The economics of sharing macro-longevity risk*

Dirk Broeders[a,c], Roel Mehlkopf[b,d], and Annick van Ool[a,c]

[a] *De Nederlandsche Bank*

[b] *Cardano Risk Management*

[c] *Maastricht University*

[d] *Tilburg University*

February 2019

Abstract:

Pension funds face macro-longevity risk or uncertainty about future mortality rates. We analyze macro-longevity risk sharing between cohorts in a pension scheme as a risk management tool. We show that both the optimal risk-sharing rule and welfare gains from risk sharing depend on the retirement age policy. Welfare gains from sharing macro-longevity risk measured on a 10-year horizon in case of a fixed retirement age are between 0.2 and 0.3 percent of certainty equivalent consumption after retirement. Cohorts experience, in this case, a similar impact of macro-longevity risk on post retirement consumption and it is not optimal for young cohorts to absorb risk of old cohorts. However, when the retirement age is fully linked to changes in life expectancy, welfare gains are substantially higher. The risk bearing capacity of workers is larger when they use their labor supply as a hedge against macro-longevity risk. As a result, workers absorb risk from retirees in the optimal risk-sharing rule, thereby increasing the welfare gain up to 2.7 percent.

Keywords: Macro-longevity risk, risk sharing, welfare analysis, retirement age

JEL classifications: D61, G22, J26, J32

* Email addresses: Annick van Ool (corresponding author) a.w.m.van.ool@dnb.nl, Dirk Broeders d.w.g.a.broeders@dnb.nl and Roel Mehlkopf r.j.mehlkopf@uvt.nl.

We thank Rob Bauer, Servaas van Bilsen, Antoon Pelsser, Peter Schotman, Ralph Stevens, Sweder van Wijnbergen, Willem Jan Willemse and participants of research seminars at DNB, Maastricht University and Netspar for useful comments.

1 Introduction

Macro-longevity risk is the uncertainty about future mortality rates. Mortality rates may for example decrease as a consequence of medical improvements, or may increase because of new diseases. Macro-longevity risk is a systematic risk. It affects the entire population. As a consequence, macro-longevity risk does not decrease by sharing it within a pool of participants of the same cohort. Nonetheless, sharing macro-longevity risk with other cohorts can be beneficial as the risk affects cohorts differently. Macro-longevity risk differs from micro-longevity risk or the individual uncertainty about the time of death. Micro-longevity risk is an idiosyncratic risk that can be fully diversified by pooling enough participants in a pension scheme.

Macro-longevity risk has a significant impact on pension provisioning, depending on the configuration. A defined benefit (DB) pension scheme protects beneficiaries against macro-longevity risk. The risk however increases the uncertainty in the funding ratio. These changes in the funding ratio are borne by the employer and employees that contribute to the pension scheme. In a defined contribution (DC) pension scheme with a fixed annuity, pension benefits are guaranteed after retirement and macro-longevity risk is borne by the pension provider, for example the shareholders of an insurance company. In a DC pension scheme with a variable annuity pension benefits are adjusted to changes in future mortality rates. As a consequence, participants bear macro-longevity risk themselves. Retirees are especially vulnerable to macro-longevity risk because they cannot compensate lower pension benefits by working longer or saving more. A decrease in mortality not only affects retirees. Also workers might be affected through a decrease in future pension benefits or an increase in contributions to finance a longer retirement period. Hence, macro-longevity risk affects both retirees and employees. However, the risk does not affect all cohorts in the same way or by the same amount. Medical progress or diseases may affect cohorts in a different way. Furthermore, workers have more risk-absorbing capacity compared to retirees as they can adjust their labor supply in response to changes in mortality. These differences create a clear case for risk sharing. This case is strengthened by the fact that the market for macro-longevity risk is close to non-existent.

The central economic problem this paper addresses is optimal risk sharing between cohorts in a pension scheme. Collective risk sharing is a risk management tool that allocates risks to cohorts. To allocate macro-longevity risk optimally, we maximize aggregate expected utility of all participants in the situation where a social planner is present. This social planner makes decisions on behalf of the participants. In this way, we derive the Pareto-optimal risk-sharing rule and calculate the welfare gain of the Pareto improvement. We determine a fair risk compensation for cohorts who absorb macro-longevity risk of other cohorts using a utility-based fairness criterion such that all participants experience the same welfare gain.

We find that the design of the retirement age policy has a large impact on the optimal risk-sharing rule and size of welfare gains. If the retirement age is fixed, welfare gains from sharing macro-longevity risk measured on a 10-year horizon are between 0.2 and 0.3 percent of certainty equivalent consumption after retirement. In this case, the impact of macro-longevity risk on consumption after retirement is more or less equal for different cohorts. As a consequence, young cohorts do not absorb macro-longevity risk of other cohorts and welfare gains from sharing macro-longevity risk are limited. If the retirement age is linked to life expectancy by contrast, welfare gains from sharing macro-longevity risk are substantially higher, up to 2.7 percent. The risk bearing capacity of workers is larger, because they can use their labor supply as a hedge against macro-longevity risk. As a result, workers absorb risk from retirees. After all, human capital of workers increases if they work longer. Moreover, a positive risk compensation is not required for young cohorts to absorb risk of retirees.

This paper contributes to the literature on macro-longevity risk. We approach this actuarial topic from an economic perspective. To the best of our knowledge, this paper is the first to investigate Pareto-optimal risk-sharing rules of macro-longevity risk. Related papers are De Waegenaere et al. (2017), De Waegenaere et al. (2018) and papers considering group self-annuitisation schemes (GSAs), for example Piggott et al. (2005), Qiao and Sherris (2013) and Boon et al. (2017). These papers investigate sharing micro- and macro-longevity risk. In GSAs longevity risk is shared uniformly among participants in a pool. De Waegenaere et al. (2017) and De Waegenaere et al. (2018) consider ad hoc risk-sharing rules for micro- and macro-longevity risk. We consider macro-longevity risk only and determine the optimal risk-sharing rule. Moreover, we include a risk compensation which is not the case in the above mentioned papers. This paper also contributes to the literature on risk sharing. Most papers on Pareto-optimal risk sharing focus on financial risks, e.g., Gollier (2008), Cui et al. (2011) and Bovenberg and Mehlkopf (2014). We determine the Pareto-optimal risk-sharing rule for a non-financial risk, namely macro-longevity risk. Finally, we are the first to investigate the impact of different retirement age policies on sharing macro-longevity risk. Investigating different retirement age policies is relevant, as several countries link the retirement age to life expectancy. Stevens (2017) investigates the impact of retirement age policies on the individual retirement age, expected remaining lifetime at retirement and value of pension benefits, but does not consider collective risk sharing.

There are multiple alternative ways to manage macro-longevity risk. First, insurance is a risk management method in which a third party guarantees to compensate specified losses in return for a levy. For example, macro-longevity risk can be transferred to institutional investors via financial products, traded in financial markets. This process of securitization is described by Cairns et al. (2006a), Blake et al. (2006a), Ngai and Sherris (2011) and Hunt and Blake (2015). Securitization can improve welfare as it achieves a more efficient risk allocation by distributing the risk among market participants who are better able to bear the risk. Second, literature suggests that governments can establish solutions to manage macro-longevity risk by issuing longevity bonds (Brown and Orszag (2006), Blake et al. (2014)). In practice, the amount of financial products that transfer macro-longevity risk is small (Basel Committee on Banking Supervision (2013)) and insurance companies and governments are reluctant to underwrite macro-longevity risk. There are several reasons for the lack of a well-functioning market, e.g. the existence of basis risk and the fact that the government is not a natural issuer of longevity bonds because it is already exposed to longevity risk.[1] Blake et al. (2006b) divide the reasons for the lack of a well-functioning market into design issues, pricing issues and institutional issues. Since there is no well-developed market for longevity risk a replicating portfolio does not exist. Pelsser (2011) discusses and compares several methods proposed in the literature to price risks in incomplete markets. In these methods one has to define a pricing operator to determine the value of a payoff. In practice these methods are difficult to implement. Third, buy-outs and buy-ins are ways to insure macro-longevity risk (Lin et al. (2015)). A disadvantage of pension buy-outs and buy-ins is that they are expensive.[2] Fourth, natural hedging is a way to manage macro-longevity risk (Cox and Lin (2007)). Macro-longevity risk in annuity policies can be hedged with mortality risk in life insurance policies.[3] Participants living longer than expected have a negative impact on annuity policies but a positive impact on life insurance products since less participants die at a young age. However, mortality risk only provides a partial hedge to longevity risk due to the different nature of both risks and the different age groups. Moreover, the mortality risk market is more than five times smaller than the longevity risk market (EIOPA (2011)).

The remainder of this paper is organized as follows. Section 2 describes the modeling of macro-longevity risk. Section 3 explains the concept of collective risk sharing. Section 4 describes the different retirement age policies. Section 5 presents the results. Section 6 concludes and gives a policy evaluation.

[1] Basis risk arises from the different mortality experience of the population cohort covered by the mortality index and the cohort relevant to the hedger.

[2] This is a result of insurance companies being typically subject to more stringent regulation than pension funds and because any initial underfunding requires a lump-sum payment by the sponsor to reach full funding before the plan can be sold to a third party (Basel Committee on Banking Supervision (2013).

[3] In this context, mortality risk is the risk that people live shorter than expected.

2 Macro-longevity risk

We consider three sources of macro-longevity risk. These are visualized in Figure 1. The first source is stochastic variation. This is the random variation in the aggregate realized number of deaths. A stochastic mortality model captures stochastic variation.[4] The second source is parameter risk. It is the uncertainty about the true value of the parameters of the stochastic model. The third source is model risk. This is the uncertainty about the appropriateness of the mortality model. For instance, model risk can occur due to structural breaks that are not captured by the model. Medical innovations or a rapid increase of obesity can cause these structural breaks. All three sources of uncertainty can lead to mis-estimation of mortality rates. A stochastic mortality model only takes into account stochastic variation while ignoring the other sources of risk. Ideally, one wants to model macro-longevity risk including all these sources of risk.

In this paper the main source of macro-longevity risk is stochastic variation. However, we also consider a type of parameter risk. This will be discussed in more detail in Section 2.2. In a sensitivity analysis in Section 5.1.3 we address model risk by considering an alternative model for macro-longevity risk.

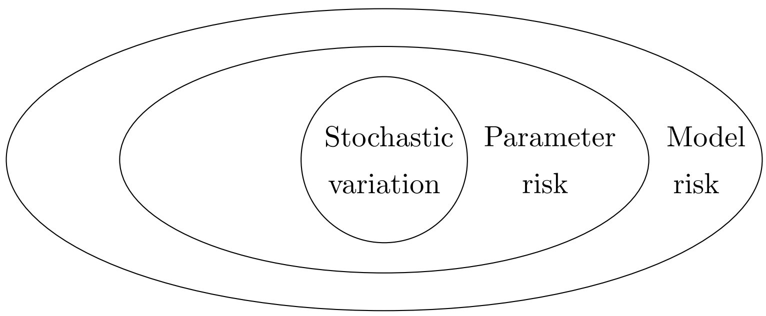

Figure 1: Sources of macro-longevity risk.

We employ the widely used Lee and Carter (1992) model which is a stochastic mortality model that allows for stochastic variation in death rates. It is fitted to historical data to forecast death rates and to quantify macro-longevity risk. Cairns et al. (2011) discuss the suitability of six stochastic mortality models for forecasting mortality and conclude that the Lee-Carter model is both reasonably robust relative to historical data and produces plausible forecasts.[5] Several academics use the Lee-Carter model to model macro-longevity risk, for example

[4] Stochastic variation in death rates of individuals within cohorts, i.e., individual uncertainty about the time of death, is excluded. We assume that cohorts are large enough so that micro-longevity risk is fully diversified.

[5] Alternative stochastic mortality models are for example the model of Renshaw and Haberman (2006) that is an extension of the Lee-Carter model including a cohort effect and the two-factor model of Cairns et al. (2006b).

Hari et al. (2008), Cocco and Gomes (2012), Stevens (2017) and De Waegenaere et al. (2017). Moreover, the model is the basis of several mortality table forecasts in practice.[6]

We discuss the Lee-Carter model in Section 2.1 and elaborate on macro-longevity risk in the Lee-Carter model in Section 2.2. In Section 4 we discuss different retirement age policies.

2.1 Lee-Carter model

The central death rate $\mu_{x,t}$ for a cohort of age x in year t equals

$$\mu_{x,t} = \frac{D_{x,t}}{E_{x,t}}, \tag{1}$$

where $D_{x,t}$ is the number of deaths in year t among the people in the cohort of age x and $E_{x,t}$ is the number of people in the cohort of age x in year t.

The Lee-Carter model estimates the log central death rates with the following expression[7]

$$\ln(\mu_{x,t}) = \alpha_x + \beta_x \kappa_t + \epsilon_{x,t}, \tag{2}$$

where α_x is an age-specific constant, κ_t is a time trend and β_x represents the sensitivity of the log central death rates to the time trend. The time trend reflects the development of death rates over time. This trend is generally downward implying an increasing life expectancy over time. The error term $\epsilon_{x,t}$ is normally distributed with mean zero and age-dependent variance $\sigma^2_{\epsilon,x}$.

