Intergenerational Risk Sharing under Loss Averse Preferences

Boes, Mark-Jan; Siegmann, A.H.

citation for published version (APA)
Boes, M-J., & Siegmann, A. H. (2018). Intergenerational Risk Sharing under Loss Averse Preferences. *Journal of Banking and Finance*, *92*, 269-279. https://doi.org/10.1016/j.jbankfin.2016.08.001

Intergenerational risk sharing under loss averse preferences

Mark-Jan Boes*, Arjen Siegmann

Vrije Universiteit Amsterdam, Faculty of Economics and Business, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

ABSTRACT

Individual retirement savings schemes could benefit from risk-sharing mechanisms between generations that take behavioral aspects into account. We introduce a novel risk-sharing mechanism that incorporates nominal loss-aversion in two ways. First, the system avoids out-of-pocket wealth transfers by sharing only a fraction of positive returns over a high-water mark of pension assets. Secondly, payments from a generation insurance fund are targeted at nominal pension shortfalls below a reference point, which mitigates the loss experience at retirement. From a simulation of overlapping generations with stochastic asset returns and interest rates we find that the generation insurance scheme outperforms a pure individual retirement scheme by a significant margin: a similar risk of pension shortfall can be achieved with a contribution rate that is up to 20% lower. The efficiency gains vary with the extent of risk sharing over generations but remain large for sensible parameter values.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Individual retirement accounts suffer from the problem that asset returns and interest rates are volatile and unpredictable. Under defined-contribution (DC) or cash-balance systems, this means that the participants run considerable risks of a less-than-desirable pension outcome. These risks can be mitigated by having a risk-sharing mechanism that spreads risks between generations, see Gollier (2008), Beetsma and Bovenberg (2009), Cui et al. (2011), Beetsma et al. (2012) and Novy-Marx and Rauh (2014). The efficiency is present in existing defined benefit (DB) plans, as long as the opacity of those plans is not used to camouflage underfunding and take excessive investment risks, see Andonov et al. (2016).

In contrast to the economic efficiency of intergenerational solidarity, public support for risk-sharing between generations is declining. The current "pension crisis" in the Netherlands, the introduction of individual retirement accounts in the UK, and the deterioration of defined-benefit (DB) schemes in the US are due to suspicions from old and young generations that the pension system does not benefit them to the extent that it should. From the perspective of limited rationality and trust in financial institutions, this is understandable: if participants cannot assess whether their pension contributions are beneficial to their own retirement outcome, support for the system declines.

An important consideration with respect to intergenerational insurance is loss aversion, i.e., people are more sensitive to losses than to gains. The reluctance to realize small losses leads to suboptimal choices in investment and savings decisions, see Odean (1998), Weber and Camerer (1998), Genesove and Mayer (2001), Feng and Seasholes (2005), Coval and Shumway (2005) and Frazzini (2006). Given the evidence, loss-averse preferences should feature in the design of a pension system, see also Broadbent et al. (2006).

In this paper, we analyze a generation-insurance pension scheme with a risk sharing mechanism that incorporates the concept of loss aversion. Participants are loss averse with respect to two anchors: (1) the final pension outcome and (2) their intermediate pension wealth. Our design allows that shortfalls of pension outcomes relative to a reference point of retirement benefits are mitigated through a generation insurance fund. Moreover, the proposed system offers the flexibility to collect insurance premia only in situations where intermediate pension wealth is at an all-time high.

The first feature, mitigating pension shortfalls, is established by having a generation-insurance fund that (partly) compensates participants for shortfalls in the pension outcome at the retirement age. The pension outcome is the level of benefits divided by the average lifetime wage and a shortfall is the extent to which the outcome is below 70%. So, a shortfall below 70% leads to a compensation payment from the insurance fund (with some maxima for the outflow of the fund), while a surplus above 70% is unchanged.

* Corresponding author.
 E-mail addresses: m.boes@vu.nl (M.-J. Boes), a.h.siegmann@vu.nl (A. Siegmann).

The second feature, avoiding out-of-pocket contributions to generation insurance, is accomplished by using a return-sharing mechanism: insurance premiums are collected as a fraction of positive returns, but only when assets are above a high-water mark. So, premia are only paid in times which participants will unequivocally perceive as "good": returns are positive, and the assets in the individual account are exceeding the previous highest level. This makes the insurance fund an attractive proposition for participants who are averse to suffering nominal losses.

Collecting premia by "shaving off" some of the investment return in good times is similar to the Save More Tomorrow scheme of Thaler and Benartzi (2004). In their scheme, employees are induced to save more by having a default fraction of wage increases being put in a savings account. It assumes employees are loss averse in nominal terms, so that they are reluctant to a decrease in their nominal income, see Kahneman and Tversky (1979), Kahneman et al. (1991). Since the contribution to the savings plan is only increased after a wage rise, the change in take home pay for a worker remains non-negative.

A problem with existing mechanisms for risk-sharing in the pension domain is their opacity: participants cannot easily compute what they pay into the system and what they get from it. Van Rooij et al. (2011) and Van Rooij et al. (2012) find that the extent of financial literacy (understanding basic concepts) has a significant impact on retirement planning and savings outcomes. Brown and Weisbenner (2014) find that many people hold wrong beliefs about pension plan parameters, which hurts their retirement outcome. This problem is even worse if free financial advice is not used by the people who need it most, see Bhattacharya et al. (2012). One of the advantages of our proposed system is that the features of the risk-sharing mechanism are transparent and easy to understand. The insurance scheme can be explained as being similar to health insurance, where everyone pays contributions, but payments are only made when a participant incurs treatment costs. Likewise, the generation insurance fund takes in insurance premiums from all participants, but only pays out in case of pension shortfalls. The transparency and the understandability of the proposed system helps people in their personal financial planning.

The system we analyze can be characterized as a collective individual defined contribution (CDC) scheme. The scheme resembles aspects of DB-schemes such as collective risk-sharing and mandatory saving. Moreover, people save individually and, consequently, there should not be discussions about ownership rights of capital like those that occur within DB-schemes. Obviously, in our setup discussions could possibly emerge about the ownership of the insurance fund but the clarity of having a given size of the insurance fund facilitates understanding of the trade-offs involved.

To test the efficiency of the risk sharing model we perform simulations of the performance of our plan using stochastic asset returns and interest rates in an overlapping-generations model. The approach of a Monte-Carlo simulation enables us to model a realistic evolution of returns, contributions and payouts of the insurance fund. We evaluate pension outcomes with loss averse value functions as well as a classical power utility function, to have a broad view of the gains in the efficiency of pension outcomes.

The simulation outcomes show that under a linear loss-averse utility function our generation insurance scheme with a contribution rate of 10% achieves the same certainty equivalent pension as an individual DC-scheme with a contribution rate of 10.8%. If only downside risk on the final pension outcome is important then a contribution level of 12.7% in an individual system is needed to match the downside risk of our insurance scheme. These outcomes confirm the benefits to using a generation insurance fund that spreads risks over generations.

The difference in contribution rates depends on the parameters used for the generation insurance scheme and the risk measure used to match pension outcomes, but remains large and significant in all cases. Moreover, we show that the outcomes are not driven by a nonlinear shift in wealth between generations. The results confirm that a simple scheme that uses behavioral insights for its design can achieve a large share of the benefits that come from sharing risks between generations.

