

Ed Westerhout, Jan Bonenkamp and Peter Broer

Collective Versus Individual Pension Schemes

A Welfare-Theoretical Perspective

Collective versus Individual Pension Schemes: a Welfare-Theoretical Perspective*

Ed Westerhout[†] Jan Bonenkamp[‡] Peter Broer[§]

October 24, 2014

Abstract

Collective pension contracts allow for intergenerational risk sharing with the unborn. They therefore imply a higher level of social welfare than individual accounts. Collective pension contracts also imply a sub-optimal allocation of consumption across time periods and states of nature however. Hence, collective pension contracts also reduce social welfare. This paper explores the welfare effects of a number of collective pension contracts, distinguishing between the two welfare effects. We find that collective schemes can be either superior or inferior to individual schemes.

1 Introduction

Pension reform is high on the political agenda. The population is ageing, the world economy undergoes a change of globalization and, at the firm level, pension liabilities increase relative to firm size. These are important factors that have led countries to discuss or implement alternatives for their pension schemes, of which many can be characterized as collective defined benefit pension schemes (Bonenkamp et al., 2014). One of the options is a greater reliance on individual retirement accounts. Examples of countries that have recently moved into

*A previous version of this paper has been presented at the Netspar/CPB/Bruegel Conference Pensions in Europe, December 6 2013 in Brussels. The authors thank the discussant of their paper, Ole Beier Sorensen, for useful comments. They further thank Casper van Ewijk, Albert van der Horst, Marcel Lever and Roel Mehlkopf for discussion and suggestions.

[†]Affiliations: UvA, CPB and Netspar.

[‡]Affiliations: CPB and Netspar.

[§]Affiliations: CPB and Netspar.

this direction are Hungary and the Slovak Republic. Another option is the introduction of individual defined contribution schemes. The US, the UK and Switzerland are examples of countries that have partly switched from collective defined-benefit (DB) arrangements to individual defined-contribution (DC) arrangements. As an illustration, the worldwide share of DC assets has grown in the last decade from 35% in 2000 to 44% in 2010 (Towers Watson, 2011). Other less radical options are reforms that maintain the collective nature of defined benefit schemes, but eliminate the pension guarantees that characterize them. The notional defined contribution scheme and the collective defined contribution schemes are examples of such schemes.

One very important element in the discussion about pension reform is how different schemes contribute to intergenerational risk sharing. This is far from trivial: pension schemes can differ a lot in the extent to which they relieve a fundamental distortion in the economy, namely the inability to trade with future, yet unborn generations. That intergenerational risk sharing can be welfare-increasing in general is well-known. Gordon and Varian (1988) and Ball and Mankiw (2007) are the classical papers. On the concrete assessment of these welfare gains for well-defined reforms, the literature is more scarce. Gollier (2008), Bucciol and Beetsma (2010) and Cui et al. (2011) are examples.

What is lacking in the literature is a focus on the contribution of pension reform to aggregate efficiency, *i.e.* welfare of the whole society. Some papers focus on the effects of pension reform upon a typical generation, for example, the generation that is born at the time of the reform or the group of future generations. However, it is dangerous to exclude some generations from the analysis as this may seriously bias results. Another approach is to focus on steady-state results. This may also give biased results as it excludes the effects during the transition to a new type of pension scheme (Sinn, 2000). A third approach adopts risk-neutral pricing to evaluate different pension schemes (CPB, 2012). This approach may be very useful to assess the effects of pension reform upon the balance between generations. For our purpose, the approach is not useful however since risk-neutral pricing by construction throws away the issue of efficiency and considers pension reform a zero-sum game (Cui et al., 2011).

This paper assesses the effects of a number of types of pension reform upon aggregate efficiency, which includes all generations: current generations, being the generations that are alive at the time of reform, and future generations, being those generations who are born thereafter. The approach is to calculate

for each generation the effects of the reform upon the expected value of his or her utility function, then convert the utility effect to an equivalent variation such as to bring the utility effects of different generations on an equal footing, and, finally, add up the equivalent variations of all generations (properly discounted if necessary). A positive sum of equivalent variations then indicates that the reform in question is a potential Pareto-improvement, *i.e.* the gains that accrue to the winning generations are larger than would be needed to compensate other generations for their losses.

Can we use our results on intergenerational risk sharing to say something about collective versus individual schemes? Here, we must be careful: collective pension schemes differ from individual schemes in many more aspects than intergenerational risk sharing alone. Collective schemes imply labour market distortions which may reduce considerably their welfare gains (Bonenkamp and Westerhout, 2014). Collective schemes are also generally larger than individual schemes, so that potentially economies of scale play a role. If participation in collective schemes is obligatory and that in private schemes voluntary, deviations from rational behavior can come into play: the participation in voluntary schemes may be sub-optimal due to myopia or self-control problems (Bodie and Prast, 2012). Competition among individual schemes may imply huge marketing and advertisement costs, but adhere more to the preferences of the participants in pension schemes (Bovenberg et al., 2007). Individual pension schemes may also exhibit more flexibility than collective schemes into shaping pension products to the preferences of heterogeneous participants. Furthermore, government policies with respect to the governance of individual schemes can be different from those that apply to collective schemes. In this paper, we deliberately abstract from all these factors in order to shed light exclusively on the aspect of intergenerational risk sharing. This allows us to highlight some important issues that until now have largely been overlooked.

A further caveat is that we focus exclusively on equity return risk. In reality, there are many more factors which are relevant for the financing of pension schemes. This implies that, although we think that equity return is the most important risk factor, our analysis is more of a first step.

We start on a fundamental level. We explore the optimal individual scheme and the optimal collective scheme that are identical, except for the fact that only the optimal collective scheme allows people to trade with the unborn. The optimal collective scheme serves as a first-best scheme. Hence, the welfare gain

it attaches to trade with unborn generations is the maximum gain that can be achieved by any collective pension scheme. Subsequently, we focus upon collective and individual schemes that resemble more the reality of pension economics. In particular, we describe collective pension schemes that do not exploit the scope for intergenerational risk sharing to its full extent, for example, by not including all future generations in the risk sharing scheme, by not including working generations in this scheme, and by applying uniform policies to different generations, even when more tailored policies would give better results. Similarly, we increase the realism of the model for individual schemes by restricting the variation over time in pension contributions, as in individual DC schemes.