The Lee Carter model assumes that the central death rates are constant during a year, i.e., $\mu_{x+s,t+s} = \mu_{x,t}(0 \leq s \leq 1)$. Therefore, we can approximate the one-year death probability $q_{x,t}$ in the following way

$$q_{x,t} \approx 1 - \exp(-\mu_{x,t}). \tag{3}$$

The one-year death probability is the probability that an individual of age x and alive at the beginning of year t dies before year $t+1$. The one-year survival probability $p_{x,t}$ equals

$$p_{x,t} = 1 - q_{x,t} \approx \exp(-\mu_{x,t}). \tag{4}$$

One-year survival probabilities can be used to calculate the probability that an individual of age x in year t is still alive after i years. This is called the cumulative survival probability $cp^t_{x,i}$

$$cp^t_{x,i} = \prod_{j=0}^{i-1} p_{x+j,t+j}. \tag{5}$$

---

[6] For example the U.S. Census Bureau and the U.S. Social Security Administration. The Actuarial Society in the Netherlands ('Koninklijk Actuarieel Genootschap') uses an alternative specification of this model.

[7] The logarithm of $\mu_{x,t}$ ensures that death rates cannot be negative. However, death rates can exceed unity but this is not a problem in practice. This can be avoided by modeling $\ln(\mu_{x,t}/(1-\mu_{x,t}))$, but in that case a linear trend in k does not imply a constant geometric rate of decline for each age-specific death rate (Lee (2000)).

The Lee-Carter model forecasts survival probabilities by estimating the time trend κ_t in (2) with a standard univariate time series model. Lee and Carter (1992) conclude after testing several $ARIMA$ specifications that the $ARIMA(0,1,0)$ model, a random walk with drift, is most appropriate to fit the data. This model equals

$$\kappa_t = c + \kappa_{t-1} + \eta_t, \tag{6}$$

where c is the drift and η_t is the error term that is normally distributed with mean zero and variance σ_η^2. The Lee-Carter model assumes that the error terms $\epsilon_{x,t}$ in (2) and η_t in (6) are independent. This independency implies that for each cohort mortality develops at an own age-specific exponential rate.

Calibration of the Lee-Carter model

In this paper we use mortality data of Dutch females from 1985 until 2014 from the Human Mortality Database to calibrate the parameters of the Lee-Carter model.[8,9] The central death rates $\mu_{x,t}$ are calculated using the number of deaths $D_{x,t}$ and number of people $E_{x,t}$ as in (1).

For very high ages no death rates are available in the database. When excluding the death rates beyond the age of 90 the expected remaining lifetime will be underestimated. We apply the method of Kannisto (1994) to extrapolate the central death rates for ages $x \in \{91, ..., 110\}$ using the death rates of younger cohorts. This method uses a logistic regression based on $\mu_{x,t}$ for ages $x \in \{80, 81, ..., 90\}$. Death rates above age $x = 110$ are assumed to be equal to the death rates at age $x = 110$.

We estimate parameters α_x, β_x and κ_t in (2) using a singular value decomposition. However, this method does not produce uniquely identified parameters. Therefore, we impose restrictions to identify the model. We use the standard identification choice of Lee and Carter (1992) that imposes the following constraints

$$\sum_{x=0}^{110} \beta_x = 1$$
$$\sum_{t=1985}^{2014} \kappa_t = 0.$$

The age-specific constant α_x is the average log central death rate of cohort of age x over time, i.e., $\alpha_x = \frac{1}{30} \sum_{t=1985}^{2014} \ln(\mu_{x,t})$. Subsequently the drift c and variance σ_η^2 in (6) are estimated using

[8] Human Mortality Database (HMD). University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany): `http://www.mortality.org/`.

[9] A calibration period of 30 years is conventional. For statistical reliability, one would prefer a longer calibration period (HMD). However, a shorter calibration period leads to a better estimate of the current trend in mortality improvements.

the κ_t's.

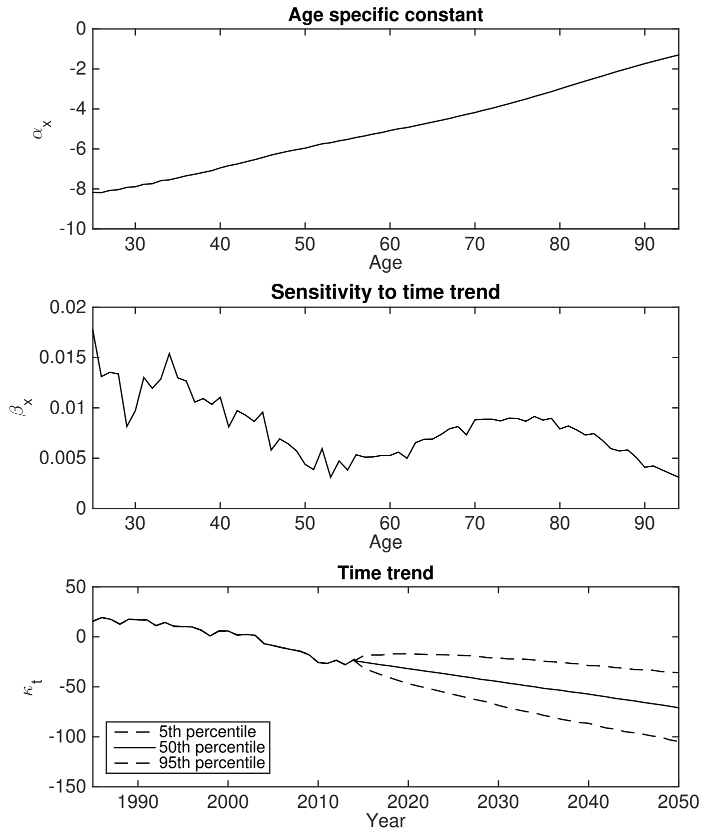

Figure 2: Parameter estimates of the Lee-Carter model using mortality data of Dutch females from 1985 until 2014. The top graph shows the age-specific constant α_x, the middle graph shows the sensitivity of death rates to the time trend β_x and the bottom graph shows the time trend κ_t. The bottom graph also contains the expected future time trend including a 90% confidence interval.

Figure 2 displays the estimates of the three key parameters in the Lee-Carter model in (2) using mortality data of Dutch females from 1985 until 2014. The top graph shows that the age-specific constant increases with age x. This implies higher death rates at higher ages. This is intuitive as older people have a higher change of dying. The middle graph shows that the sensitivity of death rates to the time trend in general decreases with age but in a non-monotonic way. A decreasing sensitivity implies that death rates for high ages are less effected by the time trend compared to death rates for young ages. The bottom graph shows that the time trend k_t decreases over time. This implies that death rates decrease over time. It is result of for example medical innovations and better nutrition. The estimated drift equals $\hat{c} = -1.3$. Each year the time trend κ_t decreases with 1.3 in expectation. The graph also contains the expected future time trend including the 90% confidence interval that is a result of the stochastic variation in the time trend.

2.2 Macro-longevity risk in Lee-Carter model

As already mentioned at the start of Section 2 the main source of macro-longevity risk in this paper is stochastic variation. Macro-longevity risk in the Lee-Carter model arises from two random variables:

- *Uncertainty in time trend*: random shock η_t in the time trend κ_t in (6). It reflects the uncertainty in the time trend, i.e., development of death rates over time. The impact of this shock on future death rates depends on the size of σ_η and β_x.
- *Uncertainty in death rates*: random shock $\epsilon_{x,t}$ in the log central death rate $\mu_{x,t}$ in (2). It reflects particular age-specific historical influences not captured by the model. The impact of this shock on future death rates depends on the size of $\sigma_{\epsilon,x}$.

We model the first source of macro-longevity risk, stochastic variation, as the aggregate effect of those two random variables. We assume that η_t and $\epsilon_{x,t}$ are independent and normally distributed. The sum of two independent normal random variables is again normally distributed

$$\left.\begin{aligned} \eta_t &\sim N(0,\sigma_\eta^2) \\ \epsilon_{x,t} &\sim N(0,\sigma_{\epsilon,x}^2) \end{aligned}\right\} \quad \Rightarrow \quad \beta_x\eta_t + \epsilon_{x,t} \sim N(0, \beta_x^2\sigma_\eta^2 + \sigma_{\epsilon,x}^2). \tag{7}$$

The trend risk η_t is multiplied with the sensitivity of to the time trend β_x because the sensitivity parameter β_x determines the impact of the time trend on death rates. Macro-longevity risk has zero mean because it is the risk that future mortality rates deviate from the best estimate mortality rates.

In this research we do not consider yearly macro-longevity shocks but consider macro-longevity risk on a 10-year horizon because a pension contract has a long horizon and we want to focus on structural changes in life expectancy only. We determine macro-longevity shocks on a 10-year horizon by summing up the independent normal random variables in (7) over 10 years

$$\sum_{i=0}^{9}(\beta_{x+i}\eta_{t+i} + \epsilon_{x+i,t+i}) \sim N\left(0, \sigma_\eta^2 \sum_{i=0}^{9}\beta_{x+i}^2 + \sum_{i=0}^{9}\sigma_{\epsilon,x+i}^2\right). \tag{8}$$

The second source of risk is parameter risk. We calibrate the parameters in the mortality model using mortality data. When more recent mortality data are available we can recalibrate the parameters. Recalibration changes the parameter estimates (Cairns (2013)). In this paper we include recalibration risk. We use the realized death rates $\mu_{x,t}$ including the trend shocks η_t and estimation shocks $\epsilon_{x,t}$ to recalibrate the parameters in (2) and (6). Subsequently, we use these recalibrated parameters to forecast future death rates. By considering recalibration risk we include the influence of parameter risk.[10]

The third source of macro-longevity risk is model risk. We initially exclude model risk in

[10] A more formal way to include parameter risk is to use standard Bayesian methods (Cairns et al. (2006b)).

our analysis and address this separately in Section 5.1.3.

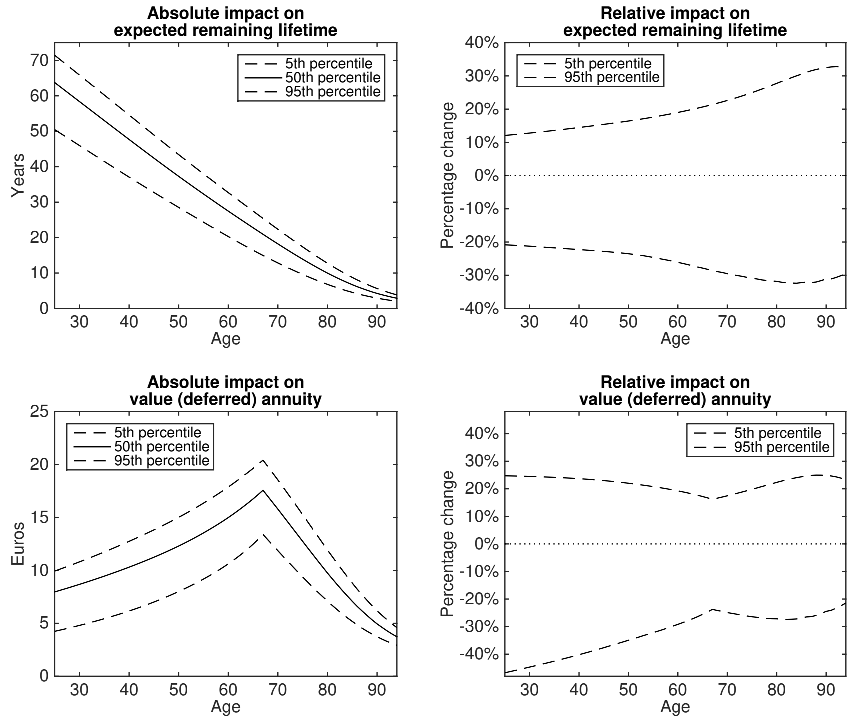

Figure 3: Impact of macro-longevity risk measured on a 10-year horizon in the Lee-Carter model on the expected remaining lifetime and the value of a (deferred) variable annuity for a Dutch female in 2014 in absolute terms (lefthand graphs) and relative change (righthand graphs) assuming a constant interest rate $r = 2\%$ and fixed retirement age $R = 67$.

Figure 3 visualizes the impact of macro-longevity risk measured on a 10-year horizon in the Lee-Carter model on the expected remaining lifetime (top graphs) and the value of a (deferred) variable annuity (bottom graphs) by displaying different percentiles of the distribution.[11] Besides the absolute impact on the expected remaining lifetime and the value of a (deferred) annuity (lefthand graphs), it is also interesting to look at the relative change of these variables (righthand graphs). We assume that the interest rate - used to determine the value of a (deferred) annuity - equals $r = 2\%$ and the retirement age equals $R = 67$. One can also make the retirement age contingent on life expectancy which is the case in several countries. This will be discussed in the next section.