The paper proceeds as follows. Section 2 describes the model. Section 3 presents the simulation results. Section 4 discusses implementation issues. Section 5 concludes.

2. Modeling pension outcomes with risk sharing

We model the participant to a retirement scheme as an agent who cares only about the pension annuity at retirement. He is not indifferent to higher pension contributions, but we evaluate that separately, from the perspective of a social planner. This setup enables us to isolate the risk-return properties of the risk-sharing scheme while avoiding complexities such as the interaction between retirement saving, consumption, health and housing wealth.

2.1. The model

The basic building block of the model is an individual retirement account. Participants pay a fixed contribution rate as a fraction of wages which accumulates in the retirement account with investment returns. The life-cycle investment mix is set by the fund and is identical for each individual of a certain age. The add-on component is an insurance fund that obtains premia from the asset returns of participants above a high-water mark, and pays out a compensation for each retiring generation who faces a pension shortfall below 70% of the average wage. The payout of the insurance fund is capped at a percentage of total assets of the fund, so that the fund is sustainable across generations.

We model $M + T$ overlapping generations that each have an active working life of T years, where M generations reach retirement. A new generation enters the working population in each year. We use t_i to denote working time for generation $i = 1, \ldots, M$. Calendar years evolve with a time index τ and since time starts at 0, we have that $\tau = t_i + i - 1$. Each generation i consists of n_{it}^{τ} employees at time τ. For notational clarity, the equations that model the evolution of a variable specific to one generation, we drop the subscript i of t_i.

A generation i at working age t builds up pension assets A_{it}, as follows:

$$A_{it} = A_{i,t-1} \cdot (1 + R_{\tau}^p) + c \cdot W_{it}, \qquad (1)$$

and

$$A_{i0} = 0 \qquad (2)$$

where R_{τ}^p is the portfolio return at calendar time τ and c is the contribution rate to the individual retirement account, and W_{it} is the wage of generation i at age t. The portfolio return is the result of a fraction invested in stocks and bond, the dynamics of which are specified below.

The fraction in stocks α_t is initially 100% and declines from year t^*. In the baseline simulation, t^* is such that the fraction stocks declines with 10%-points each year from age 50, reaching 0% stocks at age 60. In the simulations we analyze the sensitivity to this parameter. This approach follows from a life-cycle perspective, where labor income becomes a decreasing proportion of lifetime wealth with age,[1] see Campbell and Viceira (2003). We test for

---

[1] In practice, the pension scheme could work with a sensible default option with an opt-out clause to offer the opportunity to adjust the investment mix to personal tastes.

the sensitivity of the outcomes to the riskiness of the asset mix in Section 3.4.

The real wage profile $\tilde{W}_t, t = 1, \ldots, T$, is the same for every generation, i.e.,

$$W_{it} = \tilde{W}_t, \quad \text{for } i = 1, \ldots, M, \text{ and } t = 1, \ldots, T \tag{3}$$

and the average real wage is $\bar{W}$. We specify the exact age-related wage profile by using an average profile for a UK employee, see below.

Starting in year T, pension benefits are obtained by a single retiring generation. In year $\tau = T + i - 1$, the pension benefit PB_τ of the retiring generation is equal to capital divided by the average wage, $\bar{W}$, and the costs of a pension annuity, i.e.,

$$PB_\tau = A_{iT} / \left(\bar{W} \cdot a_{r_\tau} \right) \tag{4}$$

where a_{r_τ} is the annuity cost for receiving 1 in every year of retirement, using the interest rate r_τ at time τ and the survivorship table for discounting the future payments. The variable PB_τ is a pension level in terms of the fraction of average wage.

The target level of pension benefits is PB_{target} so that a pension shortfall for generation i at time $\tau = T + i - 1$ is defined as

$$PSF_\tau = \max \left(PB_{target} - PB_\tau, 0 \right) \tag{5}$$

with monetary value $PSF_\tau \cdot a_{r_\tau} \cdot \bar{W}$.

The inflow into the insurance fund consists of two components: share of excess returns and a fixed insurance contribution. The share of excess returns is a fraction of positive returns on the participant's assets above the "high-water mark", where high-water mark (HWM) is simply the previous highest level of asset. Using the HWM ensures that insurance contributions only occur when the pension assets are higher than before. Taking a fraction from positive returns ensures that declines in asset values are not made worse, and that people contribute only in good years, i.e., when returns are positive.

The high-water mark is a natural reference point, as people use historical highs to form a reference point beyond which they decode outcomes as losses. Investors of an acquiring firm are using similar frames when considering a takeover bid of a company, see Baker et al. (2012). In the hedge fund industry, high-water marks are used in the computation of performance fees, so that investors only pay these fees when assets are higher than previously reached highs and not otherwise, see Goetzmann et al. (2003).

The mechanism for collecting premia as a fraction of returns above the high-water mark is similar in spirit to the "Save More Tomorrow" program of Thaler and Benartzi (2004). Save More Tomorrow defers the saving decision to the moment of a wage rise. At that point, a fraction of the wage increase goes into the savings account, so that take-home pay is not decreasing due to increased savings. Thaler and Benartzi find that the scheme works well, with 80% of employees joining, whose subsequent savings rate increases by 10%-points over time. The reason for the success is that people are much more comfortable in putting aside a fraction of their wage for their "future self" when it does not hurt their take-home pay. The smart design of the savings system overcomes potential self-control problems, triggered by the nominal losses from an ad hoc increase in savings. Our system uses a similar principle: people give up a small fraction of the upside for their future self, in expectation.

In the model, the return-related premium is a fraction κ of any positive asset growth in the pension account that exceeds the high-water mark. The high-water mark H for generation i at age t is given by

$$H_{it} = \max \left(H_{i,t-1}, A_{it} \right). \tag{6}$$

Growth of the high-water mark is driven by growth in pension assets. Pension assets increase either through positive investment

returns or by the contribution in the personal retirement account. The return-related insurance premium is only paid when the increase in the high-water originates from positive investment returns, i.e., participants of the pension scheme only contribute in times that are perceived to be good.

A second source of inflow to the insurance fund is a fixed insurance fee f as a fraction wages and a direct payment from the participant. It fits into the frame of other types of insurance where a positive premium is paid and a payout is triggered in adverse scenarios. Given the loss-averse framework that we use, a direct insurance cost is not favored.[2] However, we include it to compare the effects to the return-based premium. The total insurance contribution, for generation i at time t is given by

$$C_{it} = f \cdot W_{it} + \begin{cases} \kappa R^p_{i+t-1} A_{i,t-1} & \text{if } A_{t-1} \geq H_{t-1} \text{ and } R_{it} > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{7}$$

In practice, the insurance premium, through either f or κ, could be made dependent on the asset level of the insurance fund. This might increase efficiency further, but at the costs of added complexity and discretion by the pension fund board, which we try to avoid. For the aim of the paper, we leave f and κ fixed.

Given the systematic risk of insurance pay-outs, the insurance fund should invest in assets that have a low (or negative) correlation with the investment returns of participants. This could be a combination of low-risk assets, such as government bonds or money market funds, assets with low beta and volatility, and possibly assets with a negative correlation with risky assets, such as funds and securities with a short exposure to the stock market. In our setup, the fund simply investments in short-term money market instruments: The assets of the insurance fund yield the short-term interest rate, r_τ, so that the evolution of the asset value of the insurance fund, IF_τ is given by

$$IF_\tau = IF_{\tau-1} \cdot (1 + r_\tau) + \sum_{i \in G_\tau} C_{it_i}, \tag{8}$$

where G_τ contains the active generations at time τ and $t_i = \tau - i + 1$. Note that, at each year $\tau = 1, \ldots, T$, the number of active generations increases with 1. After year T, the number of active generations remain constant, as one generation retires and another enters the workforce.