The real-world collective schemes differ on three accounts with the real-world individual scheme. The first is risk sharing with the unborn. The second is that real-world collective schemes apply age-independent policies in setting contributions and pension benefits, whereas individual schemes apply age-dependent policies. Moreover, the contribution and indexation policies as applied by collective schemes are not matched with each other, as would be required by the principle of consumption smoothing. Thirdly, real-world collective pension schemes typically adopt investment strategies that are based on their financial wealth, which also conflicts with the principle of consumption smoothing. Combined, on account of risk sharing real-world collective pension schemes may be either superior or inferior to real-world individual schemes.

The structure of our paper is as follows. Section 2 constructs the optimal individual scheme which serves as a benchmark. It also sets out the methodology that will be used to compare different pension schemes and discusses the numerical implementation of the model. Section 3 explores the optimal collective (first-best) pension scheme. Sections 4 and 5 focus on more realistic pension schemes. Section 4 focuses on a number of alternative collective schemes, whereas section 5 discusses the individual DC scheme. Finally, section 6 closes the paper with some concluding remarks.

2 The modelling procedure

We want to compare different types of pension schemes in terms of economic and efficiency effects. A difficulty arises when one or more of the schemes feature PAYG elements. Then, a comparison that neglects transitional effects will produce biased results (Sinn, 2000). Therefore, we adopt the following

procedure. We specify a benchmark world in which each generation saves for its own pension through an optimal individual scheme. Pension reform then implies that, at a certain point in time, a new pension scheme takes over the optimal individual schemes. Hence, any transitional effects are fully taken into account. In addition, efficiency effects of different pension schemes can be compared as all schemes are introduced under the same initial conditions.

As to this benchmark world, we assume capital markets are perfect, the pension schemes act on behalf of rational optimizing agents and do not face any borrowing constraints. This specification of the benchmark world is transparent. Further, it allows us to express welfare effects in terms of household wealth (equivalent variations), which is not possible for any arbitrary scheme.

We start with a description of our benchmark world. In particular, we specify the utility function of agents, their intertemporal budget constraint, and the economic environment in which they operate. To stay close to the literature, we will make the same assumptions as the models in the seminal papers by Merton (1969) and Samuelson (1969).

2.1 The optimal individual pension scheme

Our model features overlapping generations of households. In particular, a new cohort is born at the beginning of each year, whereas the oldest cohort dies at the same time. We abstract from demographic issues and assume all cohorts to be equally large.

The model features one risk factor, which is the rate of return on equity. The rate of return on equity follows a time-invariant lognormal distribution: $\ln(1 + r_t^e) \sim N(1 + E(r), \sigma_r^2)$. Besides risky equity, the model features risk-free bonds. The risk-free interest rate, denoted r^f , is a constant. The life cycle falls into a working and a retirement phase. Labour productivity and labour supply during the working phase are non-stochastic and independent of age and time.

Households have finite lives. They enter the economy at the age of 20 and die at the age of 85. Next to this biological age, we distinguish the economic age of a household, which is 20 years younger, and employ this measure in our figures. Households enjoy utility over consumption throughout their life-cycle. The savings and investment decisions of households are delegated to a generation-specific pension fund. The consumption outcomes that follow from the pension fund's saving and investment strategies are then taken for granted

by households. This absence of private saving could be motivated by assuming that households are unable to operate on capital markets or are unwilling to do so, given that the pension scheme acts as a perfect agent.

Households have preferences that feature a constant value of relative risk aversion. Campbell and Viceira (2002) argue that this assumption allows meaningful predictions on the equity premium. Household welfare in year t of a household who is born in year j is then described by the expected value of the following utility function:

$$U_{t,j} = \sum_{i=t}^{j+T-1} (1+\delta)^{-(i-j)} \frac{c_{i,j}^{1-\gamma}}{1-\gamma} \quad t - T + 1 \leq j \leq t \quad (2.1)$$

Here, $U_{t,j}$ denotes intertemporal utility and $c_{t,j}$ consumption in year t of a household who is born in year j . Parameter δ measures the rate of time preference and T the length of the life cycle (65 years). The parameter $\gamma > 0$ measures both the aversion to risk, γ , and the aversion to intertemporal substitution, $1/\gamma$. The discounting back to the date of birth that is implicit in equation (2.1) is not relevant here, but will be relevant in the section on the social planner pension scheme.

A number of constraints apply in the optimization problem. First, the amount of financial wealth is given by the history of the household's savings. For the generation who has just entered the economy, this is zero: all households are born with zero financial wealth: $s_{j,j} = 0$.

Second, the household is born with a given amount of human wealth:

$$h_{t,j} = \sum_{i=t}^{j+T^E-1} \frac{w}{(1+r^f)^i} \quad t - T^E + 1 \leq j \leq t \quad (2.2)$$

Here, h denotes human wealth, w denotes the wage income that applies during the working phase of the life cycle and T^E denotes the length of this phase of the life cycle (45 years). Equation (2.2) thus defines human wealth as the sum of future wage income flows, discounted at the risk-free interest rate.

Using this definition of human wealth, we can write down an equation for the accumulation of total wealth, the sum of financial and human wealth:

$$(s_{t+1,j} + h_{t+1,j}) = (1+r^f)((s_{t,j} + h_{t,j}) - c_{t,j}) + (r_t^e - r^f)s_{t,j}^e \quad (2.3)$$

Here, $s_{t,j}^e$ measures the investment in equity at age j . Equation (2.3) reflects that we measure stock variables like financial wealth s_t and the investment in

equity s_t^e at the beginning of year t ; wage income and consumption transactions also occur at the beginning of the year.

The household's optimization problem is to maximize the expected value of intertemporal utility (equation (2.1)), subject to the initial levels of financial and human wealth and the accumulation equation for total wealth. Instruments of this optimization problem are the household's consumption and the investments in equity in all years of the life cycle.