The top lefthand graph shows that the expected remaining lifetime decreases with age. E.g., at age 25 it is 64 years and 11 years at age 80. This decrease is intuitive as older people have a higher chance of dying. Moreover, we notice that the impact of macro-longevity risk also

[11] Negative (positive) macro-longevity shocks, i.e., negative (positive) random shocks in log central death rates, have a positive (negative) impact on life expectancy and annuity values. To avoid confusion we denote negative (positive) macro-longevity shocks by unexpected increases (decreases) in life expectancy.

decreases with age. E.g., the difference between the 5th and 95th percentile at age 25 is 21 years and 6 years at age 80. There are two reasons for this decreasing impact. *First*, a longevity shock has an impact on all future death probabilities. The expected remaining lifetime of young cohorts depends on more future death probabilities compared to the expected remaining lifetime of old cohorts. *Second*, the impact of both trend and estimation risk decreases with age. The sensitivity of the death rates β_x decreases with age implying a decreasing impact of the trend risk. The variance of the estimation risk $\sigma^2_{\epsilon,x}$ generally decreases with age as there is less uncertainty at higher death rates. This implies a decreasing impact of estimation risk.

The value of a deferred annuity (bottom lefthand graph) increases before retirement because of two reasons:

- The probability that a participant reaches the retirement age increases with age.
- The value of a deferred annuity is lower for young cohorts compared to cohorts just before retirement because of a larger discounting effect.

The relative change of the value of a (deferred) annuity as a result of a macro-longevity shock is in the same order of magnitude for all age cohorts. Later in this paper we will see that this explains the small welfare gains in case of collective risk sharing when the retirement age is fixed.

Another important observation in Figure 3 is that the impact of macro-longevity risk on the expected remaining lifetime and (deferred) annuity value is asymmetric. Unexpected increases in life expectancy have a smaller impact than unexpected decreases in life expectancy. This can be explained by the exponential distribution of death rates. A consequence of this asymmetry is that the expectations of future survival probabilities and therefore also the expected remaining lifetime and expected (deferred) annuity value are smaller than its forecasted values. We present a derivation in Appendix A.1.

3 Sharing macro-longevity risk

The previous section discusses the modelling of macro-longevity risk. In this section we consider the concept of collective risk sharing. Pension providers can create an internal market for macro-longevity risk. We refer to this as collective risk sharing. Collective risk sharing of macro-longevity risk can be welfare enhancing because the risk is not traded on a liquid market and cohorts are affected differently by the risk.[12] In fact, it creates a new asset that can be priced and makes the market more complete.

[12] Collective risk sharing can also be welfare enhancing if the risk is traded with future cohorts. In this paper we abstract from this dimension.

We discuss the concept of collective risk sharing in Section 3.1. We use a stylized two-agent model in Section 3.2 to derive an analytical risk sharing solution. This model gives economic intuition. Subsequently, we present a full model in Section 3.3 that consists of many cohorts representing the population of a pension fund. We use this model to share macro-longevity risk between cohorts in a pension scheme.

3.1 Risk sharing model

Suppose there are N agents with initial wealth W_i of agent i and agents are exposed to an exogenous risk factor $\tilde{y}_i$. Each individual's preferences is represented by a utility function $U_i(.)$ with positive and decreasing marginal utility. Each agent maximizes the expected utility of his consumption. In autarky consumption consists of wealth and the exposure to the risk factor

$$C_i^a = W_i + \tilde{y}_i \quad i = 1, ..., N. \tag{9}$$

In case of risk sharing the agents aggregate and subsequently redistribute the risk among themselves through the continuous risk sharing function $T_i(\tilde{y}) = T_i(\tilde{y}_1, ..., \tilde{y}_N)$. This leads to the following expression for consumption in case of risk sharing

$$C_i^s = W_i + T_i(\tilde{y}) \quad i = 1, ..., N, \tag{10}$$

under the condition that the aggregate risk is fully distributed over all agents

$$\sum_{i=1}^{N} T_i(\tilde{y}) = \sum_{i=1}^{N} \tilde{y}_i. \tag{11}$$

Risk sharing is Pareto improving compared to autarky if the welfare of at least one agent improves

$$\mathbb{E}[U(C_i^s)] > \mathbb{E}[U(C_i^a)] \quad \text{for some } i = 1, ..., N \tag{12}$$

and all other agents do not become worse off

$$\mathbb{E}[U(C_i^s)] \geq \mathbb{E}[U(C_i^a)] \quad \forall i = 1, ..., N. \tag{13}$$

A risk sharing rule $\{T_1, ..., T_N\}$ is Pareto optimal if no Pareto improvement is possible, i.e. there does not exist an alternative risk-sharing rule $\{\overline{T}_1, ..., \overline{T}_N\}$ such that

$$\mathbb{E}[U(W_i + \overline{T}_i(\tilde{y}))] \geq \mathbb{E}[U(W_i + T_i(\tilde{y}))] \quad \forall i = 1, ..., N, \tag{14}$$

and for at least one agent strict inequality holds

$$\mathbb{E}[U(W_i + \overline{T}_i(\tilde{y}))] > \mathbb{E}[U(W_i + T_i(\tilde{y}))] \quad \text{for some } i = 1, ..., N. \tag{15}$$

A Pareto optimal risk-sharing rule yields the highest welfare gain compared to autarky.

The theorem of Borch (1960) provides the following necessary and sufficient conditions for a risk sharing rule $\{T_1, ..., T_N\}$ to be Pareto optimal

$$\begin{aligned} &U_i'(W_i + T_i(\tilde{y})) = c_i U_1'(W_1 + T_1(\tilde{y})) \quad \forall i = 1, ..., N, \\ &\sum_{i=1}^{N} T_i(\tilde{y}) = \sum_{i=1}^{N} \tilde{y}_i, \end{aligned} \tag{16}$$

where $c_2, c_3, ..., c_N > 0$ can be chosen arbitrarily and $c_1 = 1$. This theorem shows that in a Pareto-optimal risk-sharing rule the ratio of marginal utilities of two different agents is equal to a constant. Borch (1960) also proofs that in a Pareto-optimal risk-sharing rule $T_i(\tilde{y})$ is a function of the aggregate risk $\sum_{i=1}^{N} \tilde{y}_i$ only. This implies that in a Pareto-optimal risk-sharing rule a pool must be formed of the aggregate risk of all participants.

The conditions for the existence of a Pareto-optimal risk-sharing rule in Borch's theorem in (16) are very weak. DuMouchel (1968) shows that if the utility functions are strictly monotonic these conditions are satisfied and thus a Pareto-optimal risk-sharing rule exists.

So far we have not further specified the risk-sharing functions $\{T_1, ..., T_N\}$. A subset of all possible risk-sharing functions $\{T_1, ..., T_N\}$ is the collection of linear risk sharing rules[13]

$$T_i(\tilde{y}) = t_{0,i} + \eta_i \sum_{i=1}^{N} \tilde{y}_i. \tag{17}$$

In this risk sharing rule the risk transfer η_i is the fraction of the aggregate risk that agent i absorbs and $t_{0,i}$ is a constant risk compensation that agent i receives ex-ante. Some agents receive a positive risk compensation which has to be financed by the other agents. Linear risk sharing rules are easy to implement and the risk compensation $t_{0,i}$ can be interpreted as a risk premium for absorbing risk. A Pareto-optimal risk-sharing rule is generally linear. Huang and Litzenberger (1985) show that a Pareto optimal risk-sharing rule is linear if the agents have the same cautiousness.[14] This condition is satisfied when the individual utility functions are member of the Hyperbolic Absolute Risk Aversion (HARA) class (Aase (2002)). This is a general class of utility functions that are often used in practice.

In this research we focus on HARA utility functions. Pareto-optimal risk-sharing rules are in this case linear. In the stylized two-agent model in Section 3.2 we assume exponential utility which belongs to the HARA class and exhibits constant absolute risk aversion (CARA). We use this utility function because of its analytical convenience and will show that the Pareto-optimal risk-sharing rule is indeed linear in aggregate risk. In the full model we assume power utility which also belongs to the HARA class and exhibits constant relative risk aversion (CRRA). We assume power utility in the full model because power utility is more

[13] Strictly speaking, these are affine risk sharing rules.

[14] The cautiousness is the derivative of the reciprocal of absolute risk aversion. It measures how quickly the coefficient of risk aversion increases as wealth goes down, see, e.g. Wilson (1968).

in line with empirical evidence. Exponential utility implies increasing relative risk aversion which contradicts empirical studies. Moreover, power utility has become the workhorse of modern macroeconomics.

The conditions in (16) do not imply a unique Pareto-optimal risk-sharing rule since the positive constants $c_2, c_3, ..., c_N$ can be chosen arbitrarily. The welfare gain from risk sharing can be distributed over the agents in different ways. However, there must be an upper limit to c_i since utility decreases with increasing c_i and agents cannot become worse off in the Pareto-optimal risk-sharing rule. One can find a unique solution within the set of Pareto-optimal risk-sharing rules by looking for an equilibrium. In this approach the agents can trade in a fictitious market. This method is used by, e.g., Krueger and Kubler (2006), Ball and Mankiw (2007) and Gottardi and Kubler (2011). An alternative way to find a unique Pareto-optimal risk-sharing rule is by making use of a social planner and using a utility-based fairness criterion. The social planner maximizes aggregate welfare and reallocates risk across agents. A social planner is used by, e.g., Gordon and Varian (1988), Gollier (2008), Cui et al. (2011) and Bovenberg and Mehlkopf (2014). The utility-based fairness criterion requires that all agents experience the same welfare gain from risk sharing. This criterion is used by, e.g., Gollier (2008) and Bovenberg and Mehlkopf (2014). We use this criterion in the full model in Section 3.3.

3.2 Stylized two-agent model

To understand how collective risk sharing works and leads to welfare gains we first consider a stylized model in which we can derive the Pareto-optimal risk-sharing rule analytically. The model consists of $N = 2$ agents which both have exponential utility with risk aversion α

$$U(C_i) = -\frac{1}{\alpha}\exp(-\alpha C_i). \tag{18}$$

Exponential utility belongs to the HARA class so the Pareto-optimal risk-sharing rule is linear in aggregate risk. For the sake of simplicity we assume both agents are exposed to the same risk factor $\tilde{y}$ but with a different exposure, i.e., $\tilde{y}_1 = \beta_1\tilde{y}$ and $\tilde{y}_2 = \beta_2\tilde{y}$. Risk $\tilde{y}$ represents the unexpected component of a macro-longevity shock that affects both agents but in a different way. We assume $\tilde{y}$ is normally distributed with zero mean and variance σ^2.

We derive the Pareto-optimal risk-sharing rule $\{T_1(\tilde{y}), T_2(\tilde{y})\}$ by maximizing the expected utility of agent 1 under the condition that agent 2 does not become worse off relative to autarky

$$\max_{\eta,t_0} \mathbb{E}[U(W_1 + T_1(\tilde{y}))] \quad \text{such that} \quad \mathbb{E}[U(W_2 + T_2(\tilde{y}))] \geq \mathbb{E}[U(W_2 + \tilde{y}_2))] \tag{19}$$
$$\text{and} \quad T_1(\tilde{y}) + T_2(\tilde{y}) = \tilde{y}_1 + \tilde{y}_2,$$

where we plug in consumption in autarky (9) and consumption in case of risk sharing (10). The maximization can also be written as follows

$$\max_{\eta,t_0} \mathbb{E}[U(W_1 + T_1(\tilde{y}))] \quad \text{such that} \quad \mathbb{E}[U(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y}))] \geq \mathbb{E}[U(W_2 + \beta_2\tilde{y}))]. \tag{20}$$

This maximization can be solved and yields the following Pareto-optimal risk-sharing rule $\{T_1(\tilde{y}), T_2(\tilde{y})\}$

$$\begin{cases} T_1(\tilde{y}) & = t_{0,1} + \eta_1(\beta_1 + \beta_2)\tilde{y} = -\frac{1}{8}\alpha\sigma^2(\beta_1 + 3\beta_2)(\beta_1 - \beta_2) + \frac{1}{2}\Big(\beta_1 + \beta_2\Big)\tilde{y}, \\ T_2(\tilde{y}) & = t_{0,2} + \eta_2(\beta_1 + \beta_2)\tilde{y} = \frac{1}{8}\alpha\sigma^2(\beta_1 + 3\beta_2)(\beta_1 - \beta_2) + \frac{1}{2}\Big(\beta_1 + \beta_2\Big)\tilde{y}. \end{cases} \tag{21}$$

A proof is provided in Appendix A.2. We can conclude that the Pareto-optimal risk-sharing rule is indeed linear in aggregate risk. This makes sense since exponential utility belongs to the HARA class. As mentioned in Section 3.3 a Pareto-optimal risk-sharing rule is linear for utility functions of the HARA class.