If a generation has a pension shortfall, the insurance fund pays out a fraction of its assets to (partly) compensate for the shortfall. Only up to a maximum fraction F of the insurance fund is paid in any year, so that the fraction of the pension shortfall paid is

$$IPP_\tau = \min \left(PSF_\tau, \frac{IF_\tau \cdot F}{a_{r_\tau} \cdot \bar{W}} \right) \tag{9}$$

If generation i has a pension shortfall at its retirement year $\tau = T + i - 1$, it receives a compensation that is at maximum equal to the shortfall, but less if the maximum amount available, $F \cdot IF_\tau$, is not enough to cover the shortfall. The payout structure of the insurance fund is similar to that of Beetsma et al. (2012), who model a linear payout below a threshold of consumption.

In the overlapping generations model, the first payout occurs when the first generation retires, i.e., at $\tau = T$. From that point onwards, the sum of wages is constant and insurance payouts can be made each year, depending on whether the retiring generation has a pension shortfall.

With the share of the payout for an individual being proportional to the monetary value of the transfer, there is no significant redistribution across low-growth and high-growth wage profiles. This is due to the fact that pension capital accumulated in

---

Panel A: Survival of retirees   Panel B: Wage profile

Fig. 1. Demographics and wage profile.
This figure shows the demographical distribution and wage profile that is used for the generations in the model. They are obtained from the Office for National Statistics (ONS). Panel A has the survival probabilities as of age 65 from the 2060 Projection of survival probabilities. (We assume 100% survival until age 68, see the text). Panel B has the UK wage profile 1997–2012, and represents the average yearly wage (in US-dollars) for a worker of each age, using an exchange rate of $1.50 per pound.

T years is roughly proportional to the average contribution. Our simulations support this (not reported). The rule of payout share equals capital share is easily communicated and is fair: low-growth earners have contributed less over their working life and will also receive less from the fund. Under the rule, the payout in terms of final pension will be similar to what they could have received when being in a separate pension fund with similar low-growth earners.

With the compensation from the insurance fund, the adjusted pension, PBA_τ for generation i becomes

$$PBA_\tau = PB_\tau + IPP_\tau \qquad (10)$$

The model results in M generations with adjusted pension benefits for which we can compute the utility. Generations $M+1, \ldots, M+T$ are active in paying contributions but do not reach retirement at the planning horizon of the model.

2.2. Age and wage profile

Generations enter the workforce at age 23 and retire at 68, i.e., $T = 45$. A retirement age of 68 is relatively high for current standards, but not implausible given the planned rises in statutory retirement ages in the UK and other countries.

To simplify the computation of the overlapping generations model, we assume a 100% survival rate until age 68. Then, for the computation of the cost of a pension annuity at 68, we use the predicted UK survival probabilities in 2060 as provided by the Office for National Statistics (ONS), see Fig. 1, panel A.

For the wage profile we take the average UK wage per age group in 2004–2012, as provided by the Annual Survey of Hours and Earnings, also from the ONS. The survey measures wages per cohort of 10 years which we interpolate and approximate using a quadratic polynomial, see Fig. 1.

In reality, wage growth is not deterministic, and possibly correlated with stock returns. However, our results in terms of the relative advantage of generation insurance remain qualitatively similar when we take this into account (results available upon request). The intuition for this is that wage growth paths and the correlation structure affect the outcomes for individual DC and generation

insurance in similar ways, so that the relative difference does not change significantly.

The wage profile in panel B of Fig. 1 is representative of wages increasing over the lifetime, leveling off, and slightly decreasing at higher ages. The decrease in wages at an older age can be seen as early retirement by a fraction of the population, accepting a lower-paid job, or working part-time.

2.3. Economic parameters and distributional assumptions

To facilitate the aggregation and comparison of outcomes across different generations, we assume that real (corrected for inflation) wage and price growth are zero. We choose parameters for interest rates and the risky returns accordingly, i.e. reflecting a real setup with zero inflation and adjusting the mean asset returns accordingly.

Stock returns are normally distributed with a real mean of 6% and volatility of 20%, which are similar to what is typically found in the literature, see Guidolin and Timmermann (2005). The real interest rate, r_τ, is normally distributed with a mean of 2% and standard deviation of 1% and is independently drawn from year to year, i.e., interest rates are i.i.d. This is the most simple representation of interest-rate uncertainty but quite sufficient for our purposes. Increasing the sophistication of modeling the short- and long-term interest rate will improve the realism of pension outcomes, but does not change the comparative performance of the two pension schemes that we analyze.

Bond returns are related to changes in the interest rates. Assuming a duration of Δ, we have a direct relation between bond returns and short-term interest rates r_τ, given by

$$R_\tau^b = 0.01 + r_\tau - \Delta \cdot (r_\tau - r_{\tau-1}) \qquad (11)$$

where the first term in (11) implies a carry return, or term premium of one percent. The relation in (11) is similar to that of Hoevenaars et al. (2008). For the rest of the paper we use a bond duration of 5.

Historically, stock returns and interest rates have shown prolonged periods of zero, positive or negative correlation, see Aslanidis and Christiansen (2014). However, we refrain from developing a model for conditional dependence or a view on the

correct correlation parameter and choose to leave the correlation at zero.

In our simulation setup we could accommodate for stochastic wage growth, correlated with stock returns, as in Lucas and Zeldes (2006), but refrain from doing so to avoid unnecessary complexity. Moreover, we find that modeling stochastic wage growth does not impact the outcomes significantly (results are available upon request).

For the annuity costs of a pension, we assume that the term-spread is equal to the inflation rate, so that the current short-term interest rate is the annuity interest rate for computing the cost of a pension annuity at retirement. The computation for the pension annuity is in Eq. (4) above.

2.4. Simulation setup

We compute the performance of the pension plan using a Monte-Carlo simulation in an overlapping generations model with yearly generations. We focus on the retirement outcome only, abstracting from labor income after retirement, consumption patterns or housing wealth that could influence the possible pension outcomes and living standards for an individual, but are beyond the scope of this analysis.

To assess the performance of generation insurance, the simulation outcomes are evaluated against the outcomes of a purely individual system. The purely individual retirement scheme has no risk sharing between generations and resembles the 401(k) system in the US or the recently re-designed second pillar retirement scheme in the UK. The simulation results for the individual-DC have the same per-scenario asset returns and asset mix as generation insurance and is obtained in the model above by setting the all insurance parameters to zero, i.e., there are no contributions to and payouts from the generation insurance fund.

To illustrate the dynamics of the pension plan, sample outcomes for two simulation runs are shown in Fig. 2.

The two simulation runs exemplify how the generation insurance leads to a lower downside risk (below 0.7 of average wage) at the cost of a lower upside potential. Both panels have higher maximums for individual DC pension outcomes, but also lower minima. The generation insurance scheme provides protection for pension outcomes below 0.7.

2.5. Evaluating pension outcomes

To evaluate the pension outcomes for different risk sharing parameters we assume the pension fund participant maximizes the utility of final pension. Final pension is the pension benefit in terms of the fraction of average wage, adjusted with a potential insurance payout, i.e., PBA_τ.

