equation 3.4). A higher market price of risk results in a higher transformed probability (as the Wang transform uses a higher number from the standard normal cumulative distribution) and therefore a higher present value for the floating leg of the swap. And because the premium has to make the present values of both legs equal, it has to increase as well. The asymptotic property of the premium can be explained by the distribution characteristics. As $(\Phi^{-1}(p) + \lambda)$ is always positive, the cumulative probability required is at the right side of the bell-shaped cdf of the standard normal distribution. The increase in cumulative probability hence decreases when λ goes up: we end up in the tail of the normal cumulative distribution.

Maturity
A longer maturity of the contract comes with a higher price of the longevity swap. This makes sense because at the initiation of the swap, the accurateness of mortality rate estimates will be better for the nearby maturities. The premium of the swap for a maturity of 1 to 40 years is shown in figure 3.8. The premium rises with the maturity, but at a decreasing rate.

Equation 3.10 can also be used to explain the pattern. The right panel of figure 3.8 presents the evolution of the different ratios of equation 3.10 when the maturity is increased. Both the ratio of the discount factors and the ratio of the mortality rates are linear, but the denominator of the scaling factor shows a shape similar to the premium.

Figure 3.8: Premium of the longevity swap for different maturities

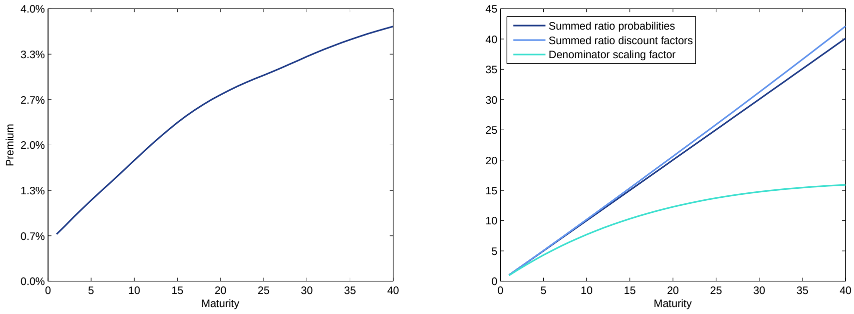

The left figure shows the premium when the maturity is increased from 1 to 40 years. The right figure shows the development of the different factors in the premium equation over time. The ratios of discount factors and mortality rates show a linear relationship. The denominator of the scaling factor has the shape of the left figure.

Figure 3.9: Premium of the swap when the survival rates are changed

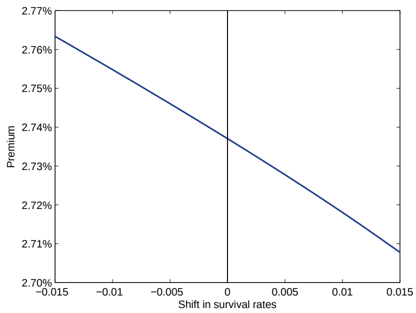

The premium of the longevity swap when all survival rates are shifted by the same amount.

Best Estimate Mortality
When estimated survival rates increase, the premium declines approximately linear; this can be seen in figure 3.9. The reason for this is found in the ratio of the probabilities: when the survival rate increases, the transformed probability increases stronger than the untransformed probability, resulting in a lower ratio. The stronger increase of $S(t)$ as compared to $H(t)$ can be explained by the fact that the untransformed probabilities increase, but the market price of risk parameter remains constant, resulting in a nonproportional increase.

The conclusion that higher survival rates come with lower premiums is confirmed by the different premiums resulting from the use of different BEM-methodologies. The alternatives used are the procedures explained in section 2.2.3: the CBS estimates, the CRC model and the (adjusted) Lee-Carter model. The respective BEMs are shown in figure 2.2; the survival rates ($1-$ the shown mortality rates) are higher for the CBS estimates than for the CRC and Lee-Carter projections, resulting in a lower premium. The differences between the three however are small, as can be seen in table 3.2.

Table 3.2: Premiums for different BEM methods

BEM method	Premium
CBS	2.7370%
CRC	2.7372%
LC	2.7373%

Premiums resulting from the Wang transform pricing method when the mortality rates from the CBS are used, and when the CRC and Lee-Carter models are applied.

Cohort

Applying the base case longevity swap to different cohorts results in a premium that increases with age (figure 3.10). This is directly linked to the findings of the previous subsection: survival rates decrease at the higher ages, and the Wang transform increases the difference between transformed and untransformed survival rates when they become lower. The right panel of figure 3.10 shows the shape of the survival rate ratio for different cohorts; it is similar to the shape of the premiums.

For a pension fund, 20-year longevity swaps for the 40-year cohort are not relevant, as the policyholders only receive payments from their retirement age onwards. Therefore, figure 3.11 shows the premiums of a forward starting swap, which starts in the retirement year of the cohort under consideration. The maturity is kept at 20 years.

The figure shows that the premiums on this forward starting swap decrease practically linear with the age of the cohort. This is not surprising, both because the uncertainty is larger as the survival rates are estimated further in the future, and also because the payments have to be discounted over a longer period, which increases the difference between the present values of the floating and fixed side payments.

Figure 3.10: Premiums of the swap for different cohorts

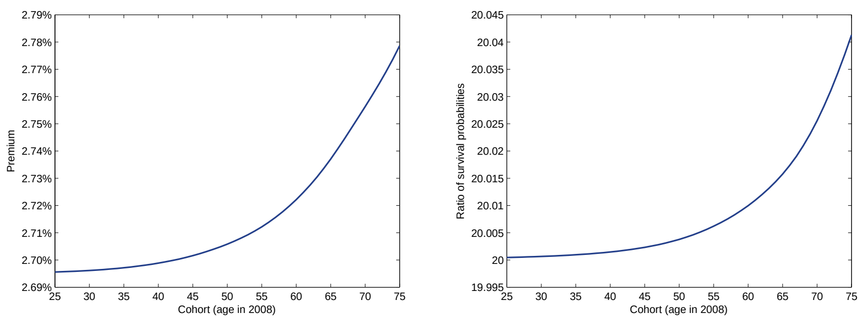

The left panel shows the level of the premium when the longevity swap is applied to different cohorts. The right panel shows the ratio of the Wang transformed to the untransformed survival rates over the different cohorts.

Figure 3.11: Premiums on a forward longevity swap on different cohorts

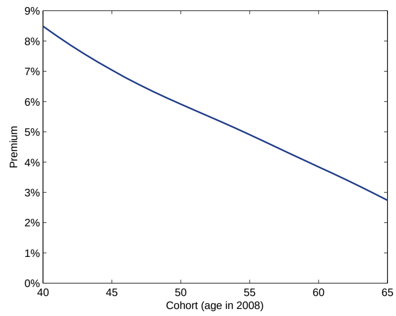

The figure shows the premiums of a forward longevity swap on different cohorts. For each cohort, the premium is determined of a swap that starts in the year the cohort turns 65 and matures 20 years later.

3.3.2 The longevity swap under the Sharpe ratio method

The premiums resulting from the Sharpe ratio method are evaluated on the influence of changes in the Sharpe ratio, its volatility, the maturity of the swap, changes in the BEM, and the age group the swap is based on.

Sharpe ratio
The value of the Sharpe ratio is not the result of a thorough analysis but just an assumption. The assumed ratio directly influences the level of the premiums, as the relationship between the two is linear. The average premium increases by 0.37% for each 0.01 extra Sharpe ratio.

Volatility
The sample period that is used for determining the volatility parameter also influences the premium levels. The use of the entire period on which data is available (1951-2006) will lead to substantial deviations from other prospective longevity markets, e.g. the US market, for which data has only been available since 1968. Figure 3.12 shows the difference in premiums resulting from the Dutch and the US population mortality data, along with the results of the Dutch data when the same sample

Figure 3.12: Comparison of NL and US sample periods

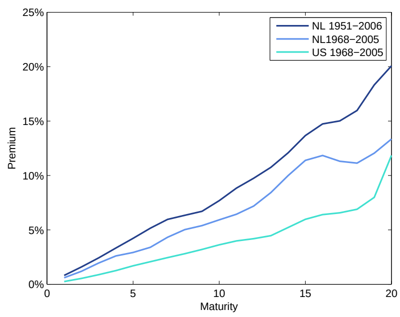

The difference in premiums with the US longevity market when the sample period of the Dutch data is as long as possible, and when the periods are equated.

Figure 3.13: Influence of sample period volatility parameter on premiums

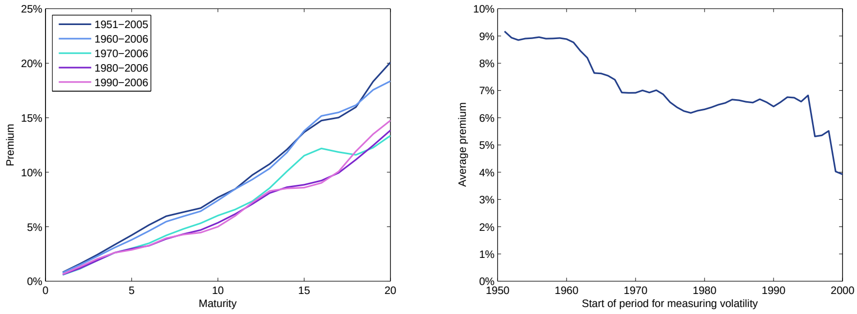

The left figure shows the different premiums resulting from different sample periods (the premiums are not continuous, but lines are drawn between the datapoints for clarity). The right figure shows the development of the average premium when the period used in the determination of the volatility is varied.

period is used as available for the US. The US premiums are substantially below the Dutch data, but this difference is reduced when the same sample periods are used.

The left panel of figure 3.13 shows the premiums for different sample periods. It appears that shorter periods (only more recent data) result in lower premiums. In order to verify this, the right panel of figure 3.13 shows the average premium when the volatility is determined by a period starting in each year possible: 1951-2006, 1952-2006, etc. The premium gradually decreases until the 1980s, when the level remains roughly constant around 6.5%. The premiums further decrease in the second half of the 1990s, but this is probably caused by the fact that the sample period is too short.

Best Estimate Mortality

A change in the Best Estimate Mortality does not influence the premium. Because the premium is stated as a percentage of the fixed payments, its absolute value fluctuates with the survival rates as they determine the payments. The percentage does not change though.

Cohort

The premium of the swap as determined by the Sharpe ratio method does not depend on survival rates. Therefore, when the swap is closed over different cohorts that are immediately in effect, the differences between the longevity swap premiums are determined by differences in the volatilities of the changes in historical mortality rates. This is shown in figure 3.14, where the right panel presents the average volatilities in question.

Just as for the Wang transformed swap premiums, we will also determine the average premium of a forward swap under the Sharpe ratio method. For each cohort, the current average premium is determined of a swap with a maturity of 20 years that starts in the year the cohort turns 65. The results can be seen in figure 3.15. Because the only changing parameter is the maturity, and the maturity is linearly linked to the age, it is not surprising that the average premium linearly decreases when an older cohort is under consideration.

Figure 3.14: Average premiums for the swap for different cohorts

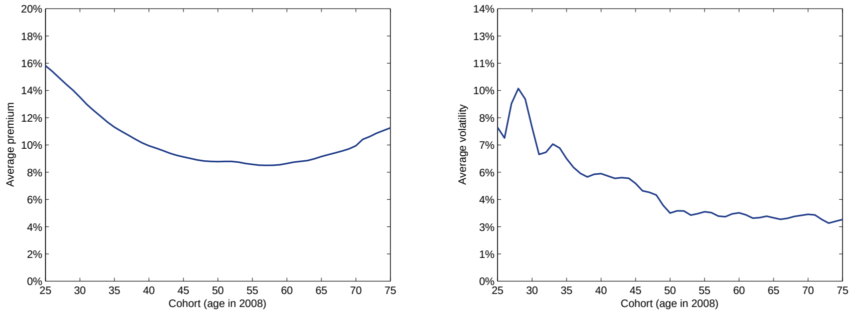

Left: The average premium of the longevity swap when it is applied to different cohorts. Right: Average volatilities of changes in historical mortality rates per age.

Figure 3.15: Average premiums for the forward longevity swap per cohort

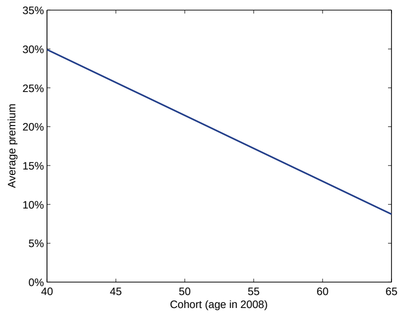

The figure shows the average premiums that apply to a forward swap that starts when the cohort turns 65.

3.4 Conclusion

This chapter examined the theory of longevity swaps. After describing the exact mechanics of the product, two pricing methods were considered. The Wang transform method was presented as a method used in practice in over-the-counter longevity swap pricing. It is based on a normal transformation of the survival probabilities and discounts the two legs of the swap back to the moment of initiation of the contract, enabling the determination of a premium over the fixed payments (the 'preset' side of the swap). The Sharpe ratio method is introduced by JP Morgan as a means to come to the price of a mortality forward product in a future mortality derivatives market. It uses an assumed Sharpe ratio to come to one premium for each payment in the swap. For both procedures, a sensitivity analysis was performed.

For the premium resulting from the Wang transform method, the effects of changes in interest rates, the market price of risk, the maturity, the survival rates and the cohort the swap is based on were considered. The influence of a parallel simultaneous change of both the risk free and the market interest rate curve is very small, but when only one curve changes, the level of the premium is modified substantially. An increase in the market price of risk comes with somewhat higher premiums. The changes in premium are not large when only considering the market prices of risk around the base case of 0.9%. An increase in maturity causes increasing premiums, as the uncertainty about survival rates further in the future is larger. Changes in the estimates of the future survival rates result in adverse changes in the premium: higher survival rates come with lower premiums, because the transformed rates increase stronger than the untransformed rates. When the swap is based on younger cohorts, the premium decreases along with the mortality rates. However, when the swap is a forward starting swap that starts to pay out at the age of 65, the premium increases for younger cohorts.