The optimal risk-sharing rule in (21) shows that in case of exponential utility the optimal risk transfer η_i is independent of the wealth of the agents and is the same for both agents. They both absorb half of the aggregate shock. Taking the expectation of the Pareto-optimal risk-sharing rule leads to a constant risk compensation $t_{0,i}$ since risk factor $\tilde{y}$ has zero mean

$$\begin{cases} \mathbb{E}[T_1(\tilde{y})] & = t_{0,1} = -\frac{1}{8}\alpha\sigma^2(\beta_1 + 3\beta_2)(\beta_1 - \beta_2), \\ \mathbb{E}[T_2(\tilde{y})] & = t_{0,2} = \frac{1}{8}\alpha\sigma^2(\beta_1 + 3\beta_2)(\beta_1 - \beta_2). \end{cases} \tag{22}$$

The constant risk compensation $t_{0,i}$ depends on the risk aversion α, the exposure of both agents to the risk factor β_i and the variance σ^2 of the risk factor $\tilde{y}$. The risk compensation $t_{0,1}$ is negative if $\beta_1 > \beta_2$. This implies that agent 1 has to pay a risk compensation to agent 2 if the exposure of agent 1 to the risk in autarky is larger compared to the exposure of agent 2. Because both agents absorb half of the aggregate shock, agent 2 wants to receive a positive risk compensation in return.

The risk compensation $t_{0,i}$ determines how the welfare gain from risk sharing is distributed among the agents. Because the inequality restriction is binding the welfare gain of risk sharing goes completely to agent 1. The optimal risk-sharing rule in (21) is not necessarily the only Pareto-optimal risk-sharing rule. In fact, there is generally a whole set of Pareto-optimal risk-sharing rules. There are also Pareto-optimal risk-sharing rules in which both agents gain from risk sharing. In that case $t_{0,2}$ is higher and thus $t_{0,1}$ must be lower as long as agent 1 does not become worse off.

We make the simplifying assumption that both agents are exposed to the same risk factor $\tilde{y}$. In case both agents are exposed to a different risk factor the optimal risk sharing rule is still linear in aggregate risk. However, the equation of the risk compensation becomes more complex.

3.3 Collective risk sharing of macro-longevity risk: full model

We extend the stylized two-agent model of Section 3.2 to a full model with many cohorts representing the population of a pension fund. In our full model, macro-longevity risk impacts survival probabilities and therefore also retirement consumption in a non-linear way based on the Lee-Carter model.

This full model consists of $N = 70$ cohorts. Cohort 1 is aged 25 and cohort 70 is aged 94.[15] We base the number of participants n_i in cohort i on the cumulative probability that a participant is still alive at age $i + 24$. So old cohorts consist of less participants compared to young cohorts. The lefthand graph in Figure 4 visualizes this population composition. Participants have identical preferences given by a power utility function with risk aversion $\gamma = 5$[16]

$$U(C_i) = \begin{cases} \dfrac{C_i^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1, \\ \ln(C_i) & \text{if } \gamma = 1. \end{cases} \tag{23}$$

Power utility belongs to the HARA class so the Pareto-optimal risk-sharing rule is linear in aggregate risk. The total wealth W_i of a participant in cohort i depends on his or her age. The righthand graph in Figure 4 visualizes the development of wealth over the life-cycle of a participant. Wealth increases during the working period as the participant contributes to the pension fund. Wealth at the start of the working period is positive because wealth consists of financial wealth and human wealth (i.e., future pension contributions).[17,18] Furthermore, we initially assume that the participants retire at age $R = 67$, the interest rate - used to determine the value a (deferred) annuity - equals $r = 2\%$.

We consider a DC pension scheme. Consumption after retirement depends on the value of an annuity. We assume the participant buys a variable annuity which value varies with future survival probabilities.[19] If for example life expectancy increases the annuity value increases. This has a negative effect on consumption after retirement and implies that macro-longevity risk is borne by the participant. The value of a (deferred) variable annuity a_x^t, that pays 1

[15] We exclude cohorts older than age 94 because the number of participants in these cohorts is very small and therefore do not influence the results significantly.

[16] Power utility has become the workhorse of macro-economics and finance and is in line with empirical studies compared to exponential utility. We justify the assumption that agents have the same risk aversion γ because collective risk sharing within a pension fund often occurs within a group of participants with similar characteristics such as education, salary, etc.

[17] Human wealth is equal to the present value of future pension contributions and not the present value of future labor income because the pension contributions are fixed in our model.

[18] We assume that labor income is fixed and the same for each cohort and each age. As a result, the impact of macro-longevity risk on replacement rates is similar to the impact of macro-longevity risk on consumption.

[19] Variable annuities can also vary with realized investment returns. Because we exclude investment risk, this is not the case in this paper.

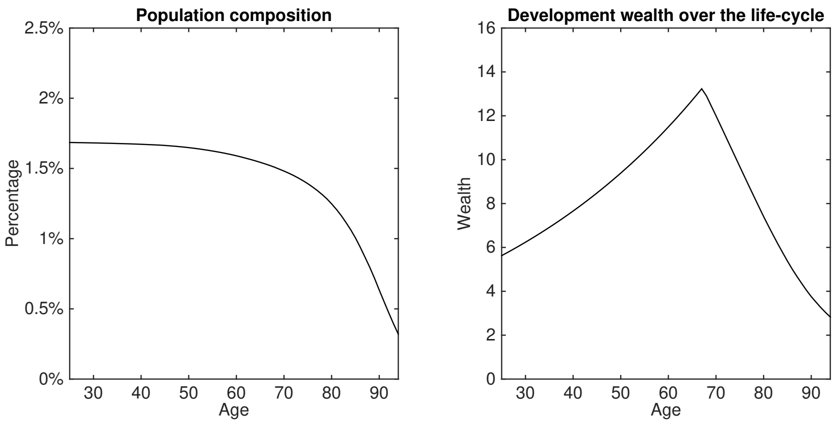

Figure 4: Population composition (lefthand graph) and development of wealth over the life-cycle of a participant (righthand graph).

dollar annually during retirement, for an individual of age x in year t is calculated as follows

$$a_x^t = \sum_{j=\max(x,R)}^{M} \frac{1}{(1+r)^{j-x}} cp_{x,j-x}^t. \tag{24}$$

In this formula R equals the retirement age, M is the maximum age an individual can reach and $cp_{x,i}^t$ is the probability of still being alive after i years as in (5). We assume a constant interest rate r.

For ease of reference we denote the value of a (deferred) annuity in (24) by $a_x^t = a_i$. A macro-longevity shock impacts future survival probabilities which influence the value of a (deferred) annuity as stated in (24). The value of a (deferred) annuity changes for cohort i from a_i to $\tilde{a}_i$ due to a macro-longevity shock in the Lee-Carter model. The expected annual consumption after retirement in autarky C_i^a for cohort i after a shock is given by

$$C_i^a = \frac{W_i}{\tilde{a}_i}. \tag{25}$$

To determine the impact of macro-longevity risk on consumption we calculate for each cohort how much money is needed (or is left) to fully compensate the impact of a macro-longevity shock.[20] We denote this by $\tilde{y}_i$

$$\frac{W_i}{a_i} = \frac{W_i + \tilde{y}_i}{\tilde{a}_i}$$

$$\tilde{y}_i = W_i\left(\frac{\tilde{a}_i}{a_i} - 1\right). \tag{26}$$

$\tilde{y}_i$ represents the amount of money to offset the effect of a macro-longevity shock on consumption in autarky. If the annuity value increases (decreases) due to an unexpected increase (decrease)

[20] In this paper we assume that consumption before retirement is fixed, i.e., a macro-longevity shock can only be absorbed by changing consumption after retirement. In case a participant can also change consumption before retirement, the impact of a macro-longevity shock on the consumption level after retirement will be smaller for workers.

in life expectancy, $\tilde{y}_i$ is positive (negative) and money is needed (left). $\tilde{y}_i$ is not the same for each cohort i because the impact of a macro-longevity shock on future death rates depends on age. We can calculate the total money needed (or left) to fully compensate the impact of a macro-longevity shock for all N cohorts. We denote this by $\widetilde{y}_T$

$$\widetilde{y}_T = \sum_{i=1}^{N} n_i \tilde{y}_i, \tag{27}$$

where n_i is the number of participants in cohort i. Macro-longevity risk is shared by distributing the aggregate macro-longevity shock $\tilde{y}_T$ among cohorts. Because power utility belongs to the HARA class the Pareto-optimal risk-sharing rule is linear. Each participant absorbs part of the aggregate macro-longevity shock η_i and receives (or pays) a risk compensation $t_{0,i}$ as in (17). Consumption in case of risk sharing thus equals

$$C_i^s = \frac{W_i + \tilde{y}_i - \eta_i \widetilde{y}_T - t_{0,i}}{\tilde{a}_i}. \tag{28}$$

We determine the Pareto-optimal risk-sharing rule by making use of a social planner who maximizes aggregate welfare and reallocates risk across agents

$$\max_{\substack{\eta_1,\eta_2,\dots\eta_N \\ t_{0,1},t_{0,2},\dots t_{0,N}}} \sum_{i=1}^{N} n_i \delta_i \mathbb{E}[U(C_i^s)], \tag{29}$$

where the following restrictions should be satisfied

$$\begin{cases} \sum\limits_{i=1}^{N} n_i \eta_i = 1, \\ \sum\limits_{i=1}^{N} n_i t_{0,i} = 0. \end{cases} \tag{30}$$

The first restriction makes sure that the aggregate risk is fully distributed over all agents. The second restriction guarantees that the total risk compensation that participants receive is paid by the other participants. The parameter δ_i in (29) represents the discount factor that the social planner uses to weigh the relative importance of the cohorts. As a result the maximization in (29) does not yield a unique Pareto-optimal risk-sharing rule. There is no unique Pareto-optimal risk-sharing rule since the welfare gain from risk sharing can be distributed over the participants in different ways. In the maximization in (20) a unique Pareto-optimal risk-sharing rule is found that allocates the complete welfare gain of risk sharing to agent 1. Such an allocation is not the most desirable risk-sharing agreement in a pension scheme. We use a utility-based fairness criterion to find a unique Pareto-optimal risk-sharing rule. We determine δ_i such that the welfare gain from risk sharing relative to autarky is the same for all participants in the pension scheme.

4 Retirement age policies

The significant increase in life expectancy during the last decades had a major impact on the sustainability of pension systems. As a response several countries are linking the state pension age to life expectancy developments. In the United Kingdom for example the government plans to link the state pension age at future dates to the projected longevity of the population in such a way that people receive state pension during a fixed proportion of adult life (Hammond et al. (2016)). Under this policy both the working and retirement period increase if life expectancy increases. In the Netherlands the retirement age is linked to life expectancy in a different way. The Dutch government implemented a law that links the retirement age to the remaining life expectancy of the population at age 65. Under this policy the absolute length of the retirement period is fixed and independent of life expectancy while the working period increases if life expectancy increases.

In this paper we focus on occupational pension schemes. The retirement age in occupational pension schemes is often equal to the state pension age. As a consequence, the retirement age policy of the government also impacts the retirement age in occupational pension schemes and thus the ability to share macro-longevity risk in occupational pension schemes. We consider three policies:

1. **Fixed retirement age** (FRA): the retirement age is fixed, i.e., the retirement age does not change after macro-longevity shocks. In this policy the length of the working period is constant. This policy supports the belief that if people live longer, they extent their retirement period. In most countries, for example in the United States and Australia, the retirement age is not linked to life expectancy.
2. **Partial adjustment of the retirement age** (PARA): the retirement age automatically adjusts to life expectancy developments in a such a way that retirement consumption remains the same.[21] This means, e.g., that if life expectancy increases (decreases) with 12 months, the retirement age should increase (decrease) with roughly 9 months.[22] In this policy consumption after retirement is constant. The adjustment only holds for working participants, since retirees cannot adjust their retirement age anymore. This policy is close to the retirement age policy in the United Kingdom.[23]
3. **Full adjustment of the retirement age** (FARA): the retirement age automatically keeps up fully with life expectancy changes. This means, e.g., that if the remaining life expectancy at retirement increases (decreases) with 12 months, the retirement age also increases (decreases) with 12 months. In this policy the length of the retirement period is constant. The adjustment holds for working participants only, since retirees cannot

[21] There are also countries in which the retirement age is not automatically linked to life expectancy but the government decides to increase the retirement age based on life expectancy improvements incidentally. We do not investigate such a policy.

[22] The exact increase (decrease) does not only depend on the size of the longevity shock but also on the impact of the longevity shock on survival probabilities at different ages and the life expectancy before the longevity shock.

[23] The retirement age adjustment in the UK proposal depends on the proportion of adult life that people receive state pension.

adjust their retirement age anymore. This policy supports the belief that if people live longer, they increase their labor supply by extending their working period. This policy is similar to the retirement age policy in the Netherlands.[24]

Stevens (2017) investigates the effect of different retirement age policies on the distribution of the (forecasted) retirement age. He concludes that if the retirement age is linked to life expectancy macro-longevity risk is effectively hedged. However, such a policy also leads to substantial uncertainty in the retirement age and length of the retirement period.