We consider loss-averse utility functions and the classical power utility function. Loss averse utility functions put an explicit penalty on realizations below a reference point and are rooted in behavioral research on decision making under uncertainty, see Kahneman and Tversky (1979). Evidence for loss aversion in many economics settings as well as among finance professionals, see Thaler and Benartzi (2004), Abdellaoui et al. (2013) and Barberis (2013).[3]

The loss averse utility functions evaluate the value of the pension outcome, PBA_τ, relative to the reference point of 70% of average wage, i.e., the objective is

$$\max v(PBA_\tau - 0.70) \tag{12}$$

---

[3] The most complete framework for decision making is prospect theory, which also involves the distortion of probabilities. However, for the purpose of the objective analysis we only consider loss aversion here.
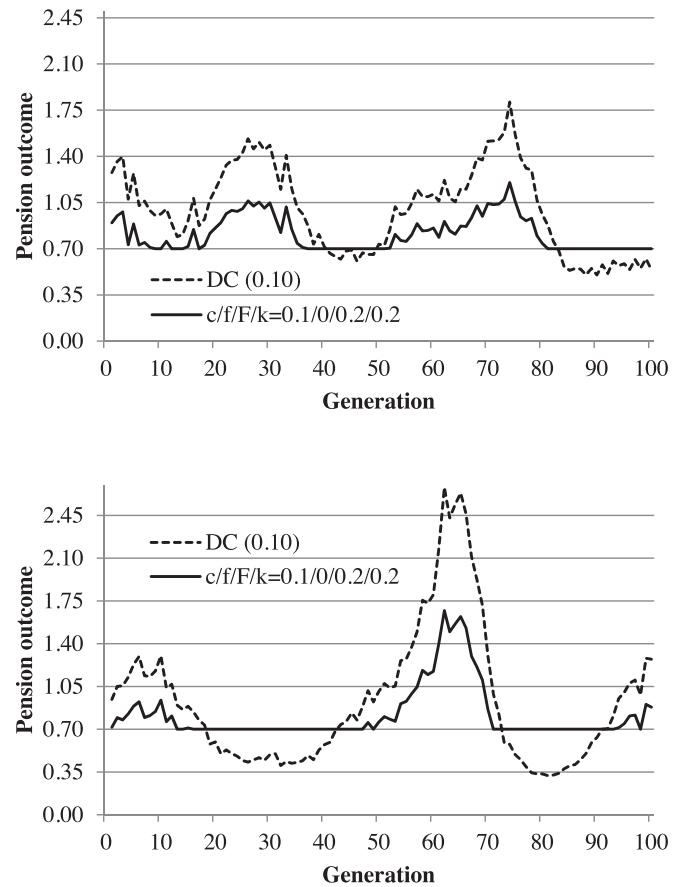

Fig. 2. Two example outcomes.
The figure shows the pension outcomes for two simulation runs. The solid lines represent the outcomes of the generation insurance model with 10% contribution, 20% return sharing and 20% maximum payout (as a fraction of the insurance fund). The dotted lines are the individual DC pension outcomes (without insurance) with 10% contribution.

where $v(\cdot)$ is a value function. The first loss-averse value function is the linear formulation given by

$$v(x) = \begin{cases} x & \text{if } x \geq 0 \\ \lambda_1 \cdot x & \text{if } x < 0, \end{cases} \tag{13}$$

where $\lambda_1 > 1$ is the degree of loss aversion, so that outcomes below 70% of average wage obtain an extra penalty. The kinked-shaped utility function is the most succinct way of modeling loss-averse preferences and is used by for example Barberis et al. (2001) to explain the equity premium. It is the linearized version of the original by Tversky and Kahneman (1992), who find risk-seeking behavior in the loss domain. We choose a loss aversion parameter $\lambda_1 = 10$ that is rather high relative to Tversky and Kahneman's value of 2.25, to account for the large disutility of below-target pension outcomes.

A potential concern with the linear loss averse utility function in (14) is that two small shortfalls are penalized similarly to one shortfall that is twice as large, i.e., it could have too little weight on large shortfalls. To take this into account, an alternative formulation uses a quadratic penalty on losses:

$$v(x) = \begin{cases} x & \text{if } x \geq 0 \\ x - \lambda_2 \cdot (a \cdot x)^2 & \text{if } x < 0, \end{cases} \tag{14}$$

where λ_2 is the coefficient of quadratic loss aversion. The use of a quadratic penalty is similar to the use of downside deviation as a risk measure in asset-liability management for pension funds, see

Boender et al. (2007). The parameter a is used to scale the outcomes so that pension outcomes below 0.65 carry a larger penalty than in the linear case, i.e., to adjust for the fractional nature of the pension outcomes. The loss aversion parameter λ_2 is again 10.

As a third loss averse specification we consider is the *probability of a pension shortfall* versus the reference point of 70% of average wage. This is the most simple representation of loss averse preferences, has a straightforward interpretation and is used by pension funds in the Asset/Liability Management process.

A potential problem with loss averse utility functions is that they lead to risk-seeking behavior in the loss domain. Even for the linearized and quadratic functions, Siegmann and Lucas (2005) find that risk-taking increases with the distance to the reference point, both in gains and in losses. It is caused by the local curvature of the loss aversion functions, that is either linear or quadratic, but not sufficiently curved to mitigate risk taking when the distance to the reference point increases. To avoid this problem we also consider a classical power utility function

$$U(x) = \frac{1}{1-\gamma} \cdot x^{1-\gamma} \qquad (15)$$

where γ is a risk aversion parameter and utility is computed over pension outcomes, i.e., the PBA_τ in Eq. (10) . Power utility is one of the most commonly used utility functions and can fit data well in many domains, see Wakker (2008). Power utility exhibits Constant Relative Risk Aversion (CRRA), so that the relative curvature of the utility function is constant across wealth. This is consistent with the idea that relative (percentage-wise) changes in wealth matter. It is closely related to the manipulation-proof performance measure of Goetzmann et al. (2007).

For the coefficient of relative risk aversion, Dalal and Arshanapalli (1993) find a value of 1.3 based on the holdings of risky assets by US-households. Guiso et al. (2013) find values distributed uniformly between 1 and 10, based on survey responses about taking a gamble that delivers either €0 or €10,000, with equal probability. However, relative risk aversions higher than 5 are seen as extremely risk averse in real-life situations. Chiappori and Paiella (2011) find a median value of 3 based on the stock holdings of Italian households, and we take that value, i.e., 3, to represent a median level of risk aversion in evaluating the power utility of pension outcomes.

In the presentation of the results, we convert utility outcomes to certainty-equivalent utility. This is the value of the pension benefit PBA_τ that gives the same utility as the average utility over all (risky) outcomes. It facilitates interpretation and provides a useful yardstick to compare the different outcomes.

3. Performance evaluation

3.1. Baseline results

As described above, we assess the performance of generation insurance we compare the descriptive statistics and utility outcomes to that of the purely individual system, i.e., with $f = \kappa = 0$. Table 1 shows our baseline results based on 1000 simulations of 100 generations.[4].

The top row in Table 1 shows the results for the individual retirement (DC)-scheme with a 10% annual contribution as a percentage of wages. This is the individual scheme against we measure the performance of the generation insurance scheme. It has an average pension outcome of 0.83 of average lifetime wage, with

a median of 0.67, but comes with large downside risks: the 5%-quantile of outcomes is 0.29 and the probability of a shortfall is 53%.