The solution to the optimization problem consists of an investment equation, a consumption equation and an equation for indirect utility. Let us define first the equity portfolio share, a^e , as the investment in equity in terms of total household wealth (after consumption):

$$a_{t,j}^e = \frac{s_{t,j}^e}{(1+r^f)(s_{t,j} + h_{t,j} - c_{t,j})} \quad (2.4)$$

The first-order condition for this equity portfolio share reads as follows:

$$E_{t-1}[(1+a_{t,j}^e(r_t^e - r^f))^{-\gamma}(r_t^e - r^f)] = 0 \quad (2.5)$$

As both the equity rate of return and the interest rate are identically and independently distributed, it follows that the equity portfolio share is age-independent. Hence, the investment in equity is proportional with total household wealth.

A more intuitive expression for the equity portfolio share is the following:

$$a_{t,j}^e = \frac{E(r) - r^f}{\gamma\sigma_r^2} \quad (2.6)$$

This expression points to the determinants of the portfolio share in equity. In particular, it states that the equity portfolio share is proportional with the equity premium and inversely proportional with both the variance of equity returns and the coefficient of relative risk aversion. In discrete time, this expression is an approximation however; it will not be used in the simulation model.¹

Consumption is in the optimum also proportional with total household wealth:

$$c_{t,j} = \mu_{t,j}(s_{t,j} + h_{t,j}) \quad (2.7)$$

The propensity to consume out of total household wealth is a function of the risk-free interest rate, the rate of time preference, the portfolio allocation over

¹The expression holds exactly in continuous time. In discrete time, it is an approximation which is better the shorter is the unit period of the model. In addition, the approximation is fairly accurate if $\gamma \gg 1$ (Draper, 2008).

equity and bonds and the distribution properties of the rate of return on equity:

$$\mu_{t,j} = \left[\sum_{i=0}^{j+T-1-t} \left(\frac{\eta}{1+r^f} \right)^i \right]^{-1} \quad (2.8)$$

where

$$\eta = \left(\left(\frac{1+r^f}{1+\delta} \right) E \left[(1+(r^e-r^f)a^e)^{-\gamma} \right] \right)^{1/\gamma} \quad (2.9)$$

Lifetime utility, to be denoted as V , is defined as the value of expected utility at the optimum. The value function expresses lifetime utility as a function of total household wealth:

$$V_{t,j} = \max_{c,a^e} (E(U_{t,j})) = \frac{\mu_{t,j}^{-\gamma} (s_{t,j} + h_{t,j})^{1-\gamma}}{1-\gamma} \quad (2.10)$$

Hidden in this formulation is the Sharpe ratio $(E(r^e) - r^f)/\sigma_r$, the mean excess return on equity in terms of its risk. It can be recovered by using equation (2.8) to write μ in terms of η , equation (2.9) to write η in terms of a^e and equation(2.6) to write a^e in terms of the Sharpe ratio. This yields that, as one would expect, utility is increasing in the Sharpe ratio.

2.2 Assessing the effects of a pension reform

Upon its introduction, the pension fund is endowed with the accumulated savings of all cohorts alive at that time:

$$s_0 = \underline{s} \equiv Q_0 \left(\sum_{j=1-T}^0 s_{0,j}^I \right) \quad (2.11)$$

Equation (2.11) displays an initial condition; the index 0 refers to the time of introduction of the pension scheme and the superscript I is attached to indicate that the variables refer to the benchmark scheme, *i.e.* the optimal individual scheme. Q_0 is an operator that picks out one of the many possible initial states for individual financial wealth. We assume that for all cohorts, initial financial wealth is equal to the mean of the corresponding distribution. Hence, we will take Q_0 equal to E_0 .

The initial amount of human wealth can be derived from the general expression for human wealth:

$$h_0 = h \equiv \sum_{i=0}^{\infty} \frac{T^E w}{(1+r^f)^i} \quad (2.12)$$

Total wealth of the pension fund accumulates over time similar to that of the pension scheme in the benchmark model.

$$(s_{t+1} + h_{t+1}) = (1 + r^f)((s_t + h_t) - c_t) + (r_t^e - r^f)s_t^e \quad (2.13)$$

The aggregate investment in equity relates to the equity portfolio share, a^e , as in equation (2.4):

$$s_t^e = a_t^e(1 + r^f)(s_t + h_t - c_t) \quad (2.14)$$

Households are required to make contributions to the pension scheme in the working phase of their life, whereas in the retirement phase they receive benefits from the scheme. Throughout the paper, we will assume that private capital markets are lacking. Therefore, actions undertaken by the pension scheme cannot be undone by participants. To the extent that the pension scheme behaves according to the preferences of its participants which we assume to be the case, this assumption is quite innocuous. In the Netherlands, the market for supplementary pension schemes are quite generous and private savings are small compared to the collective savings done by these pension schemes.

Given the absence of private savings, the contributions and benefits chosen by the pension scheme pin down consumption in all phases of the individual life cycle.

We distinguish welfare effects both for different generations and for society as a whole (the pension scheme). The latter aggregates the certainty equivalents of different generations. Basically, this is the approach adopted by Auerbach and Kotlikoff (1987), the only difference being that this paper applies this approach in a stochastic environment.

We mainly focus on the efficiency effect of policies, *i.e.* the welfare effect for society as a whole. More precisely, we are interested in the question whether a typical pension contract establishes a welfare gain or a welfare loss. One could assume that the pension scheme uses lump-sum transfers to reallocate the societal welfare gain over different generations, for example to ensure that all cohorts share in this welfare gain or to achieve that all share equally. We leave this issue aside, focussing on the contribution of pension reform to aggregate efficiency.

We base welfare effects upon the value of intertemporal utility, as expected as of the time of introduction of the pension scheme. It is convenient to distinguish

between current and future generations:

$$E_0(\tilde{U}_{0,j}) = \sum_{i=0}^{j+T-1} (1+\delta)^{-(i-j)} E_0 \left[\frac{\tilde{c}_{i,j}^{1-\gamma}}{1-\gamma} \right] \quad 1-T \leq j \leq 0 \quad (2.15)$$

$$E_0(\tilde{U}_{j,j}) = \sum_{i=j}^{j+T-1} (1+\delta)^{-(i-j)} E_0 \left[\frac{\tilde{c}_{i,j}^{1-\gamma}}{1-\gamma} \right] \quad j > 0 \quad (2.16)$$

The welfare effect for a current generation is measured at the time of introduction of the scheme over the remaining life of the generation. The welfare effect for a future generation is measured at the time of birth of this generation over his whole life.