	Working period	Retirement period	Retirement consumption	Value annuity	Wealth at retirement
FRA	constant	++	-	++	constant
PARA	+	+	constant	+	+
FARA	++	constant	+	-	++

Table 1: Impact of an unexpected increase in life expectancy on several variables for working participants in case of different retirement age policies.

Table 1 presents the impact of an unexpected increase in life expectancy on several variables for the three retirement age policies. Consumption after retirement is determined by the value of a (deferred) annuity and accumulated wealth at retirement (see (25)). The righthand graph in Figure 4 shows the development of wealth over the life-cycle in case of a fixed retirement age. If the retirement age is linked to life expectancy the development of wealth over the life-cycle is different because the participant accrues more (less) wealth by paying pension premia for a longer (shorter) period.[25] The table presents the impact for working participants only because retirees cannot adjust their retirement age as response to longevity shocks. In case of an unexpected decrease in life expectancy, the signs in Table 1 revert.

In case of a fixed retirement age the length of the working period is constant. As a result the (expected) length of the retirement period increases in case of an unexpected increase in life expectancy. The annuity value increases as a result of higher survival probabilities. Wealth at retirement remains the same. As a result retirement consumption will decrease.

In case of a partial adjustment of the retirement age both the working and retirement period are extended. The annuity value increases as a result of higher survival probabilities. The wealth at retirement also increases because the participant will work longer. The annuity

[24] The Dutch law states that the retirement age R is only adjusted in case the remaining life expectancy at age 65 increases but it remains the same if it decreases. In this paper we assume a symmetric rule, i.e., the retirement age is adjusted in case of both positive and negative shocks.

[25] We assume that the labor market functions perfectly so participants do not experience any difficulties with staying employed.

value and wealth at retirement increase such that consumption after retirement remains the same.

If the retirement age is fully adjusted the length of the retirement period is constant. The (expected) length of the working period increases in case of an unexpected increase in life expectancy. The annuity value is lower than before the longevity shock. Higher survival probabilities have a positive impact on the annuity value, but later retirement has a negative impact on the annuity value. It turns out that the latter effect outweighs. The wealth at retirement increases because the participant will work longer. As a result, retirement consumption will increase.

We use exogenous rules in the retirement age policies. An alternative is an endogenous retirement age. This is considered by, e.g., Chang (1991), Cocco and Gomes (2012) and Heijdra and Romp (2009). The participant optimizes his retirement age based on realized life expectancy improvements. In that case it is necessary to include leisure time besides consumption in the utility function to take the labor-leisure trade-off into account. Otherwise a high retirement age would always be optimal because a shorter retirement period implies a higher consumption after retirement. Cocco and Gomes (2012) investigate the impact of macro-longevity risk on the optimal saving and retirement decision in an individual life-cycle model. They conclude that individuals decide to retire later even if this entails a utility cost in terms of foregone utility of (additional) leisure. Although we do not explicitly model the labor-leisure trade-off in this paper, the retirement age policies represent different preferences regarding consumption and leisure. In case of a fixed retirement age, a life expectancy increase implies a lengthening of the retirement period (leisure) at the expense of the consumption level. In case of a partial adjusted of the retirement age both consumption and leisure (relative to labor) remain approximately equal. A full adjustment of the retirement age implies a higher consumption level at the expense of leisure.

5 Results

In this research we quantify the welfare gains from collective risk sharing in terms of aggregate certainty equivalent consumption after retirement. We use the Lee-Carter model to model macro-longevity risk using mortality data of Dutch females.[26] Table 2 presents the aggregate welfare gains for the three retirement age policies discussed in Section 4.

We observe that for each retirement age policy collective risk sharing of macro-longevity risk

[26] We focus on risk sharing between different cohorts of the same population. We do not investigate risk sharing between the sexes or between different populations. Sharing macro-longevity risk between sexes or different populations is potentially welfare improving but is out of the scope of this paper.

Fixed retirement age (FRA)	0.3%
Partial adjustment retirement age (PARA)	0.5%
Full adjustment retirement age (FARA)	2.7%

Table 2: Welfare gains in terms of aggregate certainty equivalent consumption after retirement from sharing macro-longevity risk measured on a 10-year horizon.

is welfare improving compared to autarky. The design of the retirement age policy impacts the size of welfare gains from sharing macro-longevity risk. In case of a fixed retirement age, the welfare gain equals 0.3 percent. This relatively small welfare gain is a result of the fact that in this policy the impact of macro-longevity risk on retirement consumption for different cohorts is more or less equal (Figure 3). As a result, the welfare gain from risk sharing is limited. When the retirement age is partially adjusted the welfare gain from risk sharing is higher. This is a result of the fact that the expected retirement consumption of workers is not affected by macro-longevity shocks. In case of a full adjustment of the retirement age the aggregate welfare gain increases significantly. This is a result of the large risk bearing capacity of workers. They adjust their labor supply as a hedge against macro-longevity shocks. This increases the risk appetite of the workers to provide insurance to retirees.

In this research we measure welfare gains of sharing macro-longevity risk and not welfare gains of different retirement age policies since the retirement age policy is given for both autarky and risk sharing. We do not focus on the suitability of retirement age policies. This is a different research question and requires the inclusion of leisure time besides consumption in the utility function to take the labor-leisure trade-off into account.

Figure 5 (lefthand graph) visualizes the optimal risk transfer relative to autarky for a participant in cohort i as a percentage of total risk. A positive risk transfer for cohort i means that participants in cohort i absorb risk of other cohorts. A negative risk transfer means that the own exposure to macro-longevity risk is (partly) transferred to other cohorts.[27] In case of a fixed retirement age the risk transfer increases with age for the workers until retirement and decreases with age for retirees. Macro-longevity risk of the young workers and old retirees is (partly) absorbed by the other cohorts. The development of wealth over the life-cycle (righthand graph in Figure 4) primarily explains this shape. Cohorts who have relatively more wealth can absorb more risk. The risk transfer rule in case of a fixed retirement age significantly differs from the risk transfer rule in case the retirement age is adjusted to macro-longevity shocks. The risk transfer rule in case the retirement age is partially adjusted is very similar to the risk transfer rule in case the retirement age is fully adjusted. The workers

[27] The sum of risk transfers in the graph is not exactly equal to zero because each cohort does not consist of an equal number of participants.

absorb risk and the retirees transfer risk. This makes sense because the workers adjust their labor supply to macro-longevity shocks. As a result, they are able to absorb risk of the retirees.

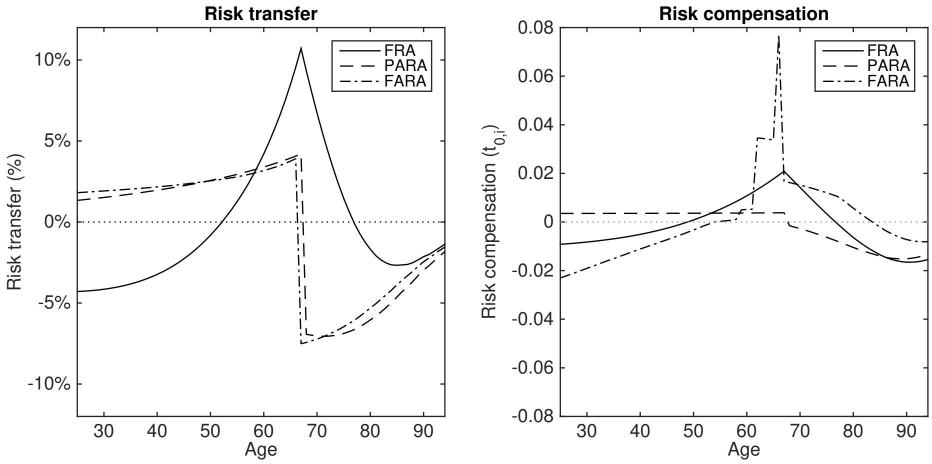

Figure 5: Optimal risk transfer relative to autarky for all cohorts as percentage of total risk (lefthand graph) and corresponding risk compensation (righthand graph) in case of sharing macro-longevity risk measured on a 10-year horizon. A positive (negative) risk transfer for cohort i means that participants in cohort i absorb risk of (transfer risk to) other cohorts. A positive (negative) risk compensation for cohort i means that participants receive (pay) a risk compensation.

The righthand graph in Figure 5 displays the risk compensation $t_{0,i}$ corresponding to the optimal risk transfer for a participant in cohort i under the utility-based fairness criterion (lefthand graph).[28] A positive risk compensation for cohort i means that participants receive a risk compensation. A negative risk compensation for cohort i means that participants pay a risk compensation. In general, cohorts who absorb risk from other cohorts receive a risk compensation and cohorts who transfer risk have to pay a risk compensation. However, this does not hold if the retirement age is fully adjusted. Young cohorts absorb risk from other cohorts but do not receive a risk premium; the risk premium is even negative. Under this policy workers adjust their labor supply as a hedge against macro-longevity shocks. This implies a reverse effect of macro-longevity shocks for workers and retirees (Table 1). As a result, a positive risk compensation is not required for young cohorts to absorb risk of retirees. A final observation is the peak in the risk compensation around age 66 in case of a fully adjusted retirement age. This peak is due to the fact that cohorts just before retirement cannot fully adjust their retirement age in case of an unexpected decrease in life expectancy, i.e., the retirement age cannot be lower than their current age. As a result, the certainty equivalent consumption of these cohorts is relatively high in autarky so risk sharing

[28] The sum of risk compensations in the graph is not exactly equal to zero because each cohort does not consist of an equal number of participants.

is less welfare improving for these cohorts. Therefore, these cohorts require a higher risk compensation.

We consider macro-longevity risk on a 10-year horizon. The welfare gains from sharing macro-longevity risk over the whole life-cycle are most likely higher. Another sidenote is that this paper applies a first-best risk-sharing rule as its benchmark for evaluating welfare effects. In practice, however, the first-best risk-sharing rule may not always be feasible. Policymakers might want to limit the maximum risk a participant can absorb to prevent very large wealth transfers in case of extreme macro-longevity shocks.

5.1 Sensitivity analyses

In this section we verify whether the welfare gains and risk-sharing rules are sensitive to mortality data and model assumptions by performing three types of sensitivity analyses:

1. **Alternative mortality data**: macro-longevity risk in the Lee-Carter model depends on the parameters in (2) and (6) that are calibrated using historical mortality data. We investigate the impact of alternative mortality data on welfare gains from risk sharing and corresponding risk-sharing rule.
2. **Alternative population compositions**: welfare gains from sharing macro-longevity risk also depend on the population composition. We will investigate the impact of alternative population compositions on welfare gains from risk sharing and corresponding risk-sharing rule.
3. **Alternative model macro-longevity risk**: instead of macro-longevity risk in the Lee-Carter model we assess the impact of alternative shocks in death rates on welfare gains from risk sharing and corresponding risk-sharing rule.

5.1.1 Alternative mortality data

Macro-longevity risk in the Lee-Carter model depends on the parameters in (2) and (6). In our main analysis we calibrate the parameters using historical mortality data of Dutch females. Using alternative mortality data changes the parameters and therefore also the size and distribution of macro-longevity shocks.

The welfare gains from risk sharing using alternative mortality data are presented in Table 3. We look at Dutch males, US females and US males. In case of a fixed retirement age or partial adjustment of the retirement age, the welfare gains do not change significantly. However, when the retirement age is fully adjusted welfare gains from risk sharing are lower compared to the mortality data of Dutch females. This especially holds for mortality data of US females. This lower welfare gain is caused primarily by lower volatility parameters in (7). A lower volatility implies smaller risk and therefore lower welfare gains from risk sharing.

Mortality data	Dutch females	Dutch males	US females	US males
Fixed retirement age (FRA)	0.3%	0.2%	0.2%	0.3%
Partial adjustment retirement age (PARA)	0.5%	0.5%	0.2%	0.3%
Full adjustment retirement age (FARA)	2.7%	1.9%	0.7%	1.5%

Table 3: Welfare gains in terms of aggregate certainty equivalent consumption after retirement from sharing macro-longevity risk measured on a 10-year horizon for alternative mortality data.

Figure 6 visualizes the optimal risk transfer relative to autarky as percentage of total risk. The black lines represent the optimal risk transfer rule using the mortality data of Dutch females and the grey lines for the alternative mortality data. We can conclude that for each retirement age policy the optimal risk transfer rule is robust to the alternative mortality data we consider.