The outcomes for the generation insurance scheme are shown in the subsequent rows, grouped by three sets of parameter values. The first group of rows varies the extent of return sharing, from 0.1 to 0.3 of positive returns over the high-water mark of assets. It leads to between 5 and 16 percentage-points lower mean outcomes than individual DC, but at a significant reduction in risks. For example, the probability of a shortfall is reduced from 53% to 37% for a return sharing of 20%. Moreover, the median pension outcome is 0.70, which is even slightly higher than the median outcome of 0.67 for individual DC.

The columns labeled 'LLA' and 'QLA' give the certainty equivalent values for linear and quadratic loss aversion, respectively. For both LLA and QLA the improvements relative to individual DC are three and four percentage points of final pension, depending on the parameters.

The final column gives the certainty-equivalent (CEQ) pension outcome under a power utility function with a risk aversion parameter of 3. The results show considerable improvements in CEQ of the generation insurance scheme relative to the individual scheme. The individual DC scheme has a CEQ of 0.52, whereas generation insurance has a CEQ of 0.57 in the first group of rows.

In the second group of simulations (rows 5–7) we vary the maximum payout that is done by the generation insurance fund. Setting it too low reduces the protective power of the fund. Setting it too high could risk emptying the insurance fund, so that later generations have no benefit from it anymore. The outcomes show that for parameter values 0.1, 0.2 and 0.3, the outcomes taken over all simulation runs vary only little.

The bottom six rows of Table 1 show the results if we allow for direct insurance contributions. We observe that the improvements in performance are similar to those with only return sharing. For example, the downside risk measures of a system with 1% insurance fee, 20% return sharing and 20% maximum payout are not better than the risk characteristics of a similar scheme without a fixed insurance fee. At the same time, the mean pension outcome (0.69) is lower than under return sharing without a fixed fee (0.72). These results provide added motivation to use return-sharing, besides the behavioral motive. For the remainder, we consider parameter sets with $f = 0$, so that out-of-pocket insurance premia are avoided.

In all, the baseline outcomes in Table 1 show that, for reasonable values of return sharing and maximum payouts, the riskiness of pension outcomes is greatly reduced. This comes at the cost of a lower mean pension outcome but if both upside potential and risk are taken into account by means of certainty equivalent pensions (CEQ), the benefits are significant.

3.2. Average contributions for similar performance

To capture the differences between the individual DC-scheme and the generation insurance scheme in one number, we compute the level of contribution for the individual DC-plan that gives equal risk as the generation insurance scheme. We consider two different return-sharing parameters (κ) and perform a grid search on the parameter c for the individual DC-scheme that gives equal risk. We match on three risk measures, namely linearized loss aversion, the probability of shortfall, and power utility. Table 2 shows the results.

Panel A has the results for matching on linear loss aversion, i.e., for a given generation insurance scheme (rows 1 and 3) we show the statistics of the individual-DC schemes (rows 2 and 4) that have the same value of linear loss aversion. For 10% return shar-

---

[4] We model in total 145 generations per simulation, but it takes 45 years for the first generation to retire, which leaves 100 generations for which we obtain a pension outcome.

Table 1
Baseline simulation results.

This table shows the simulation outcomes of 1000 simulations of the 100 generations, in terms of the final pension as a fraction of the average lifetime wage. Participants enter the pension fund at age 23 and retire at 68. Demographics and wage profiles are as in Fig. 1. The annuity interest rate at retirement is equal to the short-term interest rate. Stock returns are i.i.d. and N(0.06,0.20), short-term interest rates are i.i.d and N(0.02, 0.01). Bond returns are generated from interest rates assuming a term premium of 1% per annum and a duration of 5. The fraction in stocks starts at 100% and declines by 10 percentage points annually after the age of 50, to zero. The first four columns show the parameter values used for the simulations, i.e., the contribution rate (c), the insurance fee (f), the fraction of excess returns being shared (κ) and the maximum fraction paid out of the insurance fund (F). The statistics are computed from the pension outcomes of all generations and all simulation runs. Each set of parameter uses the same risk asset returns. P(SF) is the probability of a pension shortfall. The columns LLA, QLA and PU give the certainty equivalent pension for linear loss aversion, quadratic loss aversion and power utility with risk aversion of 3, respectively.

Contr.	Ins. fee	Sharing	Max.payout	Mean	Median	Stdev	Skew	Kurt	p5	p10	LLA	QLA	P(SF)	PU
0.10	0	0	0	0.83	0.67	0.60	3.03	18.65	0.29	0.34	0.60	0.51	0.53	0.52
0.10	0	0.1	0.2	0.78	0.70	0.44	2.82	16.63	0.32	0.38	0.63	0.54	0.41	0.57
0.10	0	0.2	0.2	0.72	0.70	0.31	2.46	14.42	0.34	0.39	0.63	0.55	0.37	0.57
0.10	0	0.3	0.2	0.67	0.70	0.23	1.77	11.43	0.34	0.40	0.63	0.55	0.37	0.57
0.10	0	0.2	0.1	0.72	0.70	0.31	2.52	14.74	0.35	0.41	0.63	0.55	0.38	0.58
0.10	0	0.2	0.2	0.72	0.70	0.31	2.46	14.42	0.34	0.39	0.63	0.55	0.37	0.57
0.10	0	0.2	0.3	0.72	0.70	0.31	2.43	14.25	0.33	0.39	0.63	0.55	0.36	0.57
0.09	0.01	0	0.2	0.79	0.70	0.52	3.25	20.73	0.32	0.36	0.61	0.52	0.50	0.55
0.08	0.02	0	0.2	0.75	0.70	0.45	3.45	22.94	0.33	0.38	0.61	0.53	0.49	0.55
0.07	0.03	0	0.2	0.70	0.70	0.37	3.61	25.20	0.33	0.38	0.61	0.53	0.50	0.55
0.09	0.01	0.2	0.2	0.69	0.70	0.27	2.53	15.67	0.35	0.40	0.62	0.55	0.41	0.57
0.08	0.02	0.2	0.2	0.66	0.70	0.23	2.37	15.85	0.35	0.40	0.61	0.54	0.45	0.56
0.07	0.03	0.2	0.2	0.63	0.70	0.20	1.97	14.45	0.35	0.39	0.60	0.54	0.50	0.55

Table 2
Equivalent no-risk-sharing contributions.

For the same economic assumptions as in Table 1, this table gives a row-wise comparison of pension outcomes with and without risk sharing, for three values of the return-sharing parameter κ (the fraction of returns above the high-water mark being contributed to the insurance fund). Each combination of two lines matches a utility or risk measure of the outcomes so that the no-risk-sharing alternative is equivalent. Panel A matches pension outcomes on linear loss aversion. Panel B matches pension outcomes on the probability of shortfall. Panel C matches certainty equivalent pension under power utility with relative risk-aversion of 3. LLA is linear loss aversion, QLA is quadratic loss aversion, P(SF) is the probability of shortfall and PU is power utility.