The price arising from the Sharpe ratio method was also evaluated on changes in its underlying parameters. The Sharpe ratio assumption linearly influences the (average) premium, just as the maturity and the historical volatility. This is not surprising, as the premium is the outcome of a multiplication of these three factors. The historical period used to determine the volatility parameter also influences the outcomes; when the years prior to 1965 are taken into account the volatility is substantially higher. The Best Estimate Mortality does not influence the premium, because it is not included in the pricing equation. The cohort that the swap is based on also does not include mortality rates; the only factor that changes is the volatility used in the pricing equation. The shapes of the premiums and the average volatilities per cohort hence are similar. When forward starting swaps are considered, the average premium linearly decreases with the age of the cohort (at the initiation of the contract).

With the information on the detailed mechanics of the longevity swaps and the exact procedures that are used to come to a price, the swap can be applied to a typical pension contract in the next chapter.

4 Hedging a pension contract with longevity swaps

This chapter applies the developed longevity swap framework to a standard group pension contract of Nationale-Nederlanden. It examines the ability of longevity swaps to reduce the longevity risk in this contract, as determined by its Economic Capital and Market Value Margin.

Broadly speaking, the typical fair value balance sheet of an insurer looks as in table 4.1.[6] The financial component of the liabilities (FCL) is the provision for future payments under the best estimate of mortality rates. In order to determine the market value of the liabilities, an uncertainty margin for the nonmarket risk has to be added to this: the Market Value Margin. The MVM is determined through the Economic Capital. The other way around, the Economic Capital is the amount of risk ING allows Nationale-Nederlanden to engage in. All present and future cashflows ING requires as compensation for this risk capital appear on the balance sheet in the Market Value Margin.

When the hedge is set in place, the swap portfolio appears on the balance sheet with a market value of zero and the levels of the EC and the MVM change. The market value of the swap and the value of the FCL only alter when the best estimate mortality rates are updated; when the update is downwards, both values increase, resulting in a modest increase of the total liabilities. A ceteris paribus decrease of the EC results in a larger amount of free own funds, which can be used for new investment activities; a c.p. decrease in MVM has the same effect. Because the MVM equals all current and future capital charges over the EC, both items will move in the same direction. The hedge is evaluated in terms of the Market Value Margin, as this number includes both current and future ECs. The initial change in EC is also taken into account, as it determines the change in the balance sheet when the hedge is set in place.

The contract is set out in section 4.1, along with the calculations for the liability it results in and the associated Economic Capital and Market Value Margin. Section 4.2 discusses the effect of applying longevity swaps to the contract, starting with one swap, and then also a particular portfolio of swaps. The results are examined in a sensitivity analysis in section 4.3.

[6]The market value of the swaps is usually placed on the asset side of the balance sheet, but we deduct it from the liabilities here for analysis purposes.

Table 4.1: Typical balance sheet of a pension provider

Assets	Surplus	
		Economic Capital
		Free own funds
	Liabilities	
		FCL
		Market Value Margin
		(-Longevity swaps)

The general contents of the balance sheet of a pension provider. FCL stands for the Financial Component of the Liability.

4.1 The group pension contract

This section introduces a standard pension contract and the longevity risk accompanied by it. The longevity risk is measured as the Economic Capital over the longevity risk sources, and more specifically as the discounted value of all capital charges required over this EC in the future: the Market Value Margin. The determination of the MVM associated to the standard pension contract is the main objective of this section. After introducing the contract, the cashflows stemming from it are determined, resulting in a value for the pension contract liability. That information enables the calculation of the Economic Capital. Subsequently, the EC is determined in time, resulting in a value for the Market Value Margin.

The contract under discussion is a Beleggings Overname Koopsom (BOK) contract: an annuity bulk pension contract to which no new liabilities will be added. It is the type of contract Nationale-Nederlanden is left with when a pension fund decides to buy out its contract and start a new contract with a different pension provider. Nationale-Nederlanden is obliged to fulfill the pension liabilities built up until that moment. A typical (standard) version of a BOK contract is used here.

Figure 4.1 shows the level of the pension rights in the contract at the start of the period under discussion (2008) for the different cohorts. The amounts shown are the payments the participants will receive each year after retiring when they are alive (old-age pensions, OP) and the payments their partners will receive each year in which they are alive while the participants have died (widow pensions, WP). It shows that most participants in the fund are males, as their rights are far higher (OPx is above OPy). The pattern of the widow pension associated to males follows a similar pattern but is shifted to the right, as women are on average younger than their spouses. The widow pension associated to women is in consequence shifted to the left.

In order to explain the patterns found, the analysis in this section is also applied to a BOK contract with relatively young and one with relatively old participants. The graphical results of this analysis can be found in appendix B. Figure B.1 shows the pension rights in these contracts.[7] As expected, the

[7] The pension rights in the old and young contract are scaled down proportionally to the total size of the standard contract in order to make them comparable. Initially, the total pension rights in the young contract are €3.8mld and those in the old contract €3.2mld, which is far higher than the €2.6mld in the standard contract.

Figure 4.1: Initial pension rights in the group pension contract per cohort

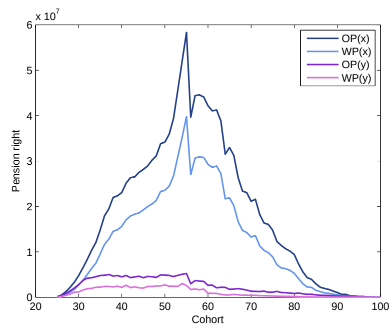

The figure shows the liabilities involved in the standard pension contract. OPx and OPy stand for the rights built up for old-age pensions of men (x) and women (y), WPx and WPy are the associated widow pensions.

younger contract has more pension rights for the lower ages, resulting in a wider distribution than the standard contract, while the older contract is strongly concentrated at the higher ages. Table B.1 lists the composition of each of the contracts to the different kinds of pensions. The portion of pensions associated to men (OPx and WPx) is higher when the participants are older, while the portion of pensions associated to women (OPy and WPy) consequently decreases.

4.1.1 Liability

This subsection further examines the contract by determining the obliged payments stemming from it. Because no new pension rights will be added to the contract, the cashflows can be computed for each point in time in the future. The sum of the present values of these cashflows equals the current pension liability.

The future cashflows for the old-age pensions are determined by translating the appropriate mortality rates ($q(x)$) to survival probabilities ($(t)p(x)$) in the way stated in appendix A. The results are multiplied with the pension rights from the year the cohort under concern turns 65 onwards, and the resulting cashflow is discounted with the swap term structure of 1-1-2008 (as Nationale-Nederlanden uses the swap curve in discounting for internal reporting purposes). In the determination of the widow pensions (WP), all participants are assumed to have a partner.[8] The survival probabilities for the widow pensions are given by the difference between the probability that the partner is alive and the probability that both the partner and the participant are alive, where an age difference of 3 years is assumed (partners are assumed to be of opposite sexes; the man is three years older than the woman). The widow pensions are paid out from the moment the pension fund participant dies onwards for the rest of the life of the partner.

Following the practice at Nationale-Nederlanden, the old-age pensions are assumed to take place on a continuous basis; the value assigned to each year is the average of the payment when it would take place at the beginning and at the end of each year. This averaging is conducted after discounting the post- and prenumerando payments. Widow pensions are assumed to be paid out at the beginning of the year.

The survival probabilities used in these calculations are deduced from the CRC 2006 table. This generation table projects mortality rates for Dutch men and women until 2056. Because not all participants in the fund and their partners will have deceased by then, the table is extended to 2100. This is conducted by assuming no further changes in mortality rates after 2056.[9]

Figure 4.2 shows the present value of the future cashflows, both by future year and by cohort, subdivided in old-age and widow pensions for males and females. When considering the future cashflow pattern by year, the old-age pensions first strongly increase as more and more participants in the fund retire, and then gradually decrease when the participants die. For the women pensions, the cashflow pattern is flat in the first few years. This is related to the pattern of the pension rights per cohort: the

[8] The reason for this is that at the age of 65, all participants can choose between either keeping their old-age pension and the widow pension for their partner separated, or summing the amounts built-up and receiving a higher old-age pension themselves (implying that there no payments are made to a partner when the participant deceases). This is in conformity with Dutch law.

[9] Statistically, it would be better to extend the model with 44 years by continuing the trend. This is not current practice at Nationale-Nederlanden, because it is not considered realistic that mortality rates will continue to decline indefinitely.

Figure 4.2: Present values of the future cashflows from the contract

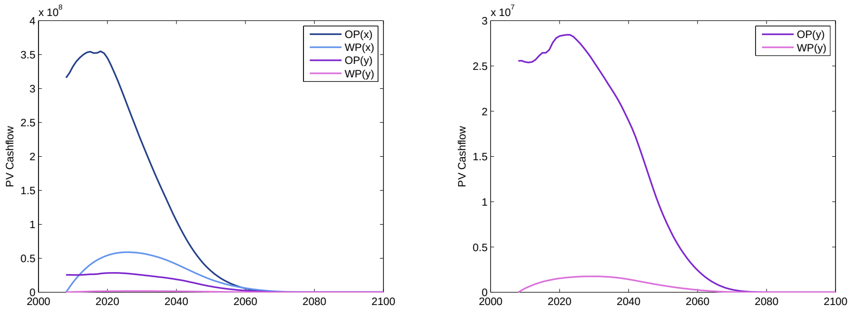

a. Present value of the cashflows per future year

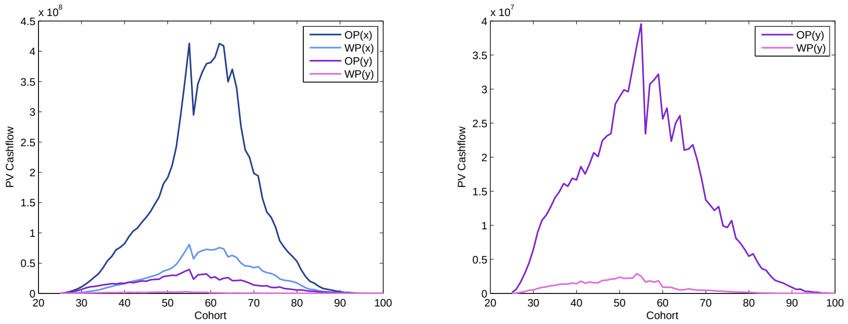

b. Present value of the cashflows per cohort

This figure shows the pattern of the future pension payments stemming from the standard BOK contract; OP stands for old-age pensions, WP for widow pensions. The left figure of panel a shows the cashflows stemming from the contract in each future year, the right figure depicts the same for only the female pensions (as their pattern is not clear in the main figure). Panel b shows the same per cohort (the current age of the participants). All numbers are discounted to 1-1-2008.

peak of the cashflow pattern is not at 65 as for the males but a few cohorts younger, hence the highest level is reached later. Also, the women pensions are relatively large at the younger ages. These two data characteristics probably stem from the fact that it has become more and more common for women to have a job in recent decades. The present value of the cashflows resulting from widow pensions starts at zero as a result of the fact that no widow pensions that are already in effect are incorporated in the standard BOK contract data.

The pattern of the cashflows per cohort roughly follows the pension rights as shown in figure 4.1. The level of the widow pension cashflows is far below that of the old-age pensions because they are the result of lower probabilities (the survival probabilities of the partner minus the survival probabilities of the participant and the partner together).

The total liability at the beginning of 2008 as determined by this methodology is €12,587mln.

Figures B.2 and B.3 show the same pictures for the introduced young and old contract, respectively. The cashflows in time follow the expected pattern: in the young portfolio, the peak in the cashflows occurs later as more participants reach the retirement age later, while the cashflows of the old portfolio start at practically their highest level. In accordance, the peaks of the different types of pensions are further away when the age is higher. The pattern of the cashflows of the old-age pensions for women in the young portfolio show the same effect as found in the standard contract: they are relatively high at the younger ages. The cashflow patterns per cohort again roughly follow the pension rights per cohort.

4.1.2 Economic Capital

A certain amount of Economic Capital has to be held over the liabilities in order to manage the risks involved in them. The concept of Economic Capital was introduced in section 2.1. Here we discuss the actual calculation of the separate sources of EC.

The EC calculations for the calamity, trend uncertainty and level uncertainty risks are determined by comparing the liability of the BOK contract arising from the use of the best estimate mortality tables, with the liability under the mortality tables arising from a particular jump or structural break in the evolution of the survival rates. The best estimate mortality table is the population CRC generation table multiplied by the experience mortality factors (the average ratio of historical mortality data of Nationale-Nederlanden and population mortality per age as determined by internal calculations; they are shown in the right panel of figure 4.3). The EC Volatility is determined in a different way.

1. Volatility
   The EC for volatility risk estimates a binomial distribution around the BEM to determine the 99,95% upper bound of the capital at risk in the upcoming year. Using the central limit theorem, this is 3.29 times the standard deviation of the capital at risk (the 99.95% confidence level of the normal distribution). The capital at risk ($R(i, s)$ (s is x or y) is determined for each cohort (i) and for both sexes (x and y); it is the amount of capital lost in case the participant dies. It is the result of the following calculations:

   - the present value of all future widow pension payments when in effect from now (this has to be paid in case the participant dies in the upcoming year);
   - minus the provision for the old-age pensions (the present value of the old-age pension payments, which will not be paid when the participant dies in the upcoming year);
   - minus the provision for the widow pensions (the present value of the widow pensions, as determined by the difference in survival probabilities for the participant and the partner; these payments are already included in the first bullet).