The welfare effect of a pension reform for a typical current generation j can be determined by comparing $E_0(\tilde{U}_{0,j})$ with $E_0(U_{0,j})$. For future generations, the comparison is between $E_0(\tilde{U}_{j,j})$ with $E_0(U_{j,j})$. In order to be able to compare the welfare effects of different generations, we convert them into equivalent variations.²

The equivalent variation of a typical cohort is defined as the amount of wealth that a household of this cohort should be given in the benchmark case (relative to initial wealth) in order to obtain the level of utility that will be achieved by participating in the pension scheme. For current generations, this equivalent variation is measured at the time of introduction of the scheme, for future generations at their time of birth. The following equations solve implicitly for this equivalent variation, EQV , for current and future generations:

$$E_0(\tilde{U}_{0,j}) = \frac{\mu_{0,j}^{-\gamma} ((s_{0,j} + h_{0,j})(1 + EQV_{0,j}))^{1-\gamma}}{1-\gamma} \quad 1-T \leq j \leq 0 \quad (2.17)$$

$$E_0(\tilde{U}_{j,j}) = \frac{\mu_{j,j}^{-\gamma} (h_{j,j}(1 + EQV_{j,j}))^{1-\gamma}}{1-\gamma} \quad j > 0 \quad (2.18)$$

Equations (2.17) and (2.18) use the indirect utility expression for utility in the benchmark case, since this relates utility to the amount of wealth in terms of the current consumption good or the consumption good at the time of birth of a future generation.

Having derived expressions for the equivalent variations of a pension reform for all current and future generations, the next step is to sum (after discounting) all the equivalent variations in order to arrive at the aggregate equivalent

²The counterpart of the equivalent variation is the compensating variation. The two may differ substantially if the arguments of the welfare function enter this function in a non-separable way. In our case, welfare is a function of consumption only, so we expect the two are quite similar.

variation, a measure of potential Pareto-efficiency. If the aggregate equivalent variation is positive, the gains of the winners of the reform are sufficiently large to *potentially* more than compensate the losers of the reform. The reform can then be said to be potentially Pareto-improving.

The literature discusses two approaches, the ex ante approach and the interim approach. The two differ with respect to the question whether some agent born in a given period but in different states of the world is the same agent. According to the concept of ex ante efficiency, he is. According to the concept of interim efficiency, he is not. Indeed, according to the interim efficiency concept, agents who are born in a given period but in different states of the world are treated as different agents.³

Here, we adopt the ex ante approach. The approach is intuitively appealing and also adopted by most of the literature. Under ex ante efficiency, the efficiency measure sums the equivalent variations of current generations and the expected values of equivalent variations of future generations. As these equivalent variations are non-stochastic, the risk-free interest rate acts as discount rate.

$$AEQV_0 = \frac{\sum_{j=1-T}^0 (s_{0,j} + h_{0,j})EQV_{0,j} + \sum_{j=1}^{\infty} (1 + r^f)^{-j} h_{j,j}EQV_{j,j}}{\sum_{j=1-T}^0 (s_{0,j} + h_{0,j}) + \sum_{j=1}^{\infty} (1 + r^f)^{-j} h_{j,j}} \quad (2.19)$$

2.3 Numerical simulations

Our calculations are based upon a particular parameter configuration. The parameter values used can be found in Table 1. For parameters that characterize the economic environment, values are chosen that are more or less realistic. Other parameters refer to a specific real-world contract and are chosen accordingly to mimic this contract. Only the value for the time preference rate δ follows from a calibration of the model. More precisely, we chose the value of δ that produces a flat profile for median consumption in the optimal individual scheme. Obviously, this value of δ will subsequently be applied to all pension schemes.

[Table 1 about here.]

³See Rangel and Zeckhauser (2001) for a clear exposition of the two concepts

For this particular choice of parameters, we show the outcomes for consumption and financial wealth in Figures 1 and 2 respectively. Figure 1 illustrates consumption smoothing over the life cycle; retirement does not affect consumption. The spread of the distribution of consumption, depicted by the 5 and 95 percent quantiles, increases over the life cycle, reflecting that households invest a constant fraction of their total wealth into equity, even after retirement.

Unlike consumption, financial wealth does have a kink at the retirement date. Only at that date, the variability of financial wealth, which is fairly constant over the life cycle, starts falling.

[Figure 1 about here.]

[Figure 2 about here.]

3 The first-best pension scheme

The benchmark scenario in which each cohort saves for its own pension through the vehicle of an individual scheme is inefficient. In particular, the scheme foregoes the gains from trading with the unborn. This inefficiency could be removed if a household could trade before it was born. This is less odd than it may seem at first sight. A pension fund can set an account for households yet to be born and trade on financial markets on behalf of these households. Yet, we consider the analysis of the optimal collective fund in the present section more like a theoretical exercise, aimed at providing an assessment of the maximum gain due to risk sharing with the unborn. Teulings and de Vries (2006) have performed a similar exercise in continuous time and arrive at a similar expression.

The maximization problem is thus similar to that in the previous section, except for the timing. Under the optimal individual scheme, the household starts to maximize at its birth. However, the first-best scheme allows the household to optimize long before its birth, *i.e.* at the time of introduction of the first-best scheme, year 0. Hence, for current generations the optimization problem is completely identical to the one studied above and we will not repeat it for convenience. For future generations, the utility function is based on consumption flows over the whole life cycle:

$$U_{0,j} = \sum_{i=j}^{j+T-1} (1 + \delta)^{-(i-j)} \frac{c_{i,j}^{1-\gamma}}{1-\gamma} \quad j > 0 \quad (3.1)$$

The pension fund creates an account for these generations at the start of time. Hence, financial wealth is zero not at these generations' years of birth j , but in year 0: $s_{0,j} = 0$. Human wealth is defined over the full life cycle:

$$h_{0,j} = \sum_{i=j}^{j+T^E-1} \frac{w}{(1+r^f)^i} \quad (3.2)$$

Instruments of this first-best maximization problem are the portfolio allocation from the start of time, $a_{i,j}^e$ $i = 0, \dots, j+T-1$, and consumption over the life cycle of the cohort, $c_{i,j}$ $i = j, \dots, j+T-1$.