Contr.	Ins. fee	Sharing	Max. payout	Mean	Median	Stdev	Skew	Kurt	p5	p10	LLA	QLA	P(SF)	PU
Panel A: Match on Linear Loss Averse utility														
0.100	0.00	0.10	0.20	0.78	0.70	0.44	2.82	16.63	0.32	0.38	0.626	0.54	0.41	0.57
0.108				0.90	0.72	0.65	3.04	18.65	0.32	0.37	0.627	0.53	0.48	0.57
0.100	0.00	0.20	0.20	0.72	0.70	0.31	2.46	14.42	0.34	0.39	0.629	0.55	0.37	0.57
0.109				0.91	0.73	0.66	3.04	18.65	0.32	0.37	0.630	0.53	0.47	0.57
Panel B: Match on shortfall probability														
0.100	0.00	0.10	0.20	0.78	0.70	0.44	2.82	16.63	0.32	0.38	0.63	0.54	0.406	0.57
0.120				1.00	0.80	0.72	3.04	18.65	0.35	0.41	0.66	0.55	0.405	0.63
0.100	0.00	0.20	0.20	0.72	0.70	0.31	2.46	14.42	0.34	0.39	0.63	0.55	0.369	0.57
0.127				1.06	0.85	0.77	3.04	18.65	0.37	0.43	0.67	0.57	0.367	0.67
Panel C: Match on Power utility														
0.100	0.00	0.10	0.20	0.78	0.70	0.44	2.82	16.63	0.32	0.38	0.63	0.54	0.41	0.567
0.108				0.92	0.73	0.66	3.04	18.65	0.32	0.38	0.63	0.53	0.48	0.567
0.100	0.00	0.20	0.20	0.72	0.70	0.31	2.46	14.42	0.34	0.39	0.63	0.55	0.37	0.574
0.109				0.95	0.76	0.69	3.04	18.65	0.33	0.39	0.64	0.53	0.47	0.572

ing, we need an individual DC contribution rate of 10.8% to have a similar utility. For 20% return sharing, the matching DC-scheme has a 10.9% contribution rate.

Panel B of Table 2 matches on the probability of shortfall, and finds that contributions need to increase with between 2 and 2.7%-points to achieve the same probability of shortfall. With the caveat that this performance measure does not take the extent of short-fall into account, generation insurance provides a large efficiency gain in terms of the probability of shortfall. The required increase in contribution is somewhat higher than for matching linear loss averse utility, and is explained by the fact that the linear loss averse utility function also takes into account the higher upside of individual-DC, whereas the probability of shortfall only considers downside risk.

Fig. 3. Performance improvement of early and late generations.
This figure shows the average differences in risk of the pension outcomes, relative to the DC scheme for the first and last 50 generations. Contribution is 10% and the insurance fee is 0. Per set of parameters (sharing percentage and maximum payout of the insurance fund), the relative improvement is shown in terms of the mean pension outcome, standard deviation, and the certainty equivalent of linear loss averse utility and power utility with a relative risk aversion coefficient of 3. The black columns have the difference compared to no-insurance (only individual accounts) for the first 50 generations. The white columns are for the last 50 generations. The results are based on 1000 simulations and parameters as for the baseline simulation results in Table 1.

Panel C has the matching on CEQ-pension for the power utility function and finds between a 0.8% to 0.9%-points contribution rate for DC in order to arrive at the same certainty equivalent. This is identical to the outcomes for linear loss aversion.

In all, the analysis in Table 2 shows an increase in contribution of the individual DC-scheme to achieve equal performance that is between 0.8%-point (same CEQ with linear loss aversion utility and power utility) and 2.7%-point (same probability of shortfall).

3.3. Shifts in intergenerational benefits

The performance of the generation insurance scheme relative to individual-DC could be biased because of a shift in intergenerational benefits. For example, generation insurance might make early generations worse off on average, and subsidize the later generations. If that is the case, the outcomes would need a better weighting over young and old generations.

To measure shifts in intergenerational benefits we compute the difference in performance of the generation insurance plan relative to the individual-DC scheme, separately for the first 50 generations and the last 50 generations. The results for three parameter configurations are in Fig. 3.

Fig. 3 uses black columns for the first 50 generations and white columns for the last 50 generations. Thus, each black/white pair of columns shows the relative improvement of the generation insurance scheme relative to DC and a large difference between the white and black column would show that the benefits do not spread out evenly over generations. The differences turn out to be small, and we conclude that it is clearly not the case that early generations empty the insurance fund, at the cost of later generations.

Note that in the case of a transition from a funded DB-type of pension system, the issue of late versus early generations could be addressed by starting with a non-empty insurance fund of a sufficient size. We discuss this together with other potential transition problems in Section 4.

3.4. Comparison of risk profiles

The generation insurance plan creates a risk profile of pension outcomes that is different from the individual-DC plan. To illustrate this, we use the investment mix parameter to match the outcomes of the two schemes based on the standard deviation of the pension outcome. The investment mix parameter is the age at which an initial 100% allocation to stocks declines with 10%-points per year.

We simulate pension outcomes for a range of ages at which the fraction in stocks starts declining,[5] and then choose the age of the individual (DC) plan that achieves equal standard deviation of the pension outcome to that of the generation insurance (CDC) plan. With the same standard deviation for the base case, we then compare average outcomes and performance metrics. Contribution rates are 10% in both plans. Generation insurance has 20% return sharing and a 20% maximum payout as a fraction of the insurance fund. The outcomes are in Table 3, for standard deviations of the pension outcomes between 0.20 and 0.40.

Table 3 shows that for the lowest standard deviation of pension outcomes (0.20), the generation insurance scheme has a 0.02 higher pension outcome and a 24%-points lower shortfall probability. Linear loss averse utility is higher by between 3%-points, in

---

[5] We keep the decline in fraction stocks to 10%-points per year.

Table 3

Comparison of risk profiles.

Difference in performance when the individual system and generation insurance are matched on having the same standard deviation of pension outcomes. Economic parameters for the simulation are as in Table 1. Results are based on 1000 simulations for 100 retiring generations. The decline in the percentage investment in stocks is 10%-points per year. The contribution rate is 10%, insurance fee is 0%, return sharing is 20% and maximum payout is 20%. The statistics are computed from the pension outcomes of all generations and all simulation runs. The age parameter for individual DC is chosen such that the standard deviation is closest to, but not exceeding the desired level. The age parameter for generation insurance is chosen such that the standard deviation is closest to, but not exceeding that of the individual DC scheme. The remaining difference in standard deviation between the two plans is in column four. The columns LLA, QLA and PU give the difference in certainty equivalent pension for linear loss aversion, quadratic loss aversion and power utility with relative risk aversion of 3, respectively. The difference in shortfall probability is in the column P(SF).

Stdev	Age parameter for generation insurance	Age parameter for individual DC	Difference in performance (Generation insurance minus individual-DC)					
			Stdev	Mean	LLA	QLA	P(SF)	PU
0.20	43	37	0.00	0.02	0.03	0.01	-0.24	0.02
0.25	47	39	0.01	0.05	0.05	0.03	-0.29	0.03
0.30	49	41	0.01	0.05	0.05	0.03	-0.28	0.03
0.35	52	43	0.02	0.06	0.06	0.03	-0.29	0.04
0.40	54	45	0.01	0.06	0.05	0.04	-0.29	0.04

Table 4

Volatility effects.

Summary statistics of pension outcomes for different values of stock volatility. For generation insurance, contribution is 10%, return sharing is 20% and maximum payout is 20%. Individual DC (panel B) has an annual contribution of 10.9%, which gives equal linear loss averse utility under a baseline volatility of 20%, see Table 2. Stdev is the standard deviation of pension outcomes. P(SF) is the probability of shortfall below 70% of average wage. The columns LLA, QLA and PU give the certainty equivalent pension for linear loss aversion, quadratic loss aversion and power utility with relative risk aversion of 3, respectively.