   Hence, for males of the cohort that is aged 45 in 2008,

$$\begin{aligned} R(45, x) =& WPx(45) \cdot a(42) - OPx \cdot (20)|a(45) \\ & - WPx \cdot ((23)|a(42) - [(23)|a(42) \cdot (23)|a(45)]) \end{aligned} \tag{4.1}$$

   where $(t)|a(x)$ is the current price of an annuity paying €1 each future year in which someone aged x is alive, starting in t years. The absence of t in the equation implies that the annuity starts immediately. In case the Capital at Risk turns out to be negative, it is set to zero. The EC is determined separately for each cohort (i). The first part of the standard deviation is associated to the pensions of male participants, the second part to the pensions for female participants.

$$\begin{aligned} EC_{vol} = & \sum_{i=25}^{100} 3.29 \cdot \\ & \sqrt{(R^2(i,x) \cdot q(i,x) \cdot (1 - q(i,x))(R^2(i,y) \cdot q(i,y) \cdot (1 - q(i,y)))} \end{aligned} \tag{4.2}$$

The EC Volatility is €89mln or 0.71% of the best estimate liability.

2. Calamity
   The extra liabilities from a one-time doubling of mortality are determined by doubling the mortality rates of the BEM table over the first year. For later years, the best estimate is used. The arising liability is compared to the liability based on the full BEM tables. This results in a total EC Calamity of -€51mln. As this number is negative, the EC Calamity is set to zero.

3. Trend uncertainty
   The procedure to come to the EC Trend starts with determining the lower bound of the population CRC best estimates (with a confidence level of 83%). The historical data the CRC best estimates are based on, are supplemented with an artificial observation for each age, which is the first datapoint on the lower bound of the best estimates. Based on the new dataset (the historical data and the artificial observation), the CRC best estimates are determined again. The new best estimates are considered to be the lower bound of the population mortality. This methodology is illustrated in the left panel of figure 4.3. The resulting mortality table is multiplied by the factor of experienced mortality for each age. [10]

   This table is used to determined the liability; the outcome is compared to the liability resulting from the use of the best estimate table. For the contract under concern, the EC Trend adds up to €270mln, or 2.14% of the best estimate liability.

4. Level uncertainty
   The EC Level is obtained by multiplying the population CRC table with the lower bound of the experience mortality factors as determined by the 83% confidence level. In the right panel of figure 4.3 can be seen that this lower bound is quite close to the best estimate of the factor.

   The economic capital for this source of risk is €13mln, or 0.1% of the liability under the best estimate.

The different sources of EC sum to €372mln, which is 2.95% of the level of the liability under the best estimate mortality table. An overview of the data is given in table 4.2. This table underlines that the ECs for calamity, trend uncertainty and level uncertainty risk are determined by the difference in liability resulting from the source specific calculations and the Best Estimate Mortality table. The largest source of risk (the main determinant of the total Economic Capital) appears to be the trend uncertainty: the risk of a general decrease in the population mortality rates.

[10] In case the model indicates a lower mortality probability for men than for women at the end of the estimation period, the factor that changes the CRC data over time is not determined by the difference between the first and the last observation, but by the difference between the last observation and the estimated probability for males at the end of the estimation period. This is determined per age. The reason for this is that, based on genetics, it is deemed impossible for men to have higher survival probabilities than women.

Figure 4.3: Illustration of methods used in the determination of ECTrend and ECLevel

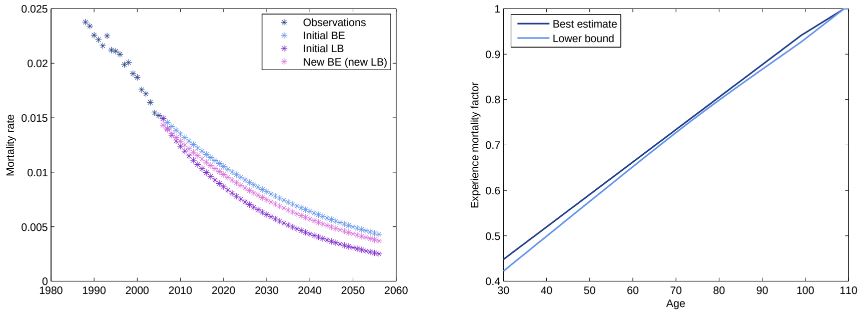

The left panel of this figure illustrates the procedure used to determine the lower bound of the population mortality for 65 year old males (for the purpose of calculating the EC for trend uncertainty): a new best estimate is determined by adding an artificial observation to the historical observations, which is the first estimate on the lower bound of the initial best estimates. The right panel shows the Experience Mortality factor for males for each age and its associated 83% lower bound: the ratio of past mortality rates experienced by insureds of Nationale-Nederlanden to the population mortality rates in the same period. This is used in the determination of the EC for level uncertainty.

Table 4.2: Overview of the Economic Capital in the standard BOK contract

		BEM	Volatility	Calamity	Trend	Level	Total
Liability	OP(x)	9,570		9,475	9,862	9,597	
	WPx	1,927		1,977	1,880	1,913	
	OPy	1,028		1,021	1,054	1,028	
	WPy	61		63	60	62	
	Total liability	12,587		12,536	12,857	12,600	
Economic Capital (mln)			89	0	270	13	372
Economic Capital (%)			0.71%	0%	2.14%	0.10%	2.95%

The Economic Capital for the different sources of risk of the BOK contract in €1,000,000, where BEM stands for the Best Estimate Mortality (the CRC table corrected for experienced mortality).

The same calculations are carried out on the data for the young and the old contract. The results can be found in tables B.2 and B.3. As could be expected based on the compilation of the contracts, the liability for male pensions (OPx and WPx) is larger when the participants in the contract are relatively older, and vice versa for the female pensions (OPy and WPy). An older contract comes with a larger portion of total payments in the beginning of the period under concern. This results in a higher total value for the liability (the cashflows are affected less by discounting) and would result in a more negative EC Calamity (more people die in the first year) when that number would not have been set at zero. The minor influence of discounting results in a higher percentage for the EC Volatility as well, as the Capital at Risk is higher. The EC Level also increases with the average age of the participants in the contract, as a result of the fact that the mortality rates of the generation tables used in the determination of the EC Level are lower, which again has a larger impact on the older contracts because of the discounting. In all contracts, the largest percentage belongs to the Economic Capital required to cover the trend uncertainty. The level of this percentage is not linked to the average age of the participants.

4.1.3 Market Value Margin

The Market Value Margin estimates the present value of the capital charges over all future ECs (from $t = 0$ to ω). For each source of risk k (volatility, calamity, trend uncertainty and level uncertainty), the MVM is calculated as follows:

$$MVM(k) = CoC \cdot \sum_{t=1}^{\omega} \frac{EC(k,t) \cdot DF(k)}{(1+r)^t} \tag{4.3}$$

In this equation, $DF(k)$ stands for the diversification factor of risk source k. This factor results from the assumption that not all shocks in the EC calculations will occur at the same time. The ING CIRM department has set all diversification factors for the sources of life risk at 40%. CoC stands for the Cost of Capital that ING's shareholders demand over the capital held in excess of the risk-free interest.[11] The CoC is set at 4%.

The MVM calculation requires a value for the EC over time. This is achieved in the same way as described in the previous subsections for the upcoming 43 years (2008-2050). Because the contract starts in 2008, the 2008 survival probabilities determine the number of people alive at the beginning of 2009 and therefore have to be taken into account when determining the liability in that year. Hence, the cashflows stemming from the contract are determined in 2008, and for each year 2008 + p, the cashflows due in the first p years are removed. By discounting the cashflows with the new interest rate curve, the liability in year 2008 + p is obtained.

This procedure is used to obtain the ECs for volatility and level uncertainty. For the EC Calamity, the mortality rates in year 2008 + p are doubled; apart from that, the same procedure is applied. The determination of the EC Trend is somewhat more cumbersome. Theoretically, it requires a new lower bound of the population mortality in each future year. Because there are no historical observations after 2006, in each year ($p - 1$) the dataset is supplemented with the first $p - 1$ estimates of the best estimate CRC table. These best estimates are based on 5 year moving averages, hence the raw data is far more volatile. Basing each new lower bound on the appended 'historical' dataset does therefore not provide plausible results. As an alternative, the proportion of the first lower bound estimate to the last historical observation is used to determine the new lower bound artificial observation in each year after the first year. For the remainder of the calculation, the usual procedure is followed.

The future EC figures are discounted with the future interest curves that are implied by the swap curve of 1-1-2008. This curve is assumed to be flat after the last observation (the 30 year swap rate). The methodology used to determine the future curves is the method of bootstrapping discount rates, as for instance explained in Buckley (2004):

$$D_{1\to 2} = \frac{D_{0\to 2}}{D_{0\to 1}} \tag{4.4}$$

where $D_{t\to T} = (\frac{1}{1+r_{t\to T}})^{T-t}$. The discount factors are calculated back by $r_{t\to T} = (\frac{1}{D_{t\to T}})^{T-t} - 1$. The resulting structure is shown in figure 4.4.

[11] It is assumed that the Economic Capital itself is invested in risk-free assets.

Figure 4.4: Implied interest rate curves 2008 - 2050

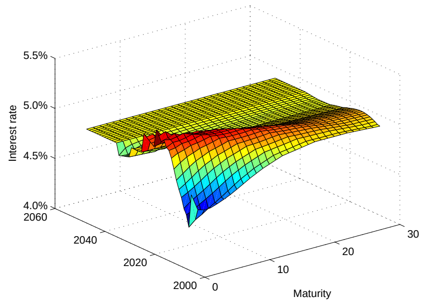

Future interest rate curves as implied by the current (1-1-2008) curve. After the last observed maturity, a flat curve is assumed.

Figure 4.5: Liability in time

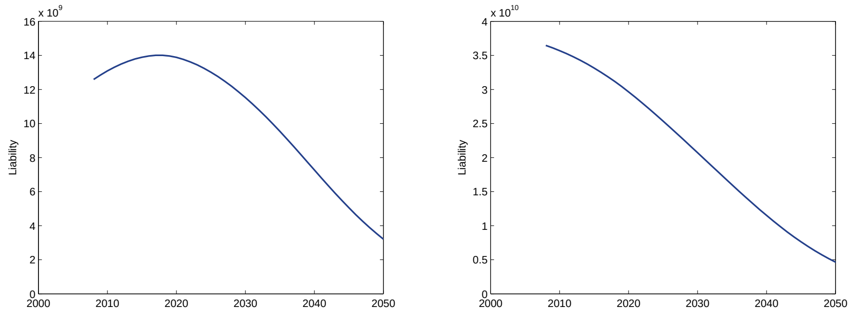

The development of the liability in time for the standard BOK contract is shown in the left figure. The right figure shows the same figure when discounting is omitted.

Using the obtained interest curves, the value of the liability (the present value of all future cashflows) is determined for each future year (2008-2050). The left graph in figure 4.5 shows the level of the liability for the standard BOK contract until 2050; figure B.4 shows these data for the young and the old contract. A comparison of the total liabilities indicates that the liability starts at a higher level when the participants in the fund are older. This is not surprising as the pension cashflows are in the more nearby future and hence their discount rates are lower. Because the liability equals the present value of all future cashflows, it should decrease steadily in time (as less cashflows remain in the future). The reason for the increase in the first years is hence caused by discounting. To illustrate this, discounting is omitted in the right panel of figure 4.5, resulting in the expected pattern. The influence of discounting is underlined further when comparing the liabilities of the different contracts. A contract with younger participants has the peak in liability later, as the peak in cashflows is later and hence the highest cashflows are discounted over a shorter period.

Figure 4.6 shows the Economic Capital until 2050, where the right figure splits it up into the components. The patterns of each of the EC sources are discussed in turn.

The EC Volatility decreases almost linearly until 2035, and is zero from that year onwards. This pattern can be explained by the development of the Capital at Risk over time. Figure 4.7 shows the Capital at Risk per cohort for the first year (2008) and year 25 (2033). In the first year, the present

Figure 4.6: Economic Capital in time

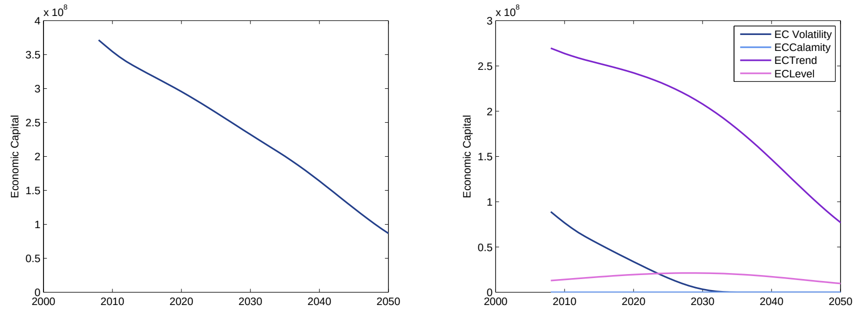

The left figure shows the total Economic Capital over the BOK contract in time. The right figure shows its components.