The accumulation equation for the future generations now consists of two parts, one before and one after the cohort's birth:

$$\begin{aligned} (s_{t+1,j} + h_{t+1,j}) &= (1+r^f)((s_{t,j} + h_{t,j})) + (r_t^e - r^f)s_{t,j} \quad 0 < j \quad (3.3) \\ (s_{t+1,j} + h_{t+1,j}) &= (1+r^f)((s_{t,j} + h_{t,j}) - c_{t,j}) + (r_t^e - r^f)s_{t,j} \\ j \leq t \leq j+T-1 \end{aligned}$$

The solution to this optimization problem has the following features. The expressions for consumption and portfolio allocation over the individual's life cycle are similar to those presented for the optimal individual model:

$$\begin{aligned} c_{t,j} &= 0 & t < j \\ c_{t,j} &= \mu_{t,j}(s_{t,j} + h_{t,j}) & j \leq t \leq j+T-1 \end{aligned} \quad (3.4)$$

$$E_{t-1}[(1 + a_{t,j}^e(r_t^e - r^f))^{-\gamma}(r_t^e - r^f)] = 0 \quad t \leq j+T-1 \quad (3.5)$$

$s_{j,j} + h_{j,j}$ is now a stochastic variable, of which the value is determined by the shocks and the portfolio allocation chosen in the pre-labour market entry period. During this pre-labour market entry period, consumption is zero and the portfolio allocation in terms of total wealth identical to the allocation after labour market entry (equation (3.5)).

Combining things, we derive the following expression for indirect utility of an individual whose labour market entry is N years after the pension reform in year 0:⁴

$$E_0 V_{N,N} = \frac{\mu_{N,N}^{-\gamma} h_{N,N}^{1-\gamma}}{1-\gamma} \exp\left\{ \frac{(1-\gamma^2)}{2\gamma^2} \left(\frac{E(r^e) - r^f}{\sigma_r} \right)^2 N \right\} \quad (3.6)$$

Comparing the expression in equation (3.6) for the first-best scheme with its counterpart for the optimal individual scheme (equation (2.10)) shows that the

⁴We make an approximation here, which is that $\log(1+x)$ equals x , where x is a short-hand expression for $(r_t^e - r^f)(E(r_t^e - r^f))/(\gamma\sigma_r^2)$

first-best scheme gives higher welfare. Furthermore, the gain is increasing in N , which is intuitive, as increasing N adds possibilities to trade with other generations. In addition, the expression shows that the gain is increasing in the square of the Sharpe ratio, like in Teulings and de Vries (2006) and reminiscent of its role in determining utility in the optimal individual pension scheme.

The expression for the corresponding equivalent variation can be derived easily from equation (3.6), using the definition of the equivalent variation for future generations (equation (2.18)):

$$EQV_{N,N} = exp\left\{\frac{(1+\gamma)}{2\gamma^2}\left(\frac{E(r^e) - r^f}{\sigma_r}\right)^2 N\right\} - 1 \quad (3.7)$$

The expression for the equivalent variation can be used to derive the effect of γ . It follows that a higher risk aversion reduces the welfare gain from trading with the unborn. This is as expected as risk aversion and the Sharpe ratio play opposite roles in the portfolio allocation between equity and bonds (see the expression below equation (2.5)).

[Figure 3 about here.]

[Figure 4 about here.]

Figure 3 presents the equivalent variations for current and future generations for the first-best scheme. For all current generations, these equivalent variations are zero. For these generations, the possibility to trade already exists so the first-best scheme has no additional value. For future generations, the possibility to trade before labour market entry is positive and the more positive, the later the generation is born. The corresponding equivalent variations are dated at the time of labour market entry of future generations as in equation (3.7). In order to calculate the aggregate equivalent variation, the equivalent variations for future generations need to be discounted to the time of pension reform, $t = 0$. If we would plot the curve of discounted equivalent variations, this curve would bend towards zero for future generations, implying that the sum of discounted equivalent variations converges. Our estimate of the sum of discounted equivalent variations, the aggregate equivalent variation, is quite large: 14.4 percent of initial wealth.

Figure 4 displays the 5%, 50% and 95% quantiles of the distribution of consumption of the cohort that is born at the time of pension reform, *i.e.* $t = 0$ and that enters the labour market at time $t = 20$. For comparison, the

corresponding quantiles for the optimal individual scheme are also included. The two distributions differ in two respects. The first is the volatility of consumption. Under the optimal individual scheme, consumption is riskless at the start of the life cycle and becomes gradually riskier over time. Under the first-best scheme, consumption is stochastic already in the beginning of the life cycle. This is due to the risk sharing across generations, which characterizes the first-best scheme.

The same risk sharing induces the first-best scheme to invest more in risky equity in order to exploit the equity premium. The difference in demand for equity is due to the fact that the first-best scheme features a higher base for risk taking. As under the first-best scheme future generations absorb part of today's shocks, it is optimal for the scheme to invest more in equity. As a consequence, the first-best scheme earns a higher average rate of return. The effect of this is also visible in Figure 4: median consumption is higher under the first-best scheme.

4 Real-world collective contracts

The first-best contract is a useful study object. It specifies the maximum of welfare gains that can be obtained through intergenerational risk sharing. Contracts that can be found in the real world differ strongly from this first-best contract, however. Real-world contracts are often constrained by supervisory policies to not let financial assets become too large or too small, relative to pension liabilities. Furthermore, investment policies are typically specified in terms of financial wealth rather than total wealth. It thus remains to be seen how large the benefits of intergenerational risk sharing in a real-world contract will be.

4.1 The basic real-world collective contract

We now discuss a real-world pension contract. This contract more or less matches the contract that is representative of Dutch second-pillar pension contracts. It differs from the first-best contract described above in four (related) aspects: the setting of pension contributions, the setting of pension benefits, the portfolio allocation and the degree of intergenerational risk sharing.