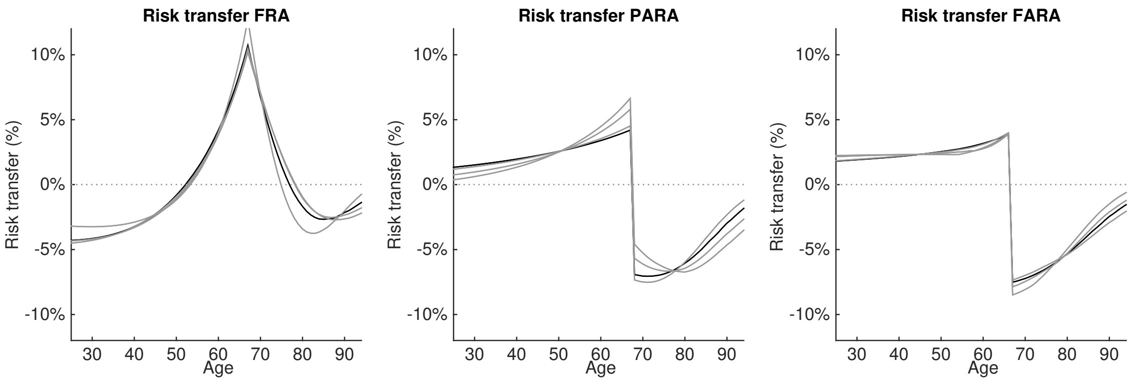

Figure 6: Optimal risk transfer relative to autarky as percentage of total risk in case of sharing macro-longevity risk measured on a 10-year horizon. The black lines represent the risk transfer rules based on Dutch females and the grey lines represent the risk transfer rules using alternative mortality data.

5.1.2 Alternative population compositions

We determined welfare gains in Table 2 and risk transfers and risk compensations in Figure 5 for a population composition of an entire country (lefthand graph in Figure 4). In practice the population composition of a pension fund is generally not equal to this standard population composition. Therefore, it is interesting to also consider alternative population compositions: a population composition of a green and grey pension fund. We assume that the green pension fund has a relatively young population. We approximate this by assuming that the number of participants in a cohort decreases with 1 percent per age year compared to the standard population composition. In the grey pension fund the number of participants in a cohort increases with 1 percent per age year compared to the standard population composition. The standard and alternative population compositions are displayed in Figure 7.

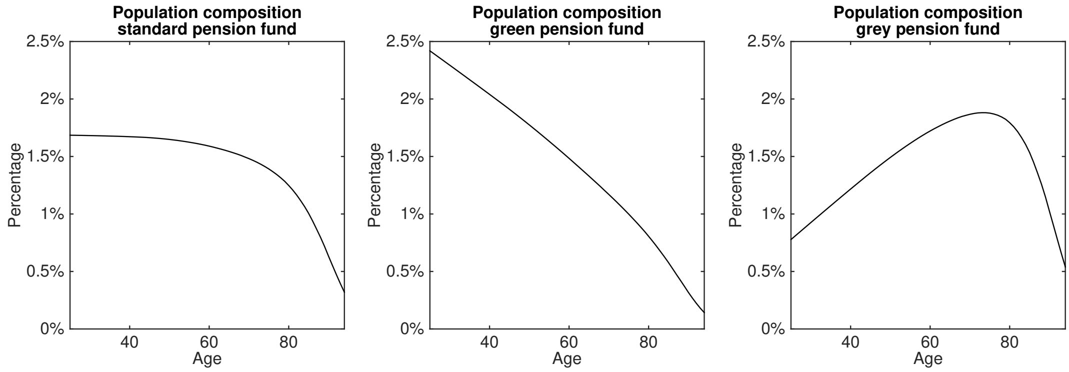

Figure 7: Different population compositions: a standard, green and grey pension fund.

Population composition	**Standard**	**Green**	**Grey**
Fixed retirement age (FRA)	0.3%	0.3%	0.3%
Partial adjustment retirement age (PARA)	0.5%	0.4%	0.7%
Full adjustment retirement age (FARA)	2.7%	2.2%	2.7%

Table 4: Welfare gains in terms of aggregate equivalent consumption after retirement from sharing macro-longevity risk measured on a 10-year horizon for alternative population compositions.

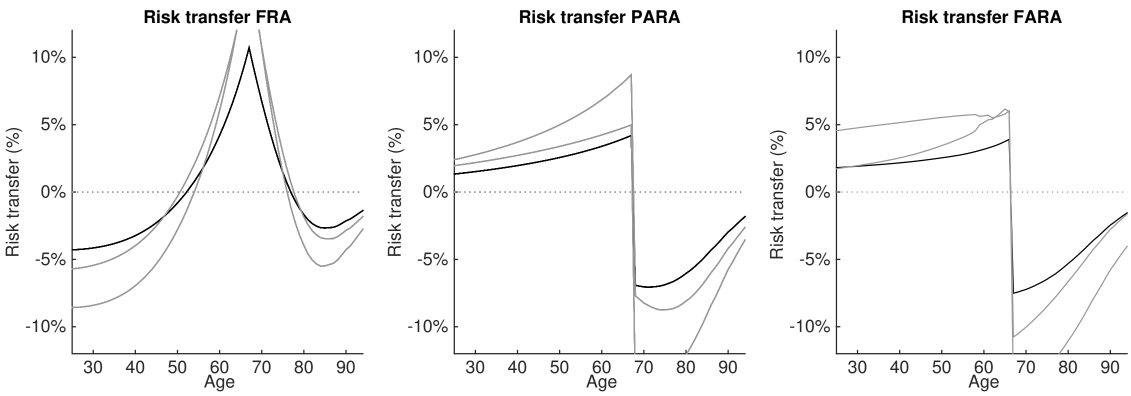

Figure 8: Optimal risk transfer relative to autarky for each cohort as percentage of total risk in case of sharing macro-longevity risk measured on a 10-year horizon. The black lines represent the original risk transfers and the grey lines represent the optimal risk transfers using alternative population compositions.

Table 4 presents the welfare gains from risk sharing using alternative population compositions. The welfare gains are not significantly different from the welfare gains for the standard population composition, even if the retirement age is fully adjusted. Figure 8 visualizes the optimal risk transfer relative to autarky as percentage of total risk. The black lines represent the optimal risk transfer rules using the original population composition and the grey lines represent the risk transfer rules using alternative population compositions. The shape of the risk transfer rule is reasonably robust to the population composition but the percentage of total risk an individual participant absorbs or transfers can be different in case of alternative population compositions. A different population composition leads to a different ratio between the individual macro-longevity shock and total macro-longevity shock. This impacts the optimal risk transfer as percentage of total risk.

5.1.3 Alternative model macro-longevity risk

Several academics use the Lee-Carter model to model macro-longevity risk. Moreover, it is the basis of several mortality table forecasts in practice. However, the model is not a perfect representation of reality because there is uncertainty about structural breaks. For example, medical innovations can cause structural breaks that are not captured by the Lee-Carter model. Therefore it is interesting to also look at the impact of alternative shocks in the death rates.

There is no scientific consensus on the development of future survival probability at old ages. Buettner (2002) suggests that there are two alternative views about the future survival probability at old ages: compression versus expansion. In case of mortality compression mortality continues to decline over a widening range of adult ages, but meets natural limits for very advanced ages. This development implies that the survival probability approaches a rectangle (Figure 9). Einmahl et al. (2017) and Dong et al. (2016) find evidence for the existence of a maximum age. In case of mortality expansion mortality continues to decline for all ages, i.e., there is no maximum age. Wilmoth (2000) and Oeppen and Vaupel (2002) argue that there is indeed no maximum age. Wilmoth (2000) states that, based on available demographic evidence, the human life span shows no sign of approaching a certain limit imposed by biology or other factors. There are even scientists who believe in the possible realization of longevity escape velocity. In this scenario death rates fall so fast that people's remaining life expectancy increases with time because therapies restore health faster than the rate of body deterioration due to biological ageing (De Grey (2004)).

The development of future mortality in the Lee-Carter model is in line with the mortality compression view. The sensitivity of the death rates to the time trend decreases in age x to almost zero at very high ages. An alternative shock in death rates is the macro-longevity shock in the Solvency II framework for insurers. The Solvency II capital requirements for

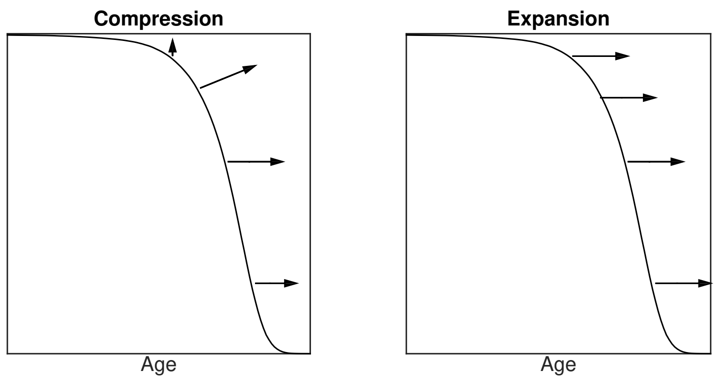

Figure 9: Different views of future survival probability: compression (lefthand graph) and expansion (righthand graph).

longevity risk are determined by applying a uniform shock, i.e., a 20 percent decrease, to all future death probabilities $q_{x,t}$.[29] For mortality risk the capital requirements are determined by applying an increase of 15 percent to all future death probabilities. The longevity shock in the Solvency II framework is in line with the expansion view because all death probabilities decrease at the same rate. Figure 10 visualizes both types of shocks, i.e., macro-longevity shocks in the Lee-Carter model and in the Solvency II framework. The graphs show that the development of future mortality in the Lee-Carter model is in line with the compression view and the Solvency II framework is in line with the expansion view.

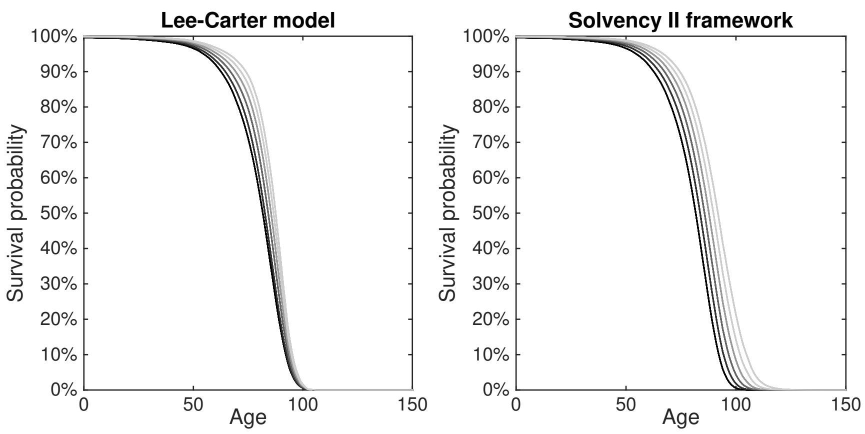

Figure 10: Impact of several consecutive macro-longevity shocks in the Lee-Carter model (lefthand graph) and in the Solvency II framework (righthand graph) on the survival probability.

The shocks for longevity and mortality risk in the Solvency II framework are deterministic, i.e., no stochastic mortality model is used to determine the distribution of future death rates. Because we have to make an assumption about the distribution of future death rates when

[29] These capital requirements are based on the 99.5% VaR of the available capital over a one-year horizon.

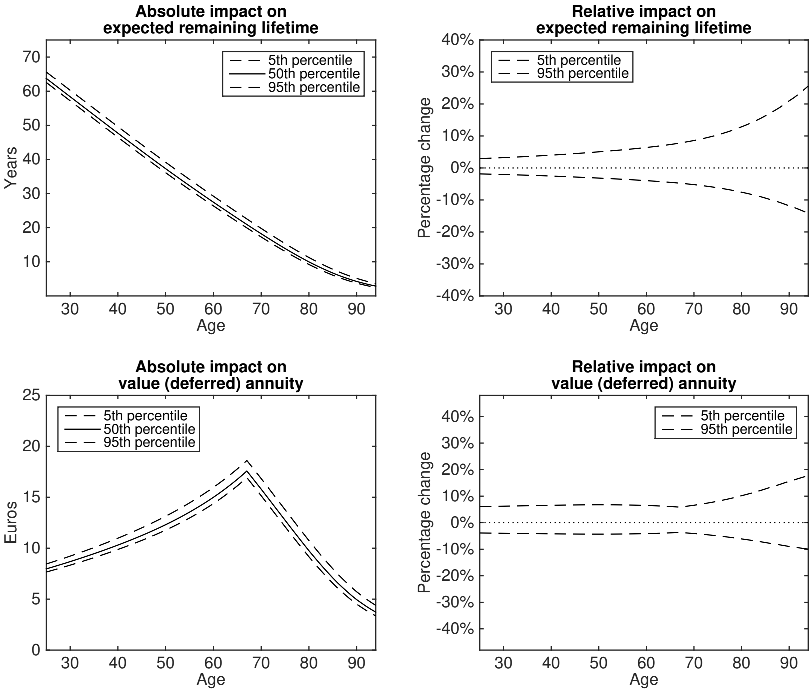

Figure 11: Impact of macro-longevity risk in the Solvency II framework on the expected remaining lifetime and the value of a (deferred) variable annuity for a Dutch female in 2014 in absolute terms (lefthand graphs) and relative change (righthand graphs) assuming a constant interest rate of 2% and fixed retirement age $R = 67$.

sharing macro-longevity risk, we assume that the shocks for longevity and mortality risk both occur with probability 50%. Figure 11 visualizes the impact of those shocks on the expected remaining lifetime and the value of a (deferred) variable annuity.