Volatility	Mean	Median	Stdev	LLA	QLA	P(SF)	PU
Panel A: Generation insurance							
0.15	0.72	0.70	0.21	0.66	0.59	0.29	0.64
0.20	0.72	0.70	0.31	0.63	0.55	0.37	0.57
0.25	0.72	0.70	0.43	0.60	0.51	0.44	0.50
Panel B: Individual DC							
0.15	0.91	0.80	0.46	0.67	0.58	0.38	0.68
0.20	0.91	0.73	0.66	0.63	0.53	0.47	0.57
0.25	0.91	0.65	0.90	0.59	0.48	0.55	0.48

terms of certainty equivalent pension. For power utility, the increase in certainty equivalent utility is 2%-point. For higher standard deviations, the extent of the differences increase accordingly. This illustrates how the generation insurance plan is not merely a risk-reduction plan, but an improvement in efficiency: at similar levels of risk (measured by standard deviation), the pension outcome is better.

3.5. Stock return volatility effects

The return-sharing scheme that we analyze works well for the given set of parameters and modeling of economic risks. However, economic parameters such as the volatility of stock returns are inherently uncertain, and the relative performance of one pension scheme to another could be sensitive to how the economy develops. Specifically, if future volatility is higher or lower than the historical average, the performance difference might be quite different, ex post. We assess the sensitivity of the two pension schemes for this effect by simulating with different volatility parameters. The contribution parameters for the simulation are set at the level where individual-DC and generation insurance have an equal average linear loss averse utility under a stock return volatility of 20%, as in Table 2. Then, we run the simulations with one lower and one higher volatility parameter. Table 4 shows the results.

A lower volatility than expected (15% instead of 20%) leads to lower pension risks in both schemes. The standard deviation of the pension outcome under generation insurance drops from 0.31 to 0.21. For the individual DC-scheme, it drops from 0.66 to 0.46. For

a higher than expected volatility (25% instead of 20%), the effect is exactly the opposite. Standard deviations increase from 0.31 to 0.43 and from 0.66 to 0.90, respectively. Moreover, the measures of certainty equivalent pension behave accordingly, with a slightly larger sensitivity for individual DC. In all, the generation insurance scheme has no excess sensitivity to ex post differences in volatility compared to the individual defined-contribution scheme.

4. Implementation issues

There are a number of practical and policy considerations that warrant discussion when considering the merits of a generation insurance scheme. We discuss the issues related to the purchase of a pension annuity, investment risk management and the transition from an existing pension plan.

4.1. The pension annuity

In practice, the pension annuity at retirement can be purchased from a commercial insurance company. However, it might be more beneficial for pensioners to have a separate non-profit entity providing the pension annuity. It can operate as a defined-benefit pension plan, and have lower costs than a commercial solution. Moreover, such an entity could provide the insurance against longevity risk. The portfolio is not necessarily risk-free, to allow for indexation ambitions to be pursued. Most of the assets would be hedged for interest-rate risk, to match the interest-rate sensitivity of the liabilities. The entity would retain some of the opaqueness of a DB-system, but with a scope limited to only retired participants. Also, having only retirees in the post-retirement plan limits the intergenerational conflicts.

Part-time employees, of employees with less than a working life of 45 years with the same pension fund, receive payouts in proportion to their capital accumulated in their individual account.

In practice, one might not want to have an insurance payout "surprise" at retirement, but rather a payout mechanism that is structured so that soon-to-retire participants can form proper expectations about the pension they will receive. For example, a prediction could be communicated about the pension level at retirement and the payout to be received, depending on the value of the individual account, the interest rate, and the assets in the insurance fund.

4.2. Risk management and portfolio choice

To facilitate a smooth conversion from a pension capital to a yearly pension, it would be beneficial to implement a form of risk management for the interest rate risk in the plan. This could be

implemented by locking-in annuity rates well before the retirement age, so that sudden changes in the pension outcome due to interest rate changes are mitigated.

Another form of risk management that is beneficial to the pension scheme is the choice of investment mix over the life cycle. Traditionally, the literature claims that young people should fully invest in risky assets, see for example Campbell and Viceira (2003). The share in risky assets declines over time, due to the diminishing share of labor income in their total discounted wealth. This is what we have used in the analyses above, and it is consistent with typical financial advice for personal financial planning. However, in contrast with that advice, a report by NEST Corporation (2010) finds that young people are actually quite sensitive to nominal losses in their pension savings and would rather not have risky assets when they are young. If these insights are true in general, one would want to adapt the asset mix in such a way that large stock investments are avoided initially, e.g., for the first five years. Another possibility is to waive the asset-related contribution to the generation insurance fund during the early working years.

4.3. Transitioning from existing pension plans

Transitioning from a DC-plan should be the easiest, as pension assets are clearly marked as such. Existing active (non-retired) participants continue to pay a contribution to their individual pension account, but by paying an insurance fee and/or the upside return sharing they obtain rights to the insurance payouts at retirement. Those rights would be proportional to the years in the new system. Given the assumption of nominal loss aversion we theoretically would prefer a setup without fixed insurance fees (i.e. $f = 0$ in our model) but to build up a reasonably sized generation insurance fund, we could imagine an implementation period in which the participants pay an insurance premium as well.

The transition from a defined-benefit (DB) system is little more complex. Active participants to a DB-system build up pension rights, not assets, and these rights have to be converted into assets. The assets then have to be split in an individual part and an insurance part. There are issues in the split in assets between (i) retired and non-retired participants, (ii) young and old participants, and (iii) the individual account and the insurance funds.

First, splitting the pension assets between retired and non-retired participants can be done on the basis of the actuarial value of the nominal pension rights. The allocation of the surplus or shortfall needs to be negotiated between the retired and active participants, and the sponsor. Such a negotiation is similar to the yearly negotiation that takes place implicitly in the DB-pension when contribution and indexation policies have to be determined. After the split, the assets of the retired generation are placed in a special payout-fund.

Second, the total pension rights of non-retired participants need to be distributed over the generations. A complication is the fact that the actuarial value of pension rights of the younger generation (say, below 40) is usually lower than the present value of the contributions paid into the fund. The reason is that under DB, pension rights accrue with a fixed percentage per year, while the actuarial value increases with age. For older participants, the yearly increase in pension rights is worth more than their contribution. This is an essential feature of a DB-fund and drives the intergenerational solidarity embedded in DB. A negotiation about the asset value of the pension rights for young and old generation will have to find the middle ground between the present value of the contributions and the actuarial value.

A final part of the transition is the allocation of assets between the individual account and the insurance fund. The desired initial size of the insurance fund determines the total fraction of assets that need to go in the insurance fund, and the contribution of each

generation is a function of the expected benefits of insurance. If the pension fund has a nominal funding ratio greater than 1, the surplus could be used to fill the insurance fund.

5. Conclusion

We analyze a risk-sharing scheme that captures benefits of intergenerational risk sharing, but is attractive for real-world participants by averting out-of-pocket transfers between generations. Moreover, the scheme retains the simplicity of an individual retirement savings scheme by having a generation insurance fund that takes in insurance contributions and a fraction of returns above a high-water mark. It pays out when a generation has a pension shortfall. Economic ownership of this fund lies with the fund's retirees.