What is similar to the pension contracts specified before is the pension fund's initial financial and human wealth. The pension fund is endowed with the

accumulated savings of all current cohorts and the human capital of all current and future cohorts:

$$s_0 = \underline{s} \equiv Q_0 \left(\sum_{j=1-T}^0 s_{0,j}^I \right) \quad (4.1)$$

$$h_0 = h \equiv \sum_{i=0}^{\infty} \frac{T^E w}{(1+r^f)^i} \quad (4.2)$$

Also similar is the accumulation over time of total wealth of the pension fund:

$$(s_{t+1} + h_{t+1}) = (1+r^f)((s_t + h_t) - c_t) + (r_t^e - r^f)s_t^e \quad (4.3)$$

In the following, we will work with an accumulation equation for financial wealth only. Using $h_t = h = ((1+r^f)/r^f)T^E w$ and $c_t = T^E(w - pp_t w) + \sum_{i=t-T+1}^{t-T^E+1} a_{t,i}$, the following expression can be derived to be equivalent to equation (4.3):

$$s_{t+1} = (1+r^f)(s_t + T^E pp_t w - \sum_{i=t-T+1}^{t-T^E+1} a_{t,i}^p) + (r_t^e - r^f)s_t^e \quad (4.4)$$

Here, a^p refers to the pensions that are paid out to retirees. We will discuss it in more detail below.

The pension contribution rate in the real-world pension contract consists of two terms: a base contribution rate and an additional rate (positive or negative) that relates to the pension fund's financial position. Let us first discuss the base contribution rate. This relates to the pension ambition, which we take to be a given fraction of the contemporaneous wage rate:

$$\bar{a}^P = \psi w \quad (4.5)$$

The base contribution rate, $\bar{p}p$, is uniform over the life cycle. It is determined by the condition that annual contributions equal the increase in pension rights of those in the workforce. Using that the contribution rate, wage rate and pension ambition are the same for all ages, the equation for the base contribution rate reads as follows:

$$\bar{p}p = \left[\left(\sum_{j=0}^{T^E-1} \sum_{i=j}^{T-T^E-1+j} (1+r^f)^{-i} \right) acc \right] / (T^E w) \quad (4.6)$$

where acc denotes the annual build-up of pension rights; it will be defined below.

The additional contribution rate relates linearly to the difference between the actual value of the pension fund's funding ratio fr and some target value.

The pension fund's funding ratio is defined as the ratio of the pension fund's financial wealth and the pension fund's liabilities (to be defined below):

$$fr_t = s_t/L_t \quad (4.7)$$

As target value, we choose 100%. Hence, the total contribution rate reads as

$$pp_t = \bar{p}p - \lambda_{pp}(fr_t - 1) \quad (4.8)$$

Equation (4.8) illustrates that the pension contribution rate equals the base contribution rate only if the funding ratio of the pension fund equals its target value of 100%. For higher (lower) levels of the funding ratio, the contribution rate is lower (higher). Deviations of the funding ratio from its target value of 100% have another implication, namely that pension benefits may exceed or remain below the pension ambition that drives the base contribution rate. This can be seen as follows.

Participants accumulate pension rights: each year of work earns them an accrual equal to acc which is based on the pension ambition \bar{a}^p . Hence, given that each household works T^E years, the accrual rate can be derived to equal

$$acc = \bar{a}^p/T^E = \psi w/T^E \quad (4.9)$$

However, pension rights may be increased above or decreased below the level that equals the sum of accruals on account of a good or bad financial position of the pension scheme. In particular, the actual evolution of pension rights is described by the following equations:

$$a_{t,i}^p = 0 \quad (4.10)$$

$$a_{t,i}^p = (1 + \pi_{t-1}^a)a_{t-1,i}^p + acc \quad i + 1 \leq t \leq i + T^E \quad (4.11)$$

$$a_{t,i}^p = (1 + \pi_{t-1}^a)a_{t-1,i}^p \quad i + T^E + 1 \leq t \leq i + T \quad (4.12)$$

Here, $a_{t,i}^p$ denotes pension rights at the beginning of year t , accumulated by the generation born in year i and π^a denotes the rate of indexation. This indexation relates to the financial position of the fund, much in the same way as the pension contribution rate:

$$\pi_t^a = \lambda_\pi(fr_t - 1) \quad (4.13)$$

The pension fund's liabilities sum the rights of the generations that are working or retired at that time. They are defined exclusive of a possible future indexation

of pensions:

$$L_t = \sum_{i=1}^{T^E} \sum_{j=t+i}^{t+i+T-T^E-1} (1+r^f)^{j-t} a_{t,i}^p + \sum_{i=-(T-T^E-1)}^0 \sum_{j=t}^{i+T} (1+r^f)^{j-t} a_{t,i}^p \quad (4.14)$$

It is little odd to speak of future pensions in terms of rights (from the participant's point of view) or liabilities (from the fund's point of view). These pensions are truly stochastic variables, indeed. In the Dutch case, it is common to use the words rights and liabilities however. Related, it is actual policy in the Netherlands to discount these future cash flows with a risk-free interest rate, whereas an appropriate discounting would use a stochastic discount factor. We follow the tradition in order to be able to apply our results to the Dutch case.

Before concluding our description of the real-world pension contract, we have to take two more steps. First is the assumption on investment policies. We take the equity investment of the pension fund as a constant fraction of its financial wealth. Note that this will in general not be optimal, since optimality requires, as seen above, that the equity investment is proportional to total wealth, rather than financial wealth. Moreover, the fraction a^e need not be optimally chosen.

$$s_t^e = a^e (s_t + T^E w - c_t) \quad (4.15)$$

Second, we need to specify the pension fund's policies regarding transition generations. We assume the following. The pension ambition of the generations that are retired at the time of introduction of the real-world pension scheme is such that at the end of their life, their financial wealth will be fully exhausted:

$$s_{0,j}^I = \sum_{i=0}^{j+T-1} (1+r^f)^{-i} a_{0,j}^p \quad 1-T \leq j \leq -T^E \quad (4.16)$$

For the generations that are working at the time of introduction of the scheme a similar assumption is made:

$$s_{0,j}^I = \sum_{i=j+T^E}^{j+T-1} (1+r^f)^{-i} a_{0,j}^p \quad 1-T^E \leq j \leq 0 \quad (4.17)$$

This completes our description of the real-world contract. Consumption of the working generations and that of retired generations relates to pension contribution and benefits as before. Note that now pension contributions are defined as the product of a contribution rate and some contribution base:

$$\begin{aligned} c_{t,i} &= w - pp_t w & i \leq t \leq T^E - 1 \\ c_{t,i} &= a_{t,i} & T^E \leq t \leq T - 1 \end{aligned} \quad (4.18)$$

Our benchmark real-world collective contract assumes $\lambda_{pp} = 0$. Hence, the contribution rate is fixed (see equation (4.8)). This scheme is representative of Dutch schemes which aim at absorbing shocks entirely by shifting the indexation of pensions.