We cannot compare the size of the impact of macro-longevity risk in the Lee-Carter model (Figure 3) and Solvency II framework (Figure 11) directly, because the shocks in the Lee-Carter model are on a 10-year horizon while the shocks in the Solvency II framework are one-off shocks. However, we can still compare the distribution of macro-longevity risk over different cohorts in both models. We notice that the relative change of the expected remaining lifetime and (deferred) annuity value per cohort (righthand figures) differ significantly. While the relative change in the Lee-Carter model decreases with age, it increases with age in the Solvency II framework. This is due to the fact that the impact of a uniform improvement of death probabilities on survival probabilities is much higher at high ages compared to low ages because death probabilities are higher at high ages. As a result, the relative change increases with age in the Solvency II framework. In the Lee-Carter model the impact of macro-longevity risk on death probabilities decreases with age.

Model	LC	SII
Fixed retirement age (FRA)	0.3%	0.3%
Partial adjustment retirement age (PARA)	0.5%	0.3%
Full adjustment retirement age (FARA)	2.7%	0.4%

Table 5: Welfare gains from sharing macro-longevity risk in terms of aggregate certainty equivalent consumption after retirement in the Lee-Carter model and in the Solvency II framework.

Table 5 shows welfare gains from risk sharing in the Solvency II framework for the three retirement age policies. We cannot compare the size of welfare gains in the Lee-Carter model and Solvency II framework directly because both shocks have a different interpretation as mentioned above. In the Solvency II framework the welfare gain does not increase significantly in case of a full adjustment of the retirement age. Recall that the high welfare gain in case the retirement age is fully adjusted in the Lee-Carter model is a result of the hedge effect of the adjusted labor supply to macro-longevity shocks for workers. In the Solvency II framework the impact of macro-longevity risk on the expected remaining lifetime (Figure 11) is small for workers. As a result, the hedge effect is much smaller in the Solvency II framework compared to the Lee-Carter model.

Figure 12 visualizes the optimal risk transfer relative to autarky for a participant in cohort i as percentage of total risk. We can conclude that for each retirement age policy the optimal risk transfer rule is reasonably robust to the alternative mortality model.

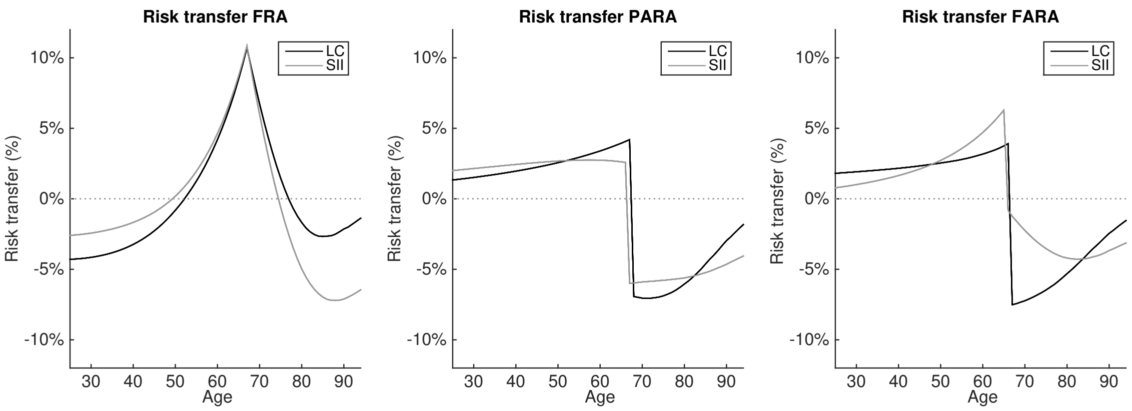

Figure 12: Optimal risk transfer relative to autarky as percentage of total risk for different retirement age policies. The black lines represent the original risk transfers in the Lee-Carter model and the grey lines represent the optimal risk transfers in the Solvency II framework.

6 Conclusion and policy evaluation

Pension funds face macro-longevity risk or uncertainty about future mortality rates. We analyze macro-longevity risk sharing between cohorts in a pension scheme as a risk management tool. We explore this economic problem as macro-longevity risk is not traded on a liquid market and cohorts are affected differently by macro-longevity risk. We derive the optimal risk-sharing rule and welfare gains from ex-ante Pareto-optimal risk-sharing rules for different retirement age policies.

We find that the design of the retirement age policy has a large impact on both the optimal risk-sharing rule and welfare gains from sharing macro-longevity risk. When the retirement age is fixed, welfare gains from sharing macro-longevity risk on a 10-year horizon are between 0.2 percent and 0.3 percent of certainty equivalent consumption after retirement. Under this policy, the impact of macro-longevity risk on retirement consumption for different cohorts is more or less equal. Young cohorts do not absorb macro-longevity risk of old cohorts in the optimal risk transfer rule. As a result, welfare gains from risk sharing are limited.

Some countries link the retirement age to life expectancy developments. If the retirement age is linked to life expectancy, welfare gains from sharing macro-longevity risk measured on a 10-year horizon are substantially higher, up to 2.7 percent. The risk bearing capacity of workers is larger, because they can use their labor supply as a hedge against macro-longevity shocks. As a result, workers absorb risk from retirees in the optimal risk transfer rule because the human capital of workers increases if they work longer. As a result, the welfare gain from risk sharing increases. The size of welfare gains from risk sharing is sensitive to the mortality data and model assumptions. This is a result of a different volatility of macro-longevity risk when using different mortality data and a different distribution of macro-longevity risk over cohorts. However, the optimal risk transfer rules are reasonably robust to the alternative mortality data and model assumptions.

The findings in this paper are relevant for pension policy, especially because of the general trend of transferring funding risks from pension scheme sponsors to pension participants (Munnell (2006) and Novy-Marx and Rauh (2014)). *First*, we determine the optimal risk-sharing rule for macro-longevity risk in this paper. In practice macro-longevity risk is shared in specific ways. In DB schemes and pooled annuity schemes, e.g., macro-longevity risk is usually shared uniformly. The results in this paper show that uniform risk sharing is suboptimal. Moreover, it is sometimes argued that workers can provide insurance to macro-longevity risk of retirees. The results in this paper show that such a risk distribution is optimal only in case the retirement age is linked to life expectancy. If the retirement age is fixed it is not optimal for young cohorts to absorb risk of retirees. *Second*, we determine a fair risk compensation for cohorts who absorb macro-longevity risk of other cohorts using a utility-based fairness criterion. In practice, there is usually no risk compensation for absorbing macro-longevity risk.

Sharing macro-longevity risk results in high welfare gains in case of a full adjustment of the retirement age. In this paper we do not make a statement about the suitability of retirement age policies. This is a different research question and involves a broader perspective. Healthy life expectancy and practical implementation are for example relevant but outside the scope of this paper. It is up to policymakers to decide whether it is appropriate to link the retirement age to life expectancy. The goal of this paper is to determine the optimal way to share macro-longevity risk between cohorts *given* a certain retirement age policy.

Sensitivity analyses show that the size of welfare gains depends on the population composition and the mortality data. For example, welfare gains from sharing macro-longevity risk are smaller for US mortality data compared to Dutch mortality data as a result of a lower volatility. An interesting area for future research is to investigate sharing macro-longevity risk between pension funds or even between countries. Van Binsbergen et al. (2014) propose sharing risks between heterogeneous pension funds by trading pension guarantees. Bodie and Merton (2002) propose swaps to achieve risk-sharing benefits of broad international diversification. Our framework is useful for further developing such instruments.

References

Aase, K. K. (2002). Perspectives of risk sharing. *Scandinavian Actuarial Journal*, 2002(2):73–128.

Ball, L. and Mankiw, N. G. (2007). Intergenerational risk sharing in the spirit of arrow, debreu, and rawls, with applications to social security design. *Journal of Political Economy*, 115(4):523–547.

Basel Committee on Banking Supervision (2013). Longevity risk transfer markets: market structure, growth drivers and impediments, and potential risks. *Bank for International Settlements*.

Blake, D., Boardman, T., and Cairns, A. (2014). Sharing longevity risk: Why governments should issue longevity bonds. *North American Actuarial Journal*, 18(1):258–277.

Blake, D., Cairns, A., Dowd, K., and MacMinn, R. (2006a). Longevity bonds: financial engineering, valuation, and hedging. *Journal of Risk and Insurance*, 73(4):647–672.

Blake, D., Cairns, A. J., and Dowd, K. (2006b). Living with mortality: Longevity bonds and other mortality-linked securities. *British Actuarial Journal*, 12(1):153–197.

Bodie, Z. and Merton, R. C. (2002). International pension swaps. *Journal of Pension Economics & Finance*, 1(1):77–83.

Boon, L.-N., Briere, M., and Werker, B. J. (2017). Longevity risk: To bear or to insure? *Working paper.*

Borch, K. (1960). The safety loading of reinsurance premiums. *Scandinavian Actuarial Journal*, 1960(3-4):163–184.

Bovenberg, L. and Mehlkopf, R. (2014). Optimal design of funded pension schemes. *Annual Review of Economics*, 6(1):445–474.

Brown, J. R. and Orszag, P. R. (2006). The political economy of government-issued longevity bonds. *Journal of Risk and Insurance*, 73(4):611–631.

Buettner, T. (2002). Approaches and experiences in projecting mortality patterns for the oldest-old. *North American Actuarial Journal*, 6(3):14–29.

Cairns, A. J. (2013). Robust hedging of longevity risk. *Journal of Risk and Insurance*, 80(3):621–648.

Cairns, A. J., Blake, D., and Dowd, K. (2006a). Pricing death: Frameworks for the valuation and securitization of mortality risk. *ASTIN Bulletin: The Journal of the IAA*, 36(1):79–120.

Cairns, A. J., Blake, D., and Dowd, K. (2006b). A two-factor model for stochastic mortality with parameter uncertainty: theory and calibration. *Journal of Risk and Insurance*, 73(4):687–718.

Cairns, A. J., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., and Khalaf-Allah, M. (2011). Mortality density forecasts: An analysis of six stochastic mortality models. *Insurance: Mathematics and Economics*, 48(3):355–367.

Chang, F.-R. (1991). Uncertain lifetimes, retirement and economic welfare. *Economica*, pages 215–232.

Cocco, J. F. and Gomes, F. J. (2012). Longevity risk, retirement savings, and financial innovation. *Journal of Financial Economics*, 103(3):507–529.

Cox, S. H. and Lin, Y. (2007). Natural hedging of life and annuity mortality risks. *North American Actuarial Journal*, 11(3):1–15.

Cui, J., De Jong, F., and Ponds, E. (2011). Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics & Finance*, 10(1):1–29.

De Grey, A. D. (2004). Escape velocity: why the prospect of extreme human life extension matters now. *PLoS Biology*, 2(6):e187.

De Waegenaere, A., Joseph, A., Janssen, P., and Vellekoop, M. (2018). Het delen van langlevenrisico. *Netspar Industry Paper*.

De Waegenaere, A., Melenberg, B., and Markwat, T. (2017). Risk sharing rules for longevity risk: impact and wealth transfers. *Netspar Industry Paper*.

Dong, X., Milholland, B., and Vijg, J. (2016). Evidence for a limit to human lifespan. *Nature*, 538(7624):257.

DuMouchel, W. H. (1968). The pareto optimality of an n-company reinsurance treaty. *Scandinavian Actuarial Journal*, 1968(3-4):165–170.

Einmahl, J., Einmahl, J. H., and de Haan, L. (2017). Limits to human life span through extreme value theory.

EIOPA (2011). *EIOPA Report on the fifth Quantitative Impact Study for Solvency II*. EIOPA.

Gollier, C. (2008). Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics*, 92(5):1463–1485.

Gordon, R. H. and Varian, H. R. (1988). Intergenerational risk sharing. *Journal of Public Economics*, 37(2):185–202.

Gottardi, P. and Kubler, F. (2011). Social security and risk sharing. *Journal of Economic Theory*, 146(3):1078–1106.

Hammond, R., Baxter, S., Bramley, R., Kakkad, A., Mehta, S., and Sadler, M. (2016). Considerations on state pension age in the united kingdom. *British Actuarial Journal*, 21(1):165–203.

Hari, N., De Waegenaere, A., Melenberg, B., and Nijman, T. E. (2008). Longevity risk in portfolios of pension annuities. *Insurance: Mathematics and Economics*, 42(2):505–519.

Heijdra, B. J. and Romp, W. E. (2009). Retirement, pensions, and ageing. *Journal of Public Economics*, 93(3-4):586–604.

Huang, C.-f. and Litzenberger, R. (1985). On the necessary condition for linear sharing and separation: a note. *Journal of Financial and Quantitative Analysis*, 20(3):381–384.

Hunt, A. and Blake, D. (2015). Modelling longevity bonds: Analysing the swiss re kortis bond. *Insurance: Mathematics and Economics*, 63:12–29.