Figure 4.7: Explanation of the EC Volatility pattern

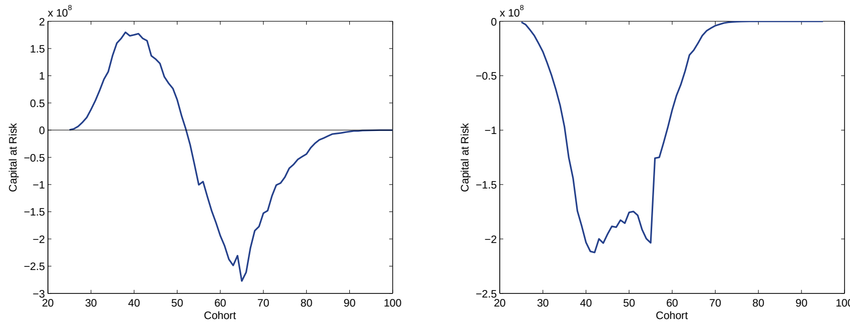

The figures show the netted cashflow (widow pensions minus provisions for old-age and widow pensions) for each cohort in 2008 (left) and 2033 (right).

value of the widow pensions of the younger cohorts is higher than the associated provisions held for deferred old-age and deferred widow pensions, resulting in a positive Capital at Risk. In the course of time, the value of the provisions increases, while the present value of the widow pensions goes down with the remaining lifetime of the potential widows decreases. In year 25, the value of the provisions is higher than that of the widow pensions, hence the Capital at Risk is zero, and so is the EC Volatility.

As expected, the EC Calamity is zero over the entire period. It can only be positive when the present value of the widow pensions that take in effect as a result of the extra mortality, is larger than that of the (deferred) old-age pensions. That can only happen for the youngest cohorts, where as a result of the discounting, the value of the old-age pensions is very low. As the EC Calamity numbers are zero at the beginning of the period, it doest not become positive in time.

The pattern of the EC Trend shows a decrease that is stronger after 2030. An explanation of this pattern is found in the methodology employed to obtain the mortality tables for the EC Trend in time: a worst case scenario is added to the (semi-)historical observations. This worst case is the first point on the lower bound of the forecasts, and hence determined by a certain percentage times the standard deviation of the observations. The relatively low mortality rates in the first years hence result in smaller absolute difference between the best estimate and the associated lower bound; the EC Trend is mostly determined by the participants with high ages. When the younger participants grow older,

Figure 4.8: Explanation of the pattern of EC Level

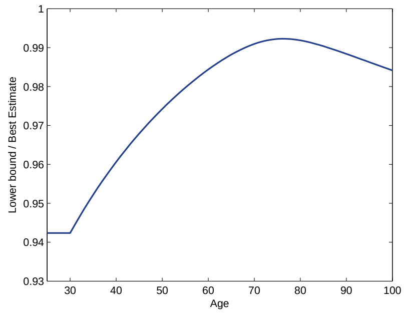

The figure shows the relationship between the lower bound and the best estimate of the experienced mortality factors for males over the ages.

their contribution to the EC Trend increases along with their mortality rates. The level of EC starts declining stronger when the older participants start to die.

The pattern of the EC Level is explained in the same way. It is the result of the difference in liability resulting from the population mortality tables times the best estimate of the experience mortality factors and the liability resulting from the population mortality tables times the lower bound of these factors. Figure 4.8 shows the ratio of the factors for males.[12] For the ages under 30, the lower bound is assumed to equal the 30-level, the pattern increases until 75, and then slowly decreases again. The ratio of the factors varies between 0.94 (at 30) and 0.99 (at 75), hence the differences are not substantial. The effect of the fact that a percentage is taken from the mortality tables is larger; the contribution to EC of young cohorts increases with their age and mortality rates.

The value of the components of Economic Capital in time enables the determination of the Market Value Margin. Equation 4.3 results in a MVM of €80mln at the beginning of 2008. Figure 4.9 shows the development of the MVM; it fades out with the risks it is associated to.

[12]The lower bound of the experienced mortality for the females is set equal to the best estimate, because internal Group Life department calculations indicate that the difference between the two is insignificant.

Figure 4.9: Market Value Margin in time

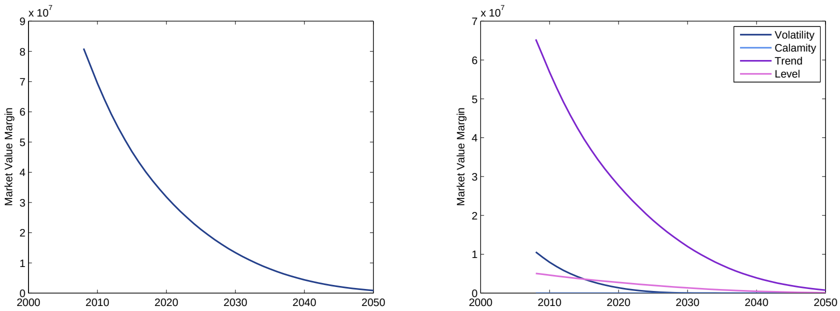

The Market Value Margin over the BOK contract in time; right the total, left the components.

EC and MVM of the young and the old BOK contract
EC calculations were also executed on the young and the old contract to enable the above analysis of the standard contract. The EC figures for these contracts can be found in appendix B in figures B.5 and B.6.

For the EC Volatility, the lower average age of the participants in the young contract results in positive Capital at Risks over a longer period and hence a larger percentage for the EC Volatility. This is reversed for the old contract. Just as in the standard contract, the EC Calamity numbers for the young and the old contract are zero in the first year, and hence also during the rest of the period. In accordance with the reasoning above, the change in pattern in the EC Trend is stronger in the young contract, while it is practically absent in the old. Hence in the young contract, we see the 'contribution' in EC of the youngest cohorts increase in the first years, and they start decreasing when the oldest cohorts fall away. The old contract already starts at this phase; the influence of younger ages is far lower. This effect is also visible for the EC Level: The peak is earlier when the average age of the participants in the contract is higher.

The patterns of the MVMs for the young and old contracts (figure B.7) exhibit the same pattern as the standard contract, they only differ in the level they start at. This difference makes sense: the younger the contract, the higher the initial MVM, as the ECs are higher during a longer period of time.

4.2 Hedging the contract with longevity swaps

This section examines the effects of hedging the contract with a portfolio of swaps, and makes an attempt to explain them. The methods of the previous section are used to determine the liability, EC and MVM of both the contract and the swap when certain shifts are applied to the mortality rates. The proposed hedge results from combining the two.

Figure 4.10 shows the effect of parallel overnight shifts in the mortality rates in the range of -15% to +15% on both the level of the liability and the Market Value Margin of the (stand alone) contract. A parallel shift in the mortality rates simply implies that the entire population generation tables for males and females are multiplied by the same shift factor. After assuring that the mortality rates never exceed 1, and equal 1 for all ages above 120, the standard calculations are applied. The liability and the MVM both increase when mortality rates are lower than expected and vice versa. This is not sur-

Figure 4.10: Change in liability and MVM for different shifts in mortality rates

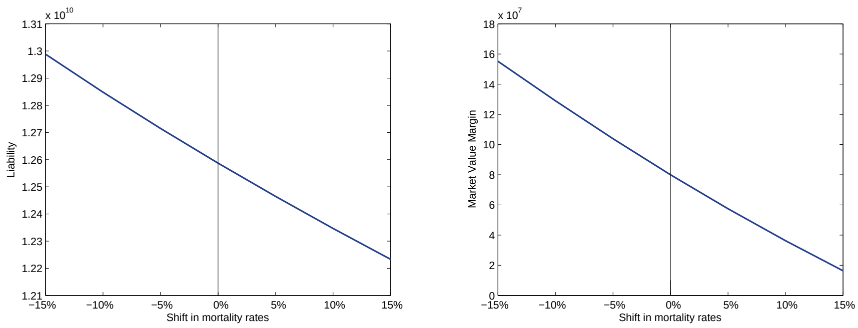

The left panel of the figure shows the level of the liability when a certain shift is applied to the mortality rates. The right panel shows the same for the Market Value Margin.

prising, as lower mortality rates result in cashflows that have to be paid over a longer period (higher liability), and the differences in the liabilities determined for the EC calculations are larger when the ECs change approximately proportionally (higher MVM).

As stated earlier, past mortality rates on average appeared to be lower than estimated in advance, resulting in changes in liability and MVM. Longevity swaps are intended to relax these changes. However, as the future is uncertain, the influence of swaps on the MVM when the mortality rates exactly equal the current best estimates might even be more relevant. When the longevity swaps cover part of the higher ECs that follow from the shock scenarios, the level of the current MVM decreases as well, making these derivatives an attractive investment.

4.2.1 One longevity swap

In order to examine the mechanics of applying a swap to the contract, we discuss the implementation of one swap before exploring the effect of an entire portfolio. As a general example, a longevity swap with a maturity of 20 years is applied to the standard BOK contract. It concerns a swap on Dutch males aged 65 in 2008, with a notional of €10mln (arbitrary number). The payer side (fixed payments) of the swap is determined using the generation tables of the CBS Best Estimates minus the premium. This premium is subtracted because we work with mortality rates here: the expected level of the payments for the receiver (the floating leg counterparty) is decreased. The cashflows arising from the swap follow from the difference between the realized population mortality rates and the CBS Best Estimates minus the premium.

We assume a premium of 2.737%, hence the price is determined by the Wang transform method. This implies that the receiver only pays when the mortality rates appear to be more than 2.737% below the best estimate; the payer (the pension provider) has to pay when the realized mortality rates equal the best estimates. Figure 4.11 shows the cashflow over time from the point of view of the payer in two cases: when the mortality rates equal the Best Estimate (the net cashflow goes to the receiver) and when the mortality rates appear to be 5% below the BEM (the payer receives the net cashflow). The net cashflows increase as a result of the higher mortality rates for higher ages, as the difference between the two cashflows is a fixed percentage (the premium and the shock).

Figure 4.11: Swap cashflows for a shift of mortality rates of 0% and -5%

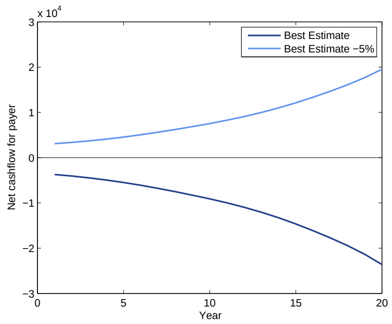

The figure shows the net cashflows from the point of view of the payer, when mortality rates equal the best estimate and when they appear to be 5% below it.

In order to determine whether the swap decreases the longevity risk in the BOK contract, a measure of its EC is required. The EC calculations of the contract are based on CRC tables, while the swap is based on the CBS best estimates. It is not possible to apply the EC procedures to the CBS tables, because the distribution underlying the CBS best estimates is unknown (making it impossible to determine lower bounds etcetera). The CRC tables are therefore 'translated' to CBS best estimates by determining the ratios between the estimates of the two tables for each age and at each point in time. Figure 4.12 shows this ratio for the first 20 years (2008 - 2027). The differences between the two tables are substantial and differ over the ages. Hence the EC scenarios are determined using the CRC tables, and the resulting generation tables are multiplied with the ratio of CBS to CRC (per age, sex and point in time) for the determination of the payout of the swap. The factors are assumed to remain constant in the years after the last available ratios of the tables.

The Economic Capital calculations are adapted to the swap as follows:

- EC Volatility: The EC Volatility in the contract equals the 99.95% upper bound of the Capital at Risk around the best estimate. In the BOK contract, the Capital at Risk equals the amount of money that is lost in case all participants die. We therefore determine the present value of the cashflows stemming from the swap in that instance. The q and $1-q$ the outcome is multiplied with, are set at the corresponding numbers in the mortality tables of the payer side of the swap (hence the CBS Best Estimates minus the premium).
- EC Calamity: The difference in the swap payout to the payer in the standard (realized population mortality, determined in CRC and converted to CBS) and the shock scenario (idem, but with double mortality rates in the first year). The payer side of the swap is of course fixed at the initial CBS population best estimates minus the premium.
- EC Trend: The calculations of the EC Trend are similar to those of the EC Calamity, where the shock scenario is of course the scenario applying to the trend uncertainty risk.
- EC Level: The swap does not have an EC for the level uncertainty. This is because this source of risk determines the effect of a change in the insureds mortality rates while the population mortality remains constant. As the swap only considers population data, the EC Level stemming from it is zero by definition.

Figure 4.12: Ratio of CBS to CRC estimates over the ages

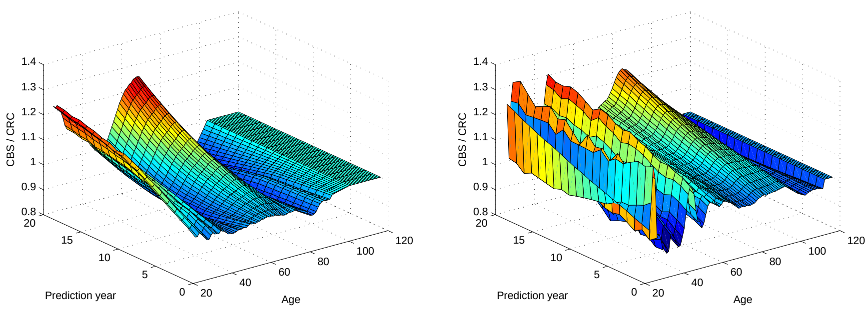

The figure shows the ratio of CBS estimates to the CRC (population) estimates over 2008-2027 per age. The left figure shows the ratio for the male tables, the right figure for female tables.

Table 4.3: Overview of hedge with one swap

	Contract	Swap	Contract with swap
EC Volatility	88.87	0.06	88.81
EC Calamity	0	0	0
EC Trend	269.59	4.28	265.31
EC Level	12.94	0	12.94
Total EC	371.4	4.34	367.06
Liability	12,587.02	1.67	12,588.69
MVM	79.98	–	79.52

This table gives an overview of the ECs in the first year, the liability, and the Market Value Margin for the BOK contract, a 20-year longevity swap on 65 year old males with a notional of €10mln, and the combination of the two. Numbers are in €1,000,000.

The Economic Capital numbers of the BOK contract are the result of differences in liability. As each of the differences in proceeds of the swap are measured at one point in time, their future payments need to be discounted at each point in time to make them equivalent to the ECs. Table 4.3 lists the Economic Capital figures for the swap under concern when the mortality rates equal the CBS best estimates. As before, the present value of the cashflows arising from the swap is negative (hence a positive liability). The EC Calamity of the swap is -€5.79mln. The reason why it is set to zero in the table is explained below.