How does this scheme perform? Some insight is given in Figure 5, which sketches the distribution of the funding ratio of the pension fund over time. The figure shows that the model converges quite quickly to a steady state, in which the funding ratio takes a value in between about 75 percent and 175 percent with 90 percent probability. Correspondingly, Figure 6 displays the development of the distribution of the indexation rate over time. The 5 percent quantile is about -2.0 percent, the 95 percent quantile about 7.5 percent. Hence, there is about 5 percent probability that an indexation cut of more than 2 percent occurs and, equally, a 5 percent probability that indexation exceeds 7.5 percent.

[Figure 5 about here.]

[Figure 6 about here.]

We depict a graph of the equivalent variations of the real-world contract in Figure 7. The figure shows that the real-world contract performs the same as the first-best contract for retired generations, but offers lower equivalent variations for working-age generations. Unborn generations do slightly better than working-age generations, but are also substantially worse off in the real-world scheme than in the first-best scheme or the individual benchmark scheme. The equivalent variation is -8.4 percent of initial wealth. Hence, the introduction of this real-world collective scheme is equivalent with a permanent reduction of consumption in the optimal individual scheme with 8.4 percent.

[Figure 7 about here.]

Figure 8 gives the consumption over the life cycle of the generation who is born at the time of pension reform, for the benchmark pension scheme and for the real-world collective scheme. It shows that consumption in the latter scheme is completely flat during the working phase, reflecting the fixed contribution policies of the collective scheme. The distribution of consumption during the retirement phase shows the opposite. This distribution is wider for the collective scheme, due to the fact that this scheme does not absorb any of the shocks that occur during the active phase of the life cycle. Median consumption is lower

for the collective scheme than for the benchmark scheme during most of the life cycle. This reflects that the collective scheme, studied in this section, is dominated by the optimal individual scheme.

[Figure 8 about here.]

4.2 Decomposition of the welfare loss

In general, the real-world collective contract and our benchmark optimal individual scheme differ in a number of respects. The collective contract allows for risk sharing with the unborn, but implies, as discussed above, a suboptimal allocation of wealth over the life cycle. Furthermore, the investment behaviour of the collective contract is suboptimal since it takes financial wealth rather than total wealth as its base. In addition, contributions and the accrual of pension entitlements under the collective contract adheres to the uniformity principle. This principle is known to be welfare-reducing, as it implies an implicit tax upon young workers and an implicit subsidy to old workers.

The consequence of this is that we have no good answer to the question what is the value of intergenerational risk sharing, when isolated from other factors? An answer to this question cannot be found by simply eliminating all other differences between the real-world collective contract and the optimal individual scheme. The reason is that these other factors have also implications for the amount of risk sharing. By eliminating these other factors, the amount of risk sharing is automatically increased or reduced. What we can do is the following. We introduce a real-world collective scheme that is in all respects identical to the one studied before, but with one exception: generations born at the time of introduction of the scheme or thereafter, do not take part in the risk sharing of shocks. Shocks are thus absorbed by the then living generations only. The measure for the aggregate equivalent variation that this simulation produces can in turn be subtracted from the aggregate equivalent measure that we calculated for the original simulation for the real-world collective scheme in order to get an estimate of the welfare effect of risk sharing with future generations.

The outcome of this additional simulation is an aggregate equivalent variation of -9.0 percent of total wealth. Subtracting this outcome from the aggregate equivalent variation for the real-world collective scheme, -8.4 percent of total wealth, yields an estimate for the gain from risk sharing of 0.6 percent of total

household wealth.

4.3 Alternative real-world collective contracts

The real-world collective contract that we have discussed above features two policy ladders: one that links the pension contribution rate to the funding ratio of the pension scheme and another one that links the rate of indexation to this funding ratio. An interesting question is what is the contribution of these two types of policy to the calculated welfare effect of the real-world collective scheme?

To answer this question, we run two additional simulations. The first one takes the real-world collective pension scheme from the previous section, but changes the value of λ_{pp} , the parameter that links the pension contribution rate to the funding ratio. Rather than 0.0, we now let this parameter take value 0.1. The second simulation is based upon the previous one, but now changes the value of the other policy parameter: the value of λ_{π} is increased from 0.1 to 0.2.

[Table 2 about here.]

Table 2 shows that the first alternative features a lower welfare loss than the benchmark real-world collective contract: 6.1 rather than 8.4. Letting the pension contribution rate respond to shocks in the funding ratio improves consumption smoothing over the life cycle. The second alternative gives an even better outcome. Halving the reaction of the indexation parameter to shocks in the funding ratio implies more risk sharing with future generations. This allows for a further reduction of the welfare loss with 2.2 percentage points to 3.9 percent of total wealth.

The two alternatives share with the benchmark real-world collective contract that they feature a lower level of social welfare than the optimal individual scheme. The question arises whether this outcome could be improved. In particular, if we would maintain the institutional setup of the real-world collective contract (the linear form of the two policy ladders, the uniformity principle and the investment strategy), but optimize over its various parameters, *i.e.* λ_{π} , λ_{pp} , μ and acc , how large would then be the welfare loss of the real-world collective scheme? Table 2 shows that in that case there is not a welfare loss, but a welfare gain. The gain is substantial: 4.4 percent of initial wealth. How is this

outcome obtained? Much more risk sharing with future generations (λ_π reduces to 0.005), a more flexible contribution rate (λ_{pp} increases to 0.17), more aggressive investment policies (μ increases to 0.7) and (compared to the benchmark real-world collective contract) a more generous pension scheme (acc increases from 1.2 to 1.5 percent a year).

These simulations make clear that a collective scheme can produce better results than an individual scheme. However, this implies a substantial deviation from current schemes as represented by what we have called the benchmark real-world collective scheme. However, there is a serious caveat. That is that it may be considered unfair to compare real-world collective schemes, with institutional features that have a political or juridical background, with the ideal form of individual scheme. A more fair comparison would be one with a real-world individual scheme rather than the optimal individual scheme.

Therefore, we develop an alternative individual scheme in the next section and compare this scheme with the optimal individual scheme and the real-world collective schemes discussed above.