Kannisto, V. (1994). Development of oldest-old mortality 1950-1990: evidence from 28 developed countries.

Krueger, D. and Kubler, F. (2006). Pareto-improving social security reform when financial markets are incomplete!? *American Economic Review*, 96(3):737–755.

Lee, R. (2000). The lee-carter method for forecasting mortality, with various extensions and applications. *North American actuarial journal*, 4(1):80–91.

Lee, R. D. and Carter, L. R. (1992). Modeling and forecasting us mortality. *Journal of the American statistical association*, 87(419):659–671.

Lin, Y., MacMinn, R. D., and Tian, R. (2015). De-risking defined benefit plans. *Insurance: Mathematics and Economics*, 63:52–65.

Munnell, A. H. (2006). Employer-sponsored plans: the shift from defined benefit to defined. *The Oxford handbook of pensions and retirement income*, 13:359–380.

Ngai, A. and Sherris, M. (2011). Longevity risk management for life and variable annuities: The effectiveness of static hedging using longevity bonds and derivatives. *Insurance: Mathematics and Economics*, 49(1):100–114.

Novy-Marx, R. and Rauh, J. D. (2014). Linking benefits to investment performance in us public pension systems. *Journal of Public Economics*, 116:47–61.

Oeppen, J. and Vaupel, J. W. (2002). Broken limits to life expectancy. *Science*, 296(5570):1029–1031.

Pelsser, A. (2011). Pricing in incomplete markets. *Netspar Panel Paper*.

Piggott, J., Valdez, E. A., and Detzel, B. (2005). The simple analytics of a pooled annuity fund. *Journal of Risk and Insurance*, 72(3):497–520.

Qiao, C. and Sherris, M. (2013). Managing systematic mortality risk with group self-pooling and annuitization schemes. *Journal of Risk and Insurance*, 80(4):949–974.

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

Stevens, R. (2017). Managing longevity risk by implementing sustainable full retirement age policies. *Journal of Risk and Insurance*, 84(4):1203–1230.

Van Binsbergen, J. H., Broeders, D., De Jong, M., and Koijen, R. S. (2014). Collective pension schemes and individual choice. *Journal of Pension Economics & Finance*, 13(2):210–225.

Wilmoth, J. R. (2000). Demography of longevity: past, present, and future trends. *Experimental gerontology*, 35(9-10):1111–1129.

Wilson, R. (1968). The theory of syndicates. *Econometrica: journal of the Econometric Society*, pages 119–132.

A Appendix

A.1 Expected survival probability

The random shocks in (7) in the log central death rates are normally distributed with mean zero, i.e., $\mathbb{E}[\beta_x \eta_t + \epsilon_{x,t}] = 0$. The following holds for the expected survival probability

$$\mathbb{E}[p_{x,t}] \approx \mathbb{E}[\exp(-\mu_{x,t})] = \mathbb{E}[\exp(-\exp(\alpha_x + \beta_x \kappa_t + \epsilon_{x,t}))] \tag{31}$$

$$\begin{aligned} \mathbb{E}[\exp(-\exp(\alpha_x + \beta_x \kappa_t + \epsilon_{x,t}))] &\leq \mathbb{E}[\exp(-\exp(\alpha_x + \beta_x c + \beta_x \kappa_{t-1} + \beta_x \eta_t + \epsilon_{x,t}))] \\ &= \exp(-\exp(\alpha_x + \beta_x c + \beta_x \kappa_{t-1} + \mathbb{E}[\beta_x \eta_t + \epsilon_{x,t}])) \end{aligned}$$

$$\exp(-\exp(\alpha_x + \beta_x c + \beta_x \kappa_{t-1} + \mathbb{E}[\beta_x \eta_t + \epsilon_{x,t}])) \leq \exp(-\exp(\alpha_x + \beta_x c + \beta_x \kappa_{t-1}) = \hat{p}_{x,t},$$

using Jensen's inequality $\mathbb{E}[f(x)] \leq f(\mathbb{E}[x])$ with $f(x) = \exp(-\exp(x))$ being a concave function for $x \leq 0$.

A.2 Derivation Pareto optimal risk-sharing rule in stylized two-agent model

The Lagrange function of the maximization problem in (20) equals

$$L(T, \lambda) = \mathbb{E}[U(W_1 + T_1(\tilde{y})) + \lambda(U(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y})) - U(W_2 + \beta_2 \tilde{y}))]. \tag{32}$$

Because $T_1(\tilde{y})$ is a continuous function we take the Fréchet-derivative

$$D_T L(T, \lambda).\tau = \mathbb{E}[(U'(W_1 + T_1(\tilde{y})) + \lambda U'(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y})))\tau(\tilde{y})]. \tag{33}$$

The first order condition should be zero for each pertubation $\tau(\tilde{y})$. This is only possible if the following holds for all values of $\tilde{y}$

$$U'(W_1 + T_1(\tilde{y})) = \lambda U'(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y})) \quad \forall \tilde{y}. \tag{34}$$

This is a non-linear equation which can be solved for $\{T_1(\tilde{y}), T_2(\tilde{y})\}$. As mentioned in Section 3.2 we assume both agents have exponential utility with risk aversion α. We plug this utility function into the first order condition

$$\begin{aligned} \exp(-\alpha(W_1 + T_1(\tilde{y})) &= \lambda \exp(-\alpha(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y})) \\ -\alpha(W_1 + T_1(\tilde{y})) &= \ln(\lambda) - \alpha(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y})) \\ -2\alpha T_1(\tilde{y}) &= \ln(\lambda) - \alpha(W_2 - W_1 + (\beta_1 + \beta_2)\tilde{y}) \\ T_1(\tilde{y}) &= -\frac{1}{2}\frac{\ln(\lambda)}{\alpha} + \frac{1}{2}\Big(W_2 - W_1 + (\beta_1 + \beta_2)\tilde{y}\Big). \end{aligned} \tag{35}$$

Because the utility function is strictly increasing, the inequality restriction in (20) is binding

$$\begin{aligned} \mathbb{E}[(U(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y}))] &= \mathbb{E}[U(W_2 + \beta_2 \tilde{y})] \\ \mathbb{E}\Big[-\frac{1}{\alpha}\exp(-\alpha(W_2 + (\beta_1 + \beta_2)\tilde{y} - T_1(\tilde{y})))\Big] &= \mathbb{E}\Big[-\frac{1}{\alpha}\exp(-\alpha(W_2 + \beta_2 \tilde{y}))\Big]. \end{aligned} \tag{36}$$

Plugging in (35) yields

$$\mathbb{E}\Big[-\frac{1}{\alpha}\exp\Big(-\alpha\Big(W_2+(\beta_1+\beta_2)\tilde{y}+\frac{1}{2}\frac{\ln\lambda}{\alpha}-\frac{1}{2}\Big(W_2-W_1+(\beta_1+\beta_2)\tilde{y}\Big)\Big)\Big)\Big]=\mathbb{E}\Big[-\frac{1}{\alpha}\exp(-\alpha(W_2+\beta_2\tilde{y}))\Big] \quad (37)$$

$$\mathbb{E}\Big[-\frac{1}{\alpha}\exp\Big(-\frac{1}{2}\ln\lambda-\alpha\Big(\frac{1}{2}(W_1+W_2)+\frac{1}{2}(\beta_1+\beta_2)\tilde{y}\Big)\Big)\Big]=\mathbb{E}\Big[-\frac{1}{\alpha}\exp(-\alpha(W_2+\beta_2\tilde{y}))\Big]$$

$$-\frac{1}{\alpha}\mathbb{E}\Big[\exp\Big(-\frac{1}{2}\ln\lambda\Big)\Big]\mathbb{E}\Big[\exp\Big(-\alpha\frac{1}{2}(W_1+W_2)\Big)\Big]\mathbb{E}\Big[\exp\Big(-\alpha\frac{1}{2}(\beta_1+\beta_2)\tilde{y}\Big)\Big]=-\frac{1}{\alpha}\mathbb{E}\Big[\exp(-\alpha W_2)\Big]\mathbb{E}\Big[\exp(-\alpha\beta_2\tilde{y})\Big]$$

$$\frac{1}{\sqrt{\lambda}}\exp\Big(-\frac{1}{2}\alpha W_1\Big)\exp\Big(-\frac{1}{2}\alpha W_2\Big)\exp\Big(\frac{1}{8}\alpha^2(\beta_1+\beta_2)^2\sigma^2\Big)=\exp(-\alpha W_2)\exp\Big(\frac{1}{2}\alpha^2\beta_2^2\sigma^2\Big)$$

$$\exp\Big(\frac{1}{2}\alpha(W_2-W_1)\Big)\exp\Big(\frac{1}{8}\alpha^2\beta_1^2\sigma^2+\frac{1}{4}\alpha^2\beta_1\beta_2\sigma^2-\frac{3}{8}\alpha^2\beta_2^2\sigma^2\Big)=\sqrt{\lambda}$$

$$\exp\Big(\alpha(W_2-W_1)\Big)\exp\Big(\frac{1}{4}\alpha^2\beta_1^2\sigma^2+\frac{1}{2}\alpha^2\beta_1\beta_2\sigma^2-\frac{3}{4}\alpha^2\beta_2^2\sigma^2\Big)=\lambda$$

Plugging λ into (35) yields

$$\begin{aligned}T_1(\tilde{y})&=-\frac{1}{2}\Big(W_2-W_1\Big)+\frac{3}{8}\alpha\beta_2^2\sigma^2-\frac{1}{8}\alpha\beta_1^2\sigma^2-\frac{1}{4}\alpha\beta_1\beta_2\sigma^2+\frac{1}{2}\Big(W_2-W_1\Big)+\frac{1}{2}\Big(\beta_1+\beta_2\Big)\tilde{y} \quad (38)\\&=-\frac{1}{8}\alpha\sigma^2(\beta_1+3\beta_2)(\beta_1-\beta_2)+\frac{1}{2}\Big(\beta_1+\beta_2\Big)\tilde{y},\end{aligned}$$

and the optimal risk-sharing rule for agent 2 equals

$$T_2(\tilde{y})=\frac{1}{8}\alpha\sigma^2(\beta_1+3\beta_2)(\beta_1-\beta_2)+\frac{1}{2}\Big(\beta_1+\beta_2\Big)\tilde{y}. \quad (39)$$

A.3 Definitions

Parameter	Definition
α_x	Age-specific constant in log central death rates
Annuity value (a_x^t)	Value of an annuity that pays 1 dollar annually during retirement for an individual of age x in year t
Autarky	Situation without risk sharing
C_i^a	Consumption after retirement in autarky for a participant in cohort i
C_i^s	Consumption after retirement in case of risk sharing for a participant in cohort i
Certainty equivalent consumption	Guaranteed consumption level that someone would accept rather than a higher uncertain consumption
Central death rate ($\mu_{x,t}$)	Average yearly death rate of an individual of age x in year t
Cumulative survival probability ($cp_{x,i}^t$)	Probability that an individual of age x in year t is still alive after i years
c	Drift in time trend
Equivalent variation (EQV_i)	Amount of wealth which agent i should be given ex-ante in autarky to obtain the same ex-ante welfare in case of risk sharing
Uncertainty in death rates ($\epsilon_{x,t}$)	Random variation in log central death rates
Fixed retirement age (FRA)	Constant retirement age
Full adjustment retirement age (FARA)	Retirement age keeps up fully with life expectancy
Longevity risk	Risk that people live longer than expected
Macro-longevity risk	Uncertainty about future mortality rates
Micro-longevity risk	Uncertainty about individual time of death
Mortality risk	Risk that people live shorter than expected
One-year death probability ($q_{x,t}$)	Probability that an individual of age x and alive in year t dies before year $t+1$
One-year survival probability ($p_{x,t}$)	Probability that an individual of age x and alive in year t is still alive in year $t+1$
Parameter risk	Uncertainty in the true value of the parameters
Partial adjustment retirement age (PARA)	Retirement age adjusts to life expectancy such that the value of an annuity remains the same
β_x	Sensitivity of log central death rates to time trend
Risk compensation ($t_{0,i}$)	Financial compensation for absorbing risk for a participant in cohort i
Risk sharing	Allocate risks to cohorts via a predetermined rule
Risk-sharing rule ($t(\tilde{y})$)	Risk transfer plus risk compensation
Risk transfer (η_i)	Part of total macro-longevity shock a participant in cohort i absorbs
Stochastic variation	Random variation in the aggregate realized number of deaths
Time trend (κ_t)	Development of death rates over time
Uncertainty in trend (η_t)	Random variation in the time trend
σ_ϵ^2	Variance death rates
σ_η^2	Variance trend
W_i	Wealth of a participant in cohort i
Welfare gain	Relative increase certainty equivalent consumption after retirement
$\tilde{y}_i$	Amount of money needed to offset effect of macro-longevity shock for a participant in cohort i
$\tilde{y}_T$	Amount of money needed to offset effect of macro-longevity shock for all cohorts