Figure 4.13 shows the development of the Economic Capital figures over time. The patterns of the EC Trend and EC Volatility gradually decrease towards zero. The EC Calamity however is strongly negative, pulling the total Economic Capital below zero as well. The striking pattern of the total EC is also explained by that of the EC Calamity. For the contract EC determination, the EC Calamity was set to zero in case it was negative. As the EC of the swap is deducted from the EC of the contract per source of risk, it would be best to deduct the (negative) EC Calamity of the swap from the (negative) EC Calamity of the contract, before determining whether the result is negative or not. However, this figure shows that the EC Calamity as determined in this way is far out of proportion and its influence is not representative; it causes the swap to increase the total EC of the contract when applied to it. Because this is not realistic, this source of risk is ignored in the rest of the calculations.

Figure 4.13: Economic Capital of the longevity swap

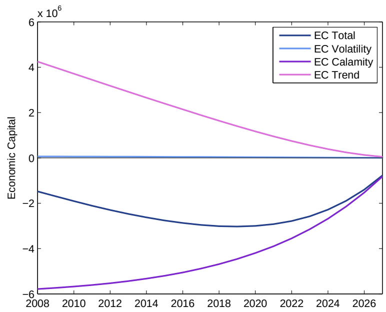

The figure shows the evaluation of the different ECs for the 20 year longevity swap on 65 year old males over its term.

Combining the swap with the pension contract implies that the ECs of the swap are subtracted from the ECs of the contract. The Market Value Margin this results in is €460,000 lower as compared to the MVM of the contract in itself; this is a reduction of 0.6% of the initial MVM. Besides, the hedge reduces the EC in the first year with 1.17%.

4.2.2 A portfolio of longevity swaps

This subsection expands the analysis in the previous subsection to a portfolio of longevity swaps. There are two ways to hedge the pension contract: hedge changes in cashflows (cashflow hedge) or hedge changes in the associated liability (value hedge). The hedging methods determine the notionals of the swaps differently and hence result in different conclusions. Both hedging methods are discussed in turn to enable which one suits the purpose of the hedge best. That method is analyzed further in a sensitivity analysis.

Cashflow hedge
In a cashflow hedge, the notionals of the swaps are set such that changes in cashflows are covered by the proceeds of the swap portfolio. A cashflow hedge only includes swaps on cohorts older than 65, because including younger cohorts would involve the risk of a paying obligation without hedging a cashflow. For each age above 65 and both sexes, the notional is a product of the following factors.

- *The old-age pension right*
  The basis of the notional is the value of the old-age pension right associated to the specific cohort and sex. The accompanying widow pension is not taken into account, because it is associated to another cohort and sex.

- *The average best estimate experience mortality factor*
  The liability of the contract is determined by the best estimate mortality, while the proceeds of the swap come from the population mortality table. The difference between the two consists of the best estimate experience mortality factor. This implies for each age x at each future time t:

  $$q_{ins}(x,t) = \alpha(x) \cdot q_{pop}(x,t) \tag{4.5}$$

  where $\alpha(x)$ is the best estimate experience mortality factor associated to age x, and q_{ins} and q_{pop} are the mortality rates of the insureds and the population, respectively. The old-age pension right is multiplied with the average of best estimate experience mortality factors of the age at hand in the future 20 ages.

- *The average ratio of CRC to CBS best estimates*
  The liability is determined by CRC tables, while the swap based on CBS best estimates. The proceeds of the swap (in CBS) influence the liability by a factor CRC / CBS (the inverse of the ratios shown in figure 4.12). Equivalent to equation 4.5, this implies:

  $$q_{CRC}(x,t) = \frac{CRC(x,t)}{CBS(x,t)} \cdot q_{CBS}(x,t) \tag{4.6}$$

  The appropriate factor the notional is multiplied with is the ratio of CRC to CBS best estimates averaged over the age and time of the cohort during the upcoming 20 years.

- *The premium*
  The picture is distorted further because of the premium (π) that is included in the swap. The premium decreases the proceeds of the swap because the mortality rates used to determine the fixed side of the swap are multiplied with $(1 - \pi)$. Therefore, the notional has to be raised by it. The premium is based on the CBS tables for the cohort and sex under concern. We use the Wang transform procedure to determine the premiums.

The notional of the swap for females currently aged 70 is hence determined as follows:

$$Notional(y, 70) = OPy(70) \cdot \alpha(y, 70) \cdot \sum_{j=0}^{19} \frac{CRC(70 + j, 2008 + j)}{CBS(70 + j, 2008 + j)} \cdot (1 - \pi) \tag{4.7}$$

In order to save computational time, we do not work with swaps for every cohort. Instead, they are summed in swaps for blocks of cohorts: for the current ages 65-69, 70-79, and 80-89. The premiums are determined using the Wang Transform method; they are averaged over the ages in the cohort block. Table 4.4 shows both the resulting notionals and the associated premiums per cohort block. The difference between the notionals for males and females is striking, but not incomprehensible when comparing the pension rights (the main building blocks of the notionals) as shown in figure 4.1.

The proceeds and ECs of the swap portfolio are determined with these notionals; and these ECs are subtracted from those of the BOK contract, resulting in a new MVM. At the same time, the expected costs of the swaps are added to the expected cashflows stemming from the contract, resulting in a new liability. Table 4.5 lists the result of the hedge in the base case scenario. The Market Value Margin of the hedged contract is €33mln below the MVM of the contract without the swap; this is a reduction of 41%. However, the liability in the first year increases with almost €415mln. Investing in this portfolio of longevity swaps is not profitable when the best estimate mortality rates are realized.

Table 4.5 also shows how the Economic Capital in the first year changes as a result of the hedge. The EC for trend uncertainty of the swap appears to be larger than this number for the contract, indicating 'overhedging' for this source of risk. As negative ECs are not allowed (cf the EC for calamity risk), the resulting value is set to zero. The total EC decreases from €371mln to €101mln. This implies a large increase of the amount of free own funds in the surplus on the balance sheet. However, Nationale-Nederlanden does not pay or receive this amount of money; no cashflows are involved. Therefore, this outcome cannot be added to the profit of the hedge.

Table 4.4: Notionals and premiums in the cashflow hedge

Cohort	Males		Females	
	Notional	Premium	Notional	Premium
65-69	119	5.12%	8	4.94%
70-79	125	5.33%	10	4.96%
80-89	41	5.64%	6	5.06%

The notionals and premiums of the longevity swaps in the cashflow hedge per cohort block. The notionals are in €1mln.

Table 4.5: The cashflow hedge

	Contract	Swap	Contract with swap
EC Volatility	88.87	1.24	87.63
EC Calamity	0	0	0
EC Trend	269.59	339.70	0
EC Level	12.94	0	12.94
Total EC	371.40	340.94	100.63
Liability	12,587.02	415.18	13,002.20
MVM	79.98	–	47.11
Hedge result			-382.31

This table gives an overview of the ECs, the liability and the Market Value Margin, for the BOK contract, the portfolio of swaps under the cashflow hedge, and the combination of the two. Numbers are in €1,000,000.

Value hedge

The value hedge aims at removing the effect of changes in mortality rates on the level of the liability. This hedge includes swaps on all ages included in the contract, because the liability is built from all future cashflows [13]). The cohorts are grouped into seven blocks: for the cohorts currently aged 25-29, 30-39, 40-49, 50-59, 60-69, 70-79, and 80-89. The notional for each block is determined as the ratio the change in liability (dL) resulting from a very small shift in the mortality rates of the cohorts at hand (dq):[14]

$$Notional = \frac{dL}{dq} \tag{4.8}$$

This notional is set separately for each block of cohorts and for males and females. Because the swap hedges the liability under changed mortality rates, this would require different notionals each year. For clarity and comparability with the cashflow hedge, we use swaps with one payment moment, at the end of its 20-year term. The Wang Transform method is employed to determine the

[13] We do not discuss how the company should deal with the fact that for the cohorts younger than 65, a change in the liability is accompanied by a change in the cash balance instead of the liability itself.

[14] In theory, it might be more consistent with the financial theory (e.g. Black-Scholes) to divide this ratio by the change in the value of the swap for a change in the mortality rates. Because of the way the value of the swap is defined, dS/dq moves linear with the changes in the mortality rates. Therefore, it has no use to take this ratio into account.

Table 4.6: Notionals and premiums in the value hedge

Cohort	Males		Females	
	Notional	Premium	Notional	Premium
25-29	2,487	4.92%	18	4.91%
30-39	30,574	4.92%	47,946	4.91%
40-49	32,963	4.94%	49,719	4.92%
50-59	22,581	4.98%	34,614	4.92%
60-69	5,325	5.08%	6,499	4.93%
70-79	332	5.33%	498	4.96%
80-89	4	5.64%	7	5.06%

The notionals and premiums of the longevity swaps in the value hedge per cohort block. The notionals are in €1mln.

accompanying premiums (which are averaged over the cohorts). The mortality rates are assumed to increase by 0.01% each year as compared to the CBS best estimate mortality (hence dq is 1.0001 times the BEM in the first year, 1.0001^2 in the second year, etcetera). The resulting notionals and premiums are shown in table 4.6.

It is striking that the required notionals for the females in the contract are generally higher than those for the males, while they were only a fraction of the male notionals in the cashflow swap. This is explained by the lower mortality rates for females, and by the fact that extra decreases have more impact as females on average live longer. The table also shows that the notionals are higher for the mid cohort blocks, which makes sense as changes in the mortality rates for these participants are into effect for a longer time than the older blocks. The notionals for the youngest blocks are smaller as a result of the large period the cashflows are discounted over.

The calculations of the liability, EC and MVM are executed for both the contract and the swap. Table 4.7 lists the numbers. The MVM decreases with €23mln to €57mln, which is a reduction of 29%. At the same time, the liability increases with €70mln. Hence, also this hedge is not cost-effective: its yield is too low to cover the expected costs. The total result of the hedge is -€42mln. The reduction of the EC in the first year is €75mln, or 20% of its original value.

Comparison
Both the cashflow hedge and the value hedge are not cost-effective: the decrease in MVM is smaller than the increase in liability that comes with it in the base case scenario. The differences between the two variants are substantial though. The value hedge places notionals on all ages in the pension contract, while the cashflow hedge only includes the ages above 65. However, the reduction in MVM is by far larger in the cashflow hedge; 41% against 29%. The liability increases further as well (€415mln versus €70mln). This is related to the fact that the cashflow hedge intends to hedge the cashflows during the term of the swaps, while the value hedge is set up to hedge the change in the liability at the end of the term of the swap. Therefore, the cashflow hedge consists of swaps that pay out every year, and the swaps in the value hedge only at $t = 20$.

The cashflow hedge does create far more room for risky activities, by taking away 73% of the Economic Capital required in the first year. This is only 20% for the value hedge. However, the EC-results

Table 4.7: The value hedge

	Contract	Swap	Contract with swap
EC Volatility	88.87	0.22	88.65
EC Calamity	0	0	0
EC Trend	269.59	75.13	194.46
EC Level	12.94	0	12.94
Total EC	371.40	75.35	296.05
Liability	12,587.02	64.99	12,652.01
MVM	79.98	–	57.00
Hedge result			-42.01

This table gives an overview of the ECs, the liability and the Market Value Margin, for the BOK contract, the portfolio of swaps under the value hedge, and the combination of the two. Numbers are in €1,000,000.

of the cashflow hedge indicate overhedging, as the EC stemming from the swap portfolio in the first year is larger than the EC of the contract without the hedge. Overhedging also explains the extremely high costs of this hedging variant.

Both hedging methods are not desirable under the current assumptions. However, the value hedge is the least unattractive when considering the expected costs of the hedge in relation to the reduction of EC and MVM and the fact that the cashflow hedge overhedges the longevity risk.

4.3 Sensitivity analysis of the results of the hedge

The previous section resulted in the conclusion that hedging the group pension contract with the specific portfolios of longevity swaps is not recommendable. In this section we determine how this result changes when the assumptions underlying the calculations are altered. For clarity, we only look at the results of the hedging method that was determined to be the least unattractive in the previous section: the value hedge. We analyze the effects of the following factors.

- Changes in mortality rates;
- Changes in interest rates;
- Changes in the premium in the swap;
- Changes in the Cost of Capital in the MVM-calculations;
- Changes in the diversification factor in the MVM-calculations;
- A different average age of the participants in the swap;
- A larger shift in the determination of the swap notionals;
- Hedging one cohort block instead of all of them;
- Hedging with a portion of the indicated notionals.

Mortality rates
Figure 4.14 shows the change in MVM of the contract resulting from the hedge when the mortality rates are shifted, as well as the increase in liability. The value hedge has positive results (the decrease in MVM is larger than the increase in liability) when mortality rates change by more than 3% as compared to the current best estimate of the population mortality rates. The figure shows that the reduction in MVM is relatively stable over the mortality rate shifts. The change in the liability, hence the expected costs of the swap, determine the 'break-even point' where the hedge becomes cost-effective.

Interest rates
In the calculations, all future numbers are discounted with the (implied) interest rate structure of January 1, 2008. In order to determine the influence of the curve used, the calculations are repeated for different interest rates: they are assumed to equal a certain fraction of the January 1 (implied) structure.

Figure 4.15 shows the effect of changes in interest rates on the result of the hedge. When the interest rates are lower, the effect of discounting is decreased and hence the liability and the MVM increase. The 'double discounting' in the MVM somewhat diminishes this effect, resulting in a slightly concave result. However, despite the fact that a higher interest rate comes with less negative results, the hedge does not become cost-effective when interest rates change.