The generation-insurance scheme provides protection in bad times which is appreciated by the efficiency measures that we evaluate. Unlike traditional DB-systems, the insurance is not financed by out-of-pocket expenses of active participants or employers but by investment returns of all individual participants in good times.

A key benefit of the behavioral design is that the rules governing the generation insurance fund can be explained in relatively simple terms. For example, it can be explained as a health insurance company that obtains insurance premiums, but only pays out to people who suffer from health problems and need treatment. Likewise, the generation insurance fund takes in insurance premiums and pays out if the investment returns or current interest rate lead to pension problems. Using a high-water mark for insurance contributions avoids nominal decreases in pension capital, while avoiding complex return-smoothing rules or direct transfers between generations.

The design has a positive externality in that the insurance capital is build up in good times and released in bad times. A pension scheme that embeds such a feature is worthwhile to pursue for the benefit of participants and public welfare. In fact, the buffer that defined-benefit pension funds should hold to withstand adverse market conditions is a reflection of the same principle. Returns in good years increase the buffer, so that shocks can be better absorbed in bad years.

A remaining question is to what extent the behavioral design comes close to the theoretical maximum for intergenerational benefits in a collective pension system. The answer to this question would give policy makers an idea of the extent of the trade-off between public support of the risk sharing mechanism and overall efficiency. We leave this as a topic for future research.

Acknowledgments

References

Abdellaoui, M., Bleichrodt, H., Kammoun, H., 2013. Do financial professionals behave according to prospect theory? An experimental study. Theory and Decision 74 (3), 411–429.
Andonov, A., Bauer, R., Cremers, M., 2016. Pension fund asset allocation and liability discount rates. Erasmus University Working paper.
Aslanidis, N., Christiansen, C., 2014. Quantiles of the realized stock–bond correlation and links to the macroeconomy. Journal of Empirical Finance 28, 321–331.
Baker, M., Pan, X., Wurgler, J., 2012. The effect of reference point prices on mergers and acquisitions. Journal of Financial Economics 106 (1), 49–71.
Barberis, N., Huang, M., Santos, T., 2001. Prospect theory and asset prices. The Quarterly Journal of Economics 116 (1), 1–53.

Barberis, N.C., 2013. Thirty years of prospect theory in economics: a review and assessment. The Journal of Economic Perspectives 27 (1), 173–195.

Beetsma, R.W.J., Bovenberg, A.L., 2009. Pensions and intergenerational risk-sharing in general equilibrium. Economica 76 (302), 364–386.

Beetsma, R.W.J., Romp, W.E., Vos, S.J., 2012. Voluntary participation and intergenerational risk sharing in a funded pension system. European Economic Review 56 (6), 1310–1324.

Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B., Meyer, S., 2012. Is unbiased financial advice to retail investors sufficient? Answers from a large field study. Review of financial studies hhr127.

Boender, C.G.E., Bovenberg, A.L., van Hoogdalem, S., Nijman, T.E., 2007. Optimal risk-sharing in private and collective pension contracts. Costs and benefits of collective pension systems. Springer, pp. 75–93.

Broadbent, J., Palumbo, M., Woodman, E., 2006. The shift from defined benefit to defined contribution pension plans - implications for asset allocation and risk management. Committee on the Global Financial System Working paper.

Brown, J.R., Weisbenner, S.J., 2014. Why do individuals choose defined contribution plans? Evidence from participants in a large public plan. Journal of Public Economics 116, 35–46.

Campbell, J.Y., Viceira, L.M., 2003. Strategic asset allocation: Portfolio choice for long-term investors. Oxford University Press.

Chiappori, P.-A., Paiella, M., 2011. Relative risk aversion is constant: evidence from panel data. Journal of the European Economic Association 9 (6), 1021–1052.

Coval, J.D., Shumway, T., 2005. Do behavioral biases affect prices. The Journal of Finance 60 (1), 1–34.

Cui, J., Jong, F.D., Ponds, E., 2011. Intergenerational risk sharing within funded pension schemes. Journal of Pension Economics and Finance 10 (01), 1–29.

Dalal, A.J., Arshanapalli, B.G., 1993. Estimating the demand for risky assets via the indirect expected utility function. Journal of Risk and Uncertainty 6 (3), 277–288.

Feng, L., Seasholes, M.S., 2005. Do investor sophistication and trading experience eliminate behavioral biases in financial markets. Review of Finance 9 (3), 305–351.

Frazzini, A., 2006. The disposition effect and underreaction to news. The Journal of Finance 61 (4), 2017–2046.

Genesove, D., Mayer, C., 2001. Loss aversion and seller behavior: Evidence from the housing market. The Quarterly Journal of Economics 116 (4), 1233–1260.

Goetzmann, W., Ingersoll, J., Spiegel, M., Welch, I., 2007. Portfolio performance manipulation and manipulation-proof performance measures. Review of Financial Studies 20 (5), 1503–1546.

Goetzmann, W.N., Ingersoll, J.E., Ross, S.A., 2003. High-water marks and hedge fund management contracts. The Journal of Finance 58 (4), 1685–1718.

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

Guidolin, M., Timmermann, A., 2005. Economic implications of bull and bear regimes in UK stock and bond returns*. The Economic Journal 115 (500), 111–143.

Guiso, L., Sapienza, P., Zingales, L., 2013. Time varying risk aversion. National Bureau of Economic Research.

Hoevenaars, R.P., Molenaar, R.D., Schotman, P.C., Steenkamp, T., 2008. Strategic asset allocation with liabilities: beyond stocks and bonds. Journal of Economic Dynamics and Control 32 (9), 2939–2970.

Kahneman, D., Knetsch, J.L., Thaler, R.H., 1991. Anomalies: the endowment effect, loss aversion, and status quo bias. The Journal of Economic Perspectives 5 (1), 193–206.

Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica: Journal of the Econometric Society 263–291.

Lucas, D., Zeldes, S.P., 2006. Valuing and hedging defined benefit pension obligations: The role of stocks revisited. Northwestern University and Columbia University Working paper.

NEST Corporation (2010). Improving consumer confidence in saving for retirement. Insight Report, available at: nestpensions.org.uk

Novy-Marx, R., Rauh, J.D., 2014. Linking benefits to investment performance in US public pension systems. Journal of Public Economics 116, 47–61.

Odean, T., 1998. Are investors reluctant to realize their losses. The Journal of finance 53 (5), 1775–1798.

Siegmann, A., Lucas, A., 2005. Discrete-time financial planning models under loss-averse preferences. Operations Research 53 (3), 403–414.

Thaler, R.H., Benartzi, S., 2004. Save more tomorrow™: using behavioral economics to increase employee saving. Journal of political Economy 112 (S1), S164–S187.

Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. Journal of Risk and uncertainty 5 (4), 297–323.

Van Rooij, M.C., Lusardi, A., Alessie, R.J., 2011. Financial literacy and retirement planning in the Netherlands. Journal of Economic Psychology 32 (4), 593–608.

Van Rooij, M.C., Lusardi, A., Alessie, R.J., 2012. Financial literacy, retirement planning and household wealth*. The Economic Journal 122 (560), 449–478.

Wakker, P.P., 2008. Explaining the characteristics of the power (CRRA) utility family. Health economics 17 (12), 1329–1344.

Weber, M., Camerer, C.F., 1998. The disposition effect in securities trading: an experimental analysis. Journal of Economic Behavior & Organization 33 (2), 167–184.