Figure 4.14: Results of the hedge for changing mortality rates

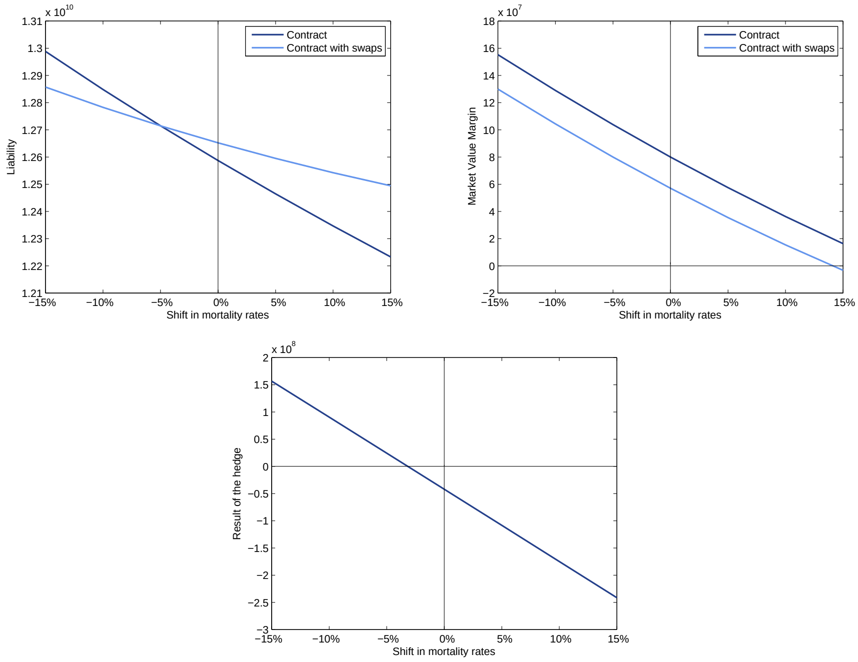

The figures show the impact of the value hedge on the liability and MVM of the BOK contract in the first year, for different shifts in mortality rates. The third graph shows the total result of the hedge: the decrease in the MVM minus the increase in the liability.

Figure 4.15: Results of the hedge for changing interest rates

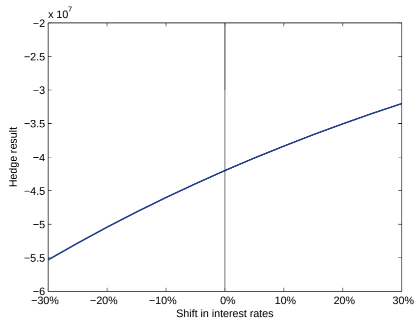

This figure shows the result of the value hedge (decrease in MVM minus increase in liability) for different shifts in interest rates.

Figure 4.16: Value hedge with Sharpe Ratio premiums

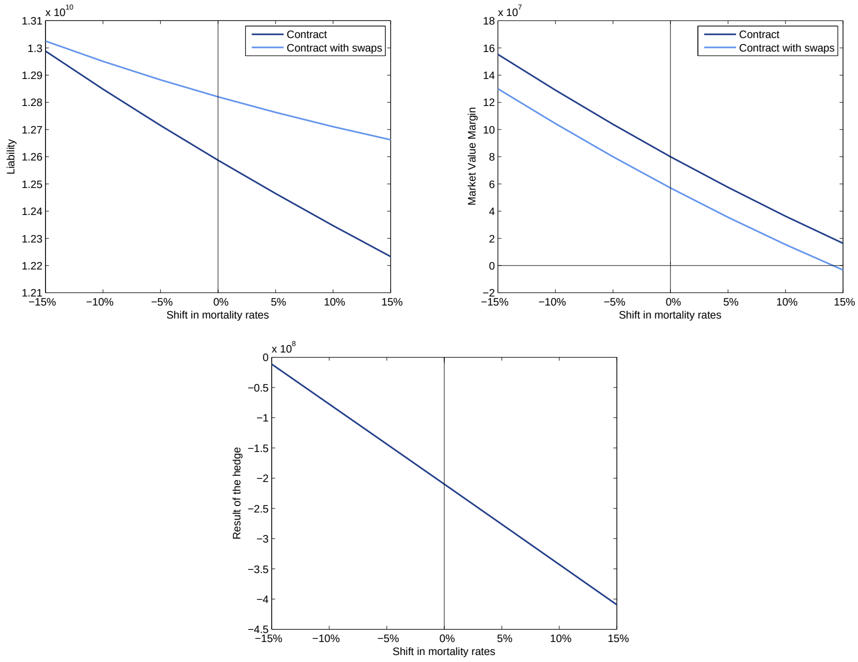

The figures show the impact of the value hedge on the liability and MVM of the BOK contract in the first year, for different shifts in mortality rates, when the premiums are determined using the Sharpe Ratio method. The third graph shows the total result of the hedge: the decrease in the MVM minus the increase in the liability.

Premium

The price of the swap is the premium π that is deducted from the expected mortality rates when the payments in the fixed leg are set. This premium was analyzed in chapter 3. Two methods were presented for its determination: the Wang Transform method and the Sharpe Ratio method. Until now, we worked with the premium following from the Wang Transform method in this section, which are around 5%. The Sharpe Ratio method results in substantially higher premiums: the average for these cohort blocks is 20%.

Figure 4.16 shows the results of the value hedge for changing mortality rates when the Sharpe Ratio method is used in the determination of the premium. The change in MVM hardly differs from the picture with the Wang premiums (the decrease is slightly stronger), however, the increase in liability is extreme. The impact of the premium is very large: with the higher premiums, the result of the hedge has worsened such that even also when the mortality rates appear to be far below the best estimate, the hedge is not cost-effective. Figure 4.17 shows how the result of the hedge changes for different premiums. We assume a flat premium for all cohort blocks, ranging from 0% to 25%. The strong influence of the premium is underlined. When the flat premium is below 3%, the hedge is cost-effective.

Figure 4.17: Results of the hedge for a changing premium

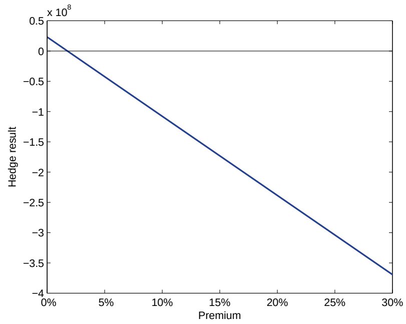

This figure shows the result of the value hedge when the premium is varied from 0% to 30%.

Cost of Capital

The price of longevity risk (Market Value Margin) is calculated by discounting the return (Cost of Capital) required over the capital set aside in the future to cover the longevity risk (Economic Capital). At ING, the Cost of Capital is fixed at 4%. It would make sense when a lower Cost of Capital leads to better results, as less return is required. However, the outcomes are opposite. Figure 4.18 shows that a higher CoC comes with a higher MVM, and also with a larger reduction of the MVM as a result of the swap. Of course, the increase in the liability is not changed. The result of the hedge is therefore better for a higher Cost of Capital; when the Cost of Capital is 11% or higher, the hedge is even cost-effective. A CoC of 11% is not realistic though, given the current 4%.

Diversification factor

The MVM calculations also include the diversification factor; a factor that should correct for the fact that the separate shocks in the EC calculations will not take place at the same time. The prescribed factor is 40%, indicating that only 40% of the current and future Economic Capital is taken into account in the determination of the Market Value Margin. Figure 4.19 shows the result of the hedge (decrease in MVM minus increase in liability) when the diversification factor is varied between 0% and 80%. The influence is linear, however the hedge is still not cost-effective, in all cases.

Figure 4.18: Results of the hedge for a changing Cost of Capital

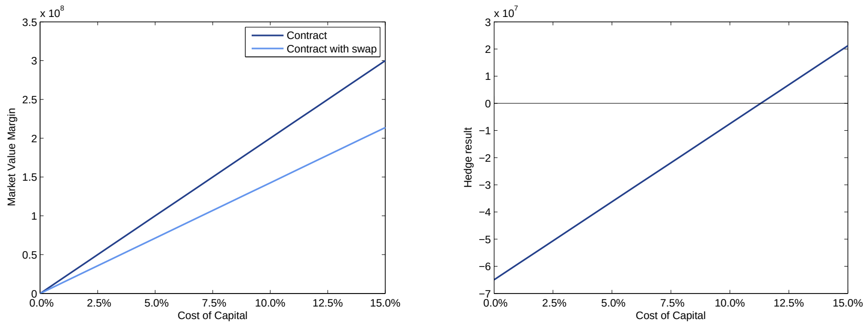

The left figure shows the MVM of the contract and the contract with the swap for different values of the Cost of Capital. The right figure shows the total result of the value hedge.

Figure 4.19: Results of the hedge for a changing diversification factor

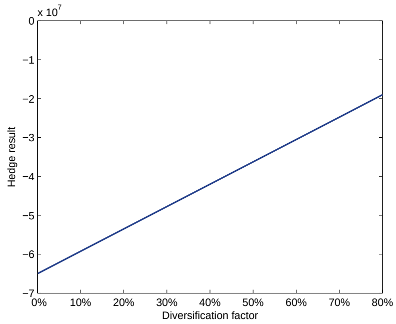

The figures show the effect of different values of the diversification factor on the result of the value hedge.

Age of the participants

In section 4.1, we used an old and a young BOK contract to explain the patterns found in the EC and MVM of the standard contract. The MVM turned out to be lower when the participants in the contract are older, as a result of the shorter time of high ECs. Figures 4.20 show what happens when the hedge is applied to the two contracts (the notionals and premiums are recalculated).

The liability and the MVM are are both proportionally higher for the younger contract, hence the difference between them is also larger. The young contract is somewhat more sensitive to changes in mortality rates, but the differences are marginal. The effect of the age of the participants in the contract is hence not very important.

Notionals

The notionals in the value hedge are determined by the change in the liability when the mortality rates change by a very small number each year. The shock scenarios of the Economic Capital are however based on larger shifts in the mortality rates. When the relationship between the mortality rates and the level of the liability is not linear, a larger shift might therefore give better results. By taking a larger shift, part of the convexity is taken into account. In the calculations so far, an 0.01% annual increase in mortality rates is assumed. Table 4.8 lists the results of the hedge in the base case scenario, for both

Figure 4.20: Hedge results for the young and the old contract

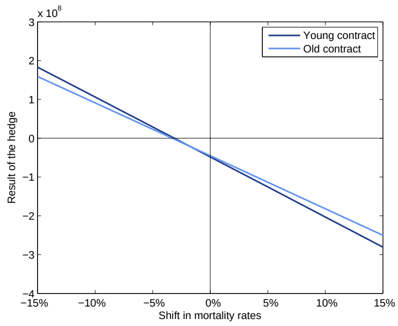

The result of the hedge in the contract with relatively young and relatively old participants, for different shifts in mortality rates.

Table 4.8: The value hedge for different shifts in the notional determination

	Contract	Contract with swaps (0.01%)	Contract with swaps (0.1%)
EC Volatility	88.87	88.65	88.66
EC Calamity	0	0	0
EC Trend	269.59	194.46	195.36
EC Level	12.94	12.94	12.94
Total EC	371.40	296.05	296.96
Liability	12,587.02	12,652.01	12,651.22
MVM	79.98	57.00	57.28
Hedge result		-42.01	-41.50

This table gives an overview of the ECs, the liability and the Market Value Margin, for the BOK contract, the BOK contract when hedged with swaps with notionals determined by an annual shock to the mortality rates of 0.01%, and the BOK contract when the notionals of the swaps are determined by shocks of 0.1%. Numbers are in €1,000,000.

the hedge with swaps with notionals of 0.01% annual shifts, and the hedge with swaps with notionals of 0.1% annual shifts, hence ten times larger. The reduction in EC and MVM and the increase in the liability are all slightly smaller, resulting in a marginally better result of the hedge. The influence of the shifts used in the determination of the notionals of the swaps is hence very small.

Cohort block

In the calculations, we hedged the entire range of ranges of the participants in the swap. However, it is possible that hedging part of the cohort blocks is cost-effective, but that this result is more than compensated by other cohort blocks with strongly negative results. Table 4.9 lists the results of the hedge when only one of the cohort blocks is hedged instead of all of them. It appears that none of the cohort blocks is cost-effective, although the younger cohorts show a relatively larger decrease in MVM and a relatively smaller increase in liability. The better results for the younger cohorts might also indicate an extra reason why the cashflow hedge performs worse than the value hedge, as the former only includes the older cohorts.

Partial hedge

The last factor considered is hedging all cohort blocks, but only with a certain portion of the calculated notionals. The results of this exercise are shown in figure 4.21. It is clear that partial hedging does not make the hedge cost-effective: the negative result of the full hedge is decreased proportionally.

Table 4.9: The result of the hedge when only one cohort block is hedged

	25-29	30-39	40-49	50-59	60-69	70-79	80-89
Increase in liability	0.081	4.737	12.626	26.350	17.211	3.881	0.104
Decrease in MVM	0.044	1.964	5.421	8.996	5.209	1.344	0.001
Hedge result	-0.037	-2.772	-7.205	-17.354	-12.002	-2.536	-0.103

This table gives an overview of the result of the value hedge when only one of the cohort blocks is hedged. Numbers are in €1,000,000.

Figure 4.21: Results of partial hedging

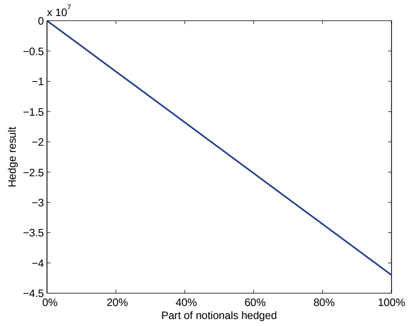

The result of the value hedge when only a certain part of the notionals is engaged in.

Conclusion
The longevity hedge as determined earlier is not cost-effective. This subsection examined the sensitivity of this result to changes in several factors.

- Shifts in mortality rates strongly influence the effect of the hedge. The costs increase strongly when mortality rates appear to be higher than estimated, but when they are more than 3% below it, the hedge is cost-effective.
- Changes in interest rates affect the magnitude of the hedge result, but do not make it profitable.
- The swap premium is an important determinant of the hedge results. Premiums under 2% result in cost-effectiveness.
- When the Cost of Capital is higher than 11% instead of the current 4%, the hedge would be attractive. Such a high CoC is not realistic though.
- Changing the diversification factor does not make the hedge profitable.
- The age of the participants in the swap hardly changes the hedge effectiveness.
- When the notionals of the swaps are determined by taking larger shifts in mortality rates, the result changes only marginally.
- None of the cohort blocks can be hedged in a cost-effective way; it has no use hedging only one or a few age groups.
- Partial hedging leads to proportional results.

We therefore conclude that only two factors can reasonably make the hedge cost-effective. Decreases in mortality rates, which is not surprising, as the hedge is set up to mitigate just that, and a lower premium, which determines the costs for the receiver party in the swap.

5 Conclusion

Does a specific portfolio of longevity swaps lower the longevity risk incorporated in group pensions? In order to answer this central question, this thesis applied swaps on the Dutch longevity indices to a standard group pension contract of Nationale-Nederlanden. The risk that mortality rates turn out to be lower than anticipated on is very important for a pension provider like Nationale-Nederlanden, as it cannot be diversified away while the potential losses are practically endless.

With the emerging market for longevity swaps, a new means to hedge this source of risk becomes available. The receiver in these derivatives is compensated for decreases in mortality rates beyond the expected decreases minus a certain margin (premium), while he has to pay money when the actual mortality rates turn out to be higher than the expectation minus the margin. Hence, a portfolio of receiver longevity swap enables a pension provider to fix its future mortality rates (as far as these are covered by the swap notionals). As measures of longevity risk are based on expected mortality rates in several shock scenarios, they change when this hedge is set in place. The most relevant risk measure in this respect is the Economic Capital for life risk: the risk capital demanded over the longevity risk in the contract (as determined by shock scenarios). The associated price measure is the Market Value Margin (MVM) for life risk. This is the cost of capital required over the current and future risk capital (ECs). But besides lowering the EC and the MVM, the hedge is expected to require cash in the future, as the margin is set at the benefit of the payer. The hedge is only attractive when the decrease in MVM is larger than this increase in the liability.

The effect of hedging with longevity swaps is determined by calculating the risk of a standard group pension contract with and without this hedge. When the notionals of the swaps are determined such that they cover as many changes in the liability of the contract as possible (value hedge), the MVM decreases by 29%, but the absolute decrease is smaller than the accompanying increase in the liability. A hedge in which the notionals aim to cover changes in cashflows results in a stronger decrease of the MVM of 41%, but also a disproportional stronger increase of the liability. At the same time, the reduction of the EC in the first year is 20% for the value hedge, and 73% for the cashflow hedge. However, the most important source of Economic Capital is overhedged; the decrease is larger than the original value. This overhedging and the high costs make the cashflow hedge the least attractive alternative of the two.

A change in several factors in the hedge calculations indicate that only two of them are able to make hedging cost-effective. The first is the level of the premium added to the swap to the benefit of the payer; when it is below 2%, the total result of the hedge is positive. The other factor is a shift in mortality rates, which of course the hedge is designed to cover for. When the mortality rates decrease by more than 3%, it is profitable to engage in the proposed hedging strategy.

It is however striking that the hedge takes away only 20% of the Economic Capital in the first year and 29% of the Market Value Margin, instead of all of it. This could be explained partially by the fact that the swaps do not affect all ECs; only volatility risk and trend uncertainty risk can be decreased, the level uncertainty risk remains. Besides, the swaps are placed on CBS mortality rates, while the EC measures use the CRC rates (corrected for experienced mortality). It remains open to question how a larger part of the longevity risk could be taken away.

Apparently, companies who have engaged in longevity swap transactions base their analysis on other measures for the price of risk. The possible losses from longevity might be larger for a stand-alone pension fund; when necessary, Nationale-Nederlanden could (temporarily) cover losses on older contracts by the incoming cashflows on contracts with the working people. Also, especially in the current period of scarcity of money, it might be more important to have as much funds available for investments as possible.

It will be very interesting to repeat this analysis in a few years, when the market is more mature and more information on the pricing is available. That information will also enable an examination of the hedge effectiveness when the environment changes during the term of the swap. Besides, a more mature market might give rise to yet a new derivatives market: longevity swaptions. These products could be more attractive, as their costs are limited while their benefits equal those of the swaps.

6 References

Actuarieel Instituut (1998). Inleiding Verzekeringstechniek, *Actuarieel Rekenaar module 3*.

Bauer, D. and J. Ruß (2006). Pricing longevity bonds using implied survival probabilities, Working Paper, Ulm University.

Beelders, O. and D. Colarossi (2004). Modelling mortality risk with extreme value theory: The case of Swiss Re's mortality-indexed bonds, *Global Association of Risk Professionals*, July/August, Issue 19, pp.26-30.

Beer, J. de (2006). Verleden en toekomst van de levensverwachting, *De Actuaris*, May, pp.26-30.

Blake, D., A.J.G. Cairns, and K. Dowd (2006). Living with mortality: Longevity bonds and other mortality-linked securites, *Presented to the Faculty of Actuaries*.

Blake, D. (2008). Survivor products: Managing longevity risk & mortality improvements, available at $http://siteresources.worldbank.org/INTLACREGTOPFINSECDEV/Resources/SurvivorProductsBlake.ppt$, Pensions Institute.

Booth, H., J. Maindonald and L. Smith (2002). Age-Time interactions in mortality projection: Applying Lee-Carter to Australia, *Population Studies*, vol.56 No.3, pp.325-336.

Broekhoven, H. van (2002). Market value of liabilities mortality risk: a practical model, *North American Actuarial Journal*, vol.6 no.2, pp.95-106.

Buckley, A. (2004). *Multinational finance*, fifth edition, pp.331-335, Financial Times - Prentice Hall.

Cairns, A.J.G., D. Blake, P. Dawson, K. Dowd (2005). Pricing the risk on longevity bonds, *Life and Pensions*, October, pp.41-44.

CEA (2008). Results and discussion on the impact assessment of the future Solvency II framework on insurance products and markets, available at $http://www.tsrsb.org.tr/NR/rdonlyres/FD2B8B3D-023A-4812-A7E3-6F090F653BDB/2872/MicrosoftWordCEA impactassessmentsurveyresults.pdf$, April 21.

Corporate Insurance Risk Management (2006). Economic Capital for Insurance Life Risk: Standard of Practice, *ING internal document*.

Corporate Insurance Risk Management (2007). Economic Capital standards: Standard of Practice, *ING internal document*.

Coughlan, G., D. Epstein, A. Ong, A. Sinha, J. Hevia-Portocarrero, E. Gingrich, M. Khalaf-Allah, P. Joseph (2007a). *LifeMetrics: A toolkit for measuring and managing longevity and mortality risks (LifeMetrics Technical Document)*, JPMorgan.

Coughlan, G., D. Epstein, A. Sinha and P. Honing (2007b). *Q-forwards, Derivatives for transferring longevity and mortality risk*, JPMorgan.

Coughlan, G. (2007). Longevity & mortality risk transfer via the capital markets, available at $http://www.icms.org.uk/downloads/QRM_talks/Guy_Coughlan.ppt$, JPMorgan.

Cox, T. (2005). Longevity protection, *De Actuaris*, September, pp.34-35.

Elandt-Johnson, R.C. and N.L. Johnson (1980). *Survival models and data analysis*, John Wiley & Sons, Inc., New York.

Gregorkiewicz, F. and R. Plat (2004). Stochastische prognosemodellen voor sterfte, *De Actuaris*, Juli, pp.18-21.

Heijningen, R. van (2006). *Sterftederivaten ter afdekking van het langlevenrisico*, ING internal document.

Hull, J.C. (2006). *Options, futures and other derivatives*, 6th edition, Pearson Education Inc., Upper Saddle River, New Jersey.

ING (2008). De geschiedenis van ING, available at http://www.ing.com.

Jong, A. de and A. van der Meulen, A (2005). Prognose van sterfte naar doodsoorzaken: model en veronderstellingen, *Bevolkingstrends* 53(2), pp.50-62, Centraal Bureau voor de Statistiek.

JPMorgan (2007). Life Metrics Index Netherlands, available at http://www.lifemetrics.com.

Lee, R.D. and Carter, L. (1992). Modeling and forecasting U.S. mortality, *Journal of the American Statistical Association*, No.87, pp.659-675.

Life & Pensions (2008). Lift-off for longevity market, *Life & Pensions*, March.

Lin and Cox (2005). Securitization of mortality risks in life annuities, *Journal of Risk & Insurance*, vol.72 No.2, pp.227-252.

Loeys, J., N. Panigirtzoglou and R.M. Ribeiro (2007). *Longevity: A market in the making*, JPMorgan Global Market Strategy.

Nationale-Nederlanden (2008). WTV Jaarverslag Nationale-Nederlanden Levensverzekering Maatschappij 2007, available at http://www.nn.nl.

Neville, L. and H. Ho (2006). Buying in bulk, *Risk Magazine* vol.19 no.5.

Olivier, P. and T. Jeffery (2004). Stochastic mortality models, *Presented to the Society of Actuaries Ireland*.

Pelsser, A.A.J. (2007). On the applicability of the Wang Transform for pricing financial risks, *accepted for publication in ASTIN Bulletin*.

Pengelly, M. (2008). Longevity: A long way to go, *Risk Magazine*, Nordic Risk, Winter.

Risk Magazine (2008). Longevity and mortality risk, available at http://http://db.riskwaters.com/ public/show

Page.html?page=480659, April 21.

Srinivasan, S. (2004). *CPI futures at the CME*, Chicago: Chicago Mercantile Exchange.

Tornij, E. (2006). Pensioentafel 2006 houdt rekening met trends, *De Actuaris*, Mei, pp.32-34.

Wackerly, D.D., W. Mendenhal III and R.L. Scheaffer (2002). *Mathematical Statistics with Applications*, 6th edition, Duxbury Thomson Learning, Pacific Grove (California).

Wang, S.S. (2002). A universal framework for pricing financial and insurance risks, *Astin Bulletin*, vol.32 no.2, pp. 213-234.

A Mortality rate measures

Mortality rates can be presented using several measures. The most commonly used measures are listed below, along with their mutual relationships. (Actuarieel Instituut, 1998; Coughlan et al, 2007a, Elandt-Johnson and Johnson, 1980)

- **Initial / annual rate of mortality ($q(x)$)**
  The probability that a person aged x dies within the next year: the number of deaths of people aged x (at the start of the year) over the year ($d(x)$) divided by the number of people alive aged x at the start of the year ($l(x)$).
  $$q(x) = \frac{d(x)}{l(x)} \tag{A.1}$$
  These numbers are corrected for migration. When they are also smoothed across ages to reduce the influence of noise and outliers, the probabilities are called graduated initial rates of mortality.

- **Initial /annual rate of survival ($p(x)$)**
  The probability that a person aged x survives over the upcoming year.
  $$p(x) = \frac{l(x+1)}{l(x)} \tag{A.2}$$
  The relationship between $q(x)$ and $p(x)$ hence is:
  $$p(x) = 1 - q(x)$$

- **Central observed rate of mortality ($m(x)$)**
  Deaths per unit of exposure over the past year. It is calculated by dividing the number of deaths over the year in the group of people aged x at the beginning of the year ($d(x)$) by the total number of years lived by all members of the group ($L(x)$).
  $$m(x) = \frac{d(x)}{L(x)} \tag{A.3}$$
  $L(x)$ is composed of the number of people surviving the whole year and the average fraction of the year lived by the people who died during the year ($f(x)$).
  $$\begin{aligned} L(x) &= l(x+1) + f(x) \cdot d(x) \\ &= [l(x) - d(x)] + f(x) \cdot d(x) \\ &= l(x) - [1 - f(x)]d(x) \\ \Leftrightarrow l(x) &= L(x) + [(1 + f(x))d(x)] \end{aligned}$$
  As the number of deaths ($d(x)$) equals the number of people alive at the beginning of the year ($l(x)$) times the probability that a person aged x dies within the upcoming year ($q(x)$), the relationship between $m(x)$ and $q(x)$ is as follows:
  $$m(x) = \frac{d(x)}{L(x)} = \frac{d(x)}{l(x) - [1 - f(x)]d(x)} = \frac{q(x)}{1 - [1 - f(x)]q(x)}$$
  And the other way around, because $d(x) = L(x) \cdot q(x)$:
  $$q(x) = \frac{d(x)}{l(x)} = \frac{d(x)}{L(x) + [1 + f(x)]d(x)} = \frac{m(x)}{1 + [1 + f(x)]m(x)}$$

When the deaths are assumed to be distributed uniformly across the year, the fraction of year the nonsurvivers live ($f(x)$) equals 0.5, and the relationship between $m(x)$ and $q(x)$ simplifies to

$$m(x) = \frac{q(x)}{1 - 0.5q(x)}$$
$$q(x) = \frac{m(x)}{1 + 0.5m(x)}$$

Alternatively, an exponential distribution can be assumed. This would imply that

$$q(x) = 1 - \exp\left(-m(x)\right)$$
$$m(x) = -\ln[1 - q(x)]$$

- **Force of mortality ($\mu(x)$)**
  The instantaneous death rate for lives aged x exactly:

  $$\mu(x) = -\frac{d\log[l(x)]}{dx} \tag{A.4}$$

  Equivalently, $\mu(x+t)$ is the instantaneous death rate for individuals aged $x+t$ exactly at time t, where t can be a fractional year. The number of lives aged $x+t$ exactly at time t is $l(x+t)$. Therefore, $\mu(x)$ and $q(x)$ are related in the following way (integral is over the upcoming year):

  $$\int_0^1 l(x+t) \cdot \mu(x+t)dt = l(x) \cdot q(x)$$

- **Life expectancy ($e(x)$)**
  The average future lifetime of an individual aged x in years. It can be written as

  $$e(x) = \sum_{t=1}^{\omega - x} (t)p(x) \tag{A.5}$$

  where ω is the maximum attainable age, usually set at 120 years[15], and $(t)p(x)$ is the relevant survival rate defined as the probability that a person aged x survives t more years:

  $$(t)p(x) = \prod_{s=0}^{t-1} (1 - q(x+s))$$

  There are two versions of life expectancy:

  - period life expectancy: future mortality changes are not incorporated in survival rates (the age 90 mortality rate that is assumed in 70 years time (for today's 20 year olds) is the same as that for a 90 year old today);
  - cohort life expectancy: future mortality changes are incorporated, using a separate set of survival rates for each year of each birth cohort. These rates are usually tabled in a generation table, where the rows stand for the age and the columns for the year. In order to obtain the survival rates for a certain age group, the relevant diagonal is selected. This is the most common (and specific) way to determine life expectancies.

[15] The official and generally accepted maximum age is 122 years. This age was attained by Jeanne Calment, a French woman who died in 1997. (De Beer, 2006)

B Determining the MVM for a young and an old contract

This appendix shows the building blocks for the determination of the Market Value Margin for a group pension contract with relatively young and one with relatively old participants. It gives the results of the calculations expounded in section 4.1 for these contracts. The graphs and tables are used in that section to help in explaining the results of the calculations for the standard BOK contract.

Figure B.1: Initial pension rights in the young and old group pension contract per cohort

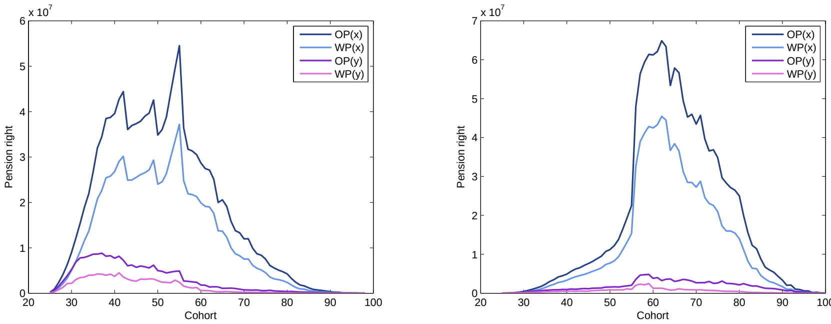

The left panel of the figure shows the liabilities that are involved in the young pension contract, the right panel shows the same for the old pension contract. OPx *and* OPy *stand for the rights built up for old-age pensions of men (*x*) and women (*y*),* WPx *and* WPy *is the associated widow pension.*

Table B.1: Composition of the pension rights for the young, standard and old BOK contract

	OP(x)	WP(x)	OP(y)	(WP(y)
Young contract	53%	35%	8%	4%
Standard contract	54%	36%	7%	3%
Old contract	57%	37%	5%	2%

The table states how the three BOK contracts are composed. Each percentage is the portion of pension rights for the specific type of pension.

Figure B.2: Present values of the future cashflows from the young contract

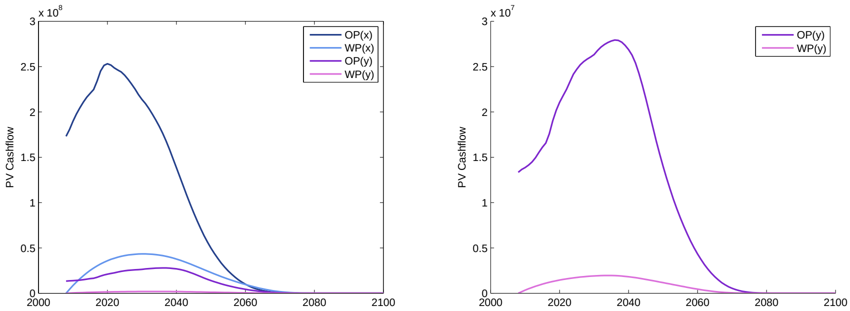

(a) Present value of the cashflows per future year

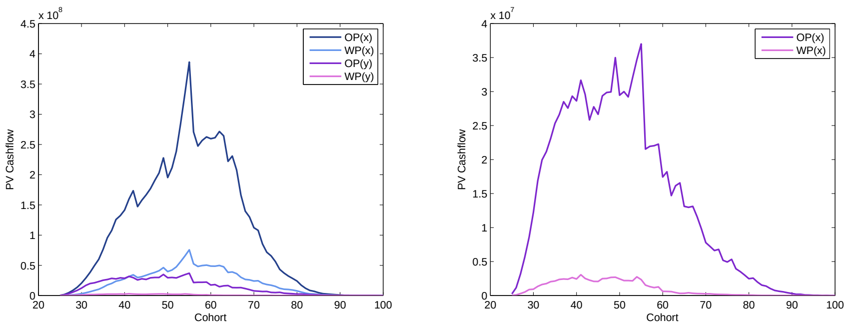

(b) Present value of the cashflow per cohort

This figure shows the pattern of the future pension payments stemming from a BOK contract consisting of relatively young participants; OP stands for old-age pensions, WP for widow pensions. The left figure of panel a shows the cashflows stemming from the contract in each future year, the right figure depicts the same for only the female pensions (as their pattern is not clear in the main figure). Panel b shows the same per cohort (the current age of the participants). All numbers are discounted to 2008.

Figure B.3: Present values of the future cashflows from the old contract

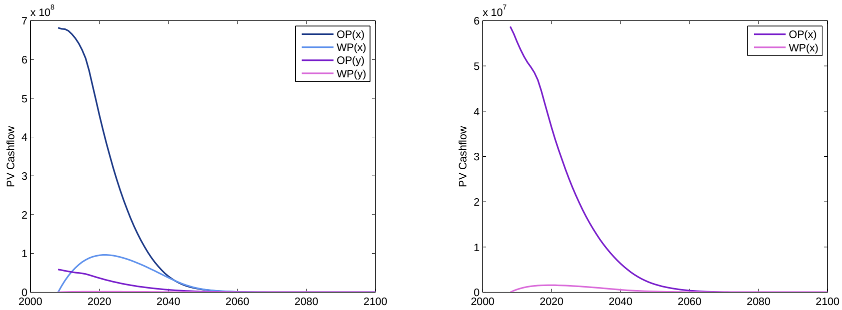

(a) Present value of the cashflows per future year

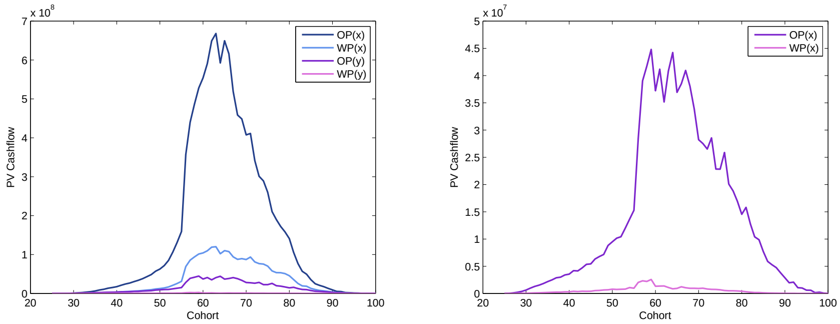

(b) Present value of the cashflow per cohort

This figure shows the pattern of the future pension payments stemming from a BOK contract consisting of relatively old participants; OP stands for old-age pensions, WP for widow pensions. The left figure of panel a shows the cashflows stemming from the contract in each future year, the right figure depicts the same for only the female pensions (as their pattern is not clear in the main figure). Panel b shows the same per cohort (the current age of the participants). All numbers are discounted to 2008.

Table B.2: Overview of the EC in the young BOK contract

		BEM	Volatility	Calamity	Trend	Level	Total
Liability	OP(x)	8,073		8,018	8,326	8,096	
	WP(x)	1,596		1,631	1,548	1,583	
	OP(y)	1,036		1,032	1,063	1,036	
	WP(y)	72		74	70	72	
	Total liability	10,778		10,754	11,008	10,787	
Economic Capital (mln)			112	0	230.6	9.7	352.3
Economic Capital (%)			1.04%	0%	2.14%	0.09%	3.27%

The Economic Capital for the different sources of risk of the young BOK contract in €mln, where the CRC(EM) table is the BEM table (the CRC table corrected for experienced mortality).

Table B.3: Overview of the EC in the old BOK contract

		BEM	Volatility	Calamity	Trend	Level	
Liability	OP(x)	11,896		11,700	12,252	11,931	
	WP(x)	2,500		2,584	2,461	2,484	
	OP(y)	1,029		1,013	1,052	1,029	
	WP(y)	41		42	41	41	
	Total liability	10,778		10,754	11,008	10,787	
Economic Capital (mln)			36	0	340	19	395
Economic Capital (%)			0.23%	0%	2.2%	0.12%	2.55%

The Economic Capital for the different sources of risk of the old BOK contract in €mln, where the CRC(EM) table is the BEM table (the CRC table corrected for experienced mortality).

Figure B.4: Liability in time

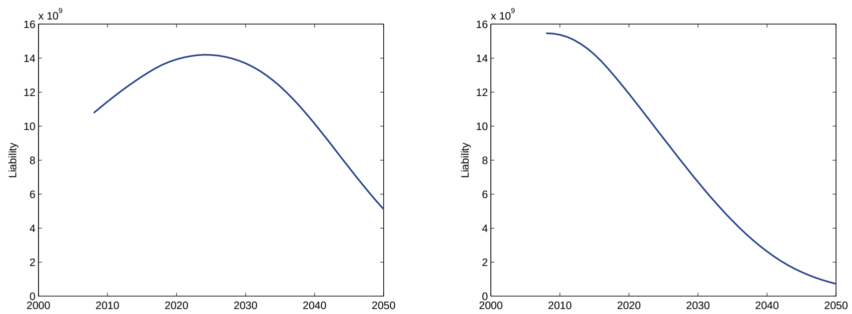

The development of the liability in time is shown for the young BOK contract (left) and its old equivalent (right).

Figure B.5: Economic Capital in time for the young BOK contract

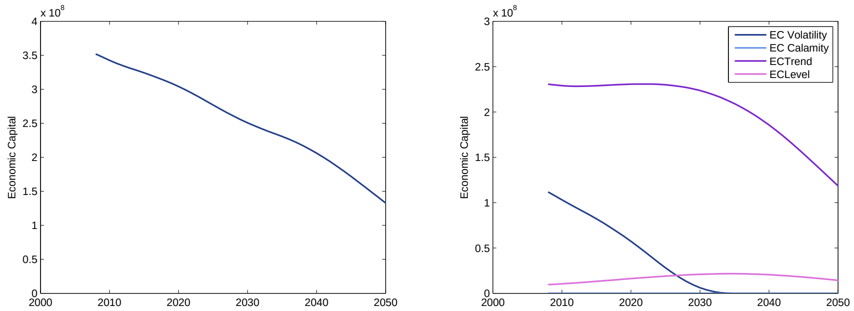

The left figure shows the total Economic Capital over the young BOK contract in time, the right figure shows its components.

Figure B.6: Economic Capital in time for the old BOK contract

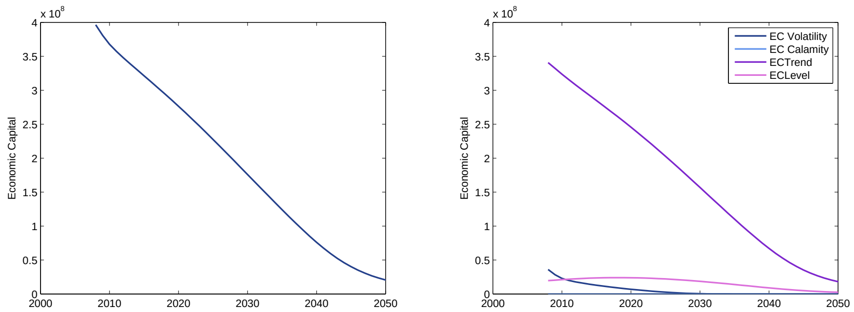

The left figure shows the total Economic Capital over the old BOK contract in time, the right figure shows its components.

Figure B.7: Market Value Margin in time for both contracts

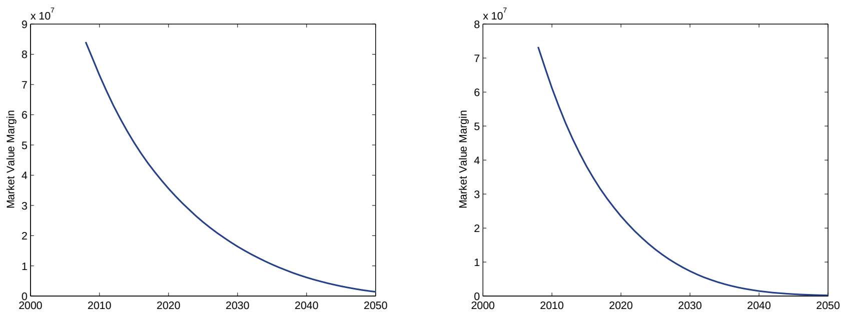

The Market Value Margin in time for the young BOK contract (left) and the old contract (right).