5 An individual DC scheme

This section develops a more realistic type of individual scheme. In particular, the individual scheme that we will discuss now differs in two respects from the optimal individual scheme. One is the pension contribution rate. This takes a fixed value and does not respond to any shocks in the pension fund's rate of return. Second, the scheme adopts a strategy of life cycle investment that can be observed in many individual DC schemes in the real world. Concretely, the pension scheme invests all of its financial wealth in equity when households enter the pension scheme and reduces the ratio of equity to financial wealth linearly to 21.4 percent at the time of retirement. The equity-financial wealth ratio stays at the value of 21.4 percent during the retirement phase. The value coincides with that of the optimal individual scheme at the time of retirement. The value of the contribution rate is chosen such as to smooth consumption over the life cycle.

What are the results? Figure 7 includes the equivalent variations of a change from the optimal individual contract to the individual DC contract that we introduced here. The equivalent variations are negative everywhere, indicating that the suboptimal investment policy and the fixed contribution policy of the

modified individual scheme are welfare-reducing. The figure also shows that the equivalent variations of the modified individual contract are less negative than those of the benchmark real-world scheme. The benchmark real-world collective contract is thus found to be inferior to the individual DC scheme.

Figure 9 gives the consumption over the life cycle of the generation born at the time of the reform, both for our benchmark scheme and the individual DC scheme. It shows that consumption is completely flat during the working phase, reflecting the fixed contribution policies of the individual DC scheme. The distribution of consumption during the retirement phase shows the opposite. This distribution is wider for the individual DC scheme, due to the fact that the individual DC scheme does not absorb any of the shocks that occur during this phase of the life cycle. Median consumption is lower than in the benchmark case. This reflects the welfare loss that is due to the combination of suboptimal investment and contribution policies.

[Figure 9 about here.]

6 Concluding remarks

In comparing collective schemes with individual schemes, we have focused on risk sharing. Yet, there are no less than three differences between the two types of pension schemes: intergenerational risk sharing, investment policies and age-dependent policy parameters. Intergenerational risk sharing relates to risk sharing with the unborn and is therefore something that is reserved for collective schemes. Investment strategies are a second source of difference. Individual schemes apply age-dependent policies. Typically, this involves reducing the equity exposure when the participant gets older. Thirdly, collective schemes apply rules specifying how contributions and pension benefits will be adjusted in response to adverse shocks that do not distinguish between different generations. Indeed, contribution rates and rates of indexation are age-independent. Furthermore, contribution and indexation policies are typically not matched with each other, which is at odds with the principle of consumption smoothing (between active and retired generations).

As risk sharing with the unborn is welfare-increasing, it favors collective schemes over individual schemes. The lack of age-dependent policies that is distinctive of collective schemes, favors individual schemes over their collective

counterparts. Taken together, collective pension schemes may therefore be either superior or inferior to individual schemes. Our calculations indeed suggest that one or the other may yield higher social welfare, depending on the specific way the schemes have been set up.

Obviously, there are caveats to our analysis. In particular, the assumption that labour income is risk-free is problematic. The same can be said of the restriction of our analysis to equity as the only risk factor. Accounting for these factors will almost surely change our numerical results. However, it is by no means clear that it would also change the basic message of our paper that real-world collective pension schemes are unsuccessful in exploiting all the possible benefits of intergenerational risk sharing. The reason for this lies in age-independent contribution, indexation and investment policies, which is something orthogonal to the set of risk factors that is included in the analysis.

An implication of our analysis that collective schemes could gain by introducing age-specific elements into their policies, *i.e.* their investment, contribution and indexation policies. This could be a way to reap the welfare gains of intergenerational risk sharing without having to bear the welfare costs of age-independent policies. Our analysis indicates that the gains from doing so could be huge.

References

- Auerbach, A. J. and L. J. Kotlikoff (1987), *Dynamic Fiscal Policy*, Cambridge University Press.
- Ball, L. and N. Mankiw (2007), “Intergenerational risk sharing in the spirit of Arrow, Debreu, and Rawls, with applications to social security design,” *Journal of Political Economy*, **115**, 523–547.
- Bodie, Z. and H. Prast (2012), “Rational pensions for irrational people: Behavioral science lessons for the Netherlands,” in L. Bovenberg, C. van Ewijk, and E. Westerhout (eds.), *The Future of Multi-Pillar Pensions*, Cambridge University Press, Cambridge, pp. 299–329.
- Bonenkamp, J., L. Meijdam, E. Ponds, and E. Westerhout (2014), “Reinventing intergenerational risk sharing,” Netspar panel paper.
- Bonenkamp, J. and E. Westerhout (2014), “Intergenerational Risk Sharing and

- Endogenous Labour Supply within Funded Pension Schemes,” *Economica*, **81**, 566–592.
- Bovenberg, A., R. Koijen, T. Nijman, and C. Teulings (2007), “Saving and investing over the life cycle and the role of collective pension funds,” *De Economist*, **155**, 347–415.
- Buccioli, A. and R. Beetsma (2010), “Inter- and intragenerational consequences of pension buffer policy under demographic, financial and economic shocks,” *CESifo Economic Studies*, **56**, 366–403.
- Campbell, J. and L. Viceira (2002), *Strategic Asset Allocation*, Oxford University Press.
- CPB (2012), “Generatie-effecten Pensioenakkoord,” CPB notitie, CPB.
- Cui, J., F. De Jong, and E. Ponds (2011), “Intergenerational risk sharing within funded pension schemes,” *Journal of Pension Economics and Finance*, **10**, 1–29.
- Draper, N. (2008), “Prudence and Robustness as Explanations for Precautionary Savings; an Evaluation,” CPB Memorandum 196, CPB.
- Gollier, C. (2008), “Intergenerational Risk-Sharing and Risk-Taking of a Pension Fund,” *Journal of Public Economics*, **92(5)**, 1463–1485.
- Gordon, R. and H. Varian (1988), “Intergenerational risk sharing,” *Journal of Public Economics*, **37**, 185–202.
- Merton, R. C. (1969), “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case,” *The Review of Economics and Statistics*, **51(3)**, 247–257.
- Rangel, A. and R. Zeckhauser (2001), “Can Market and Voting Institutions Generate Optimal Intergenerational Risk Sharing?” in J. Y. Campbell and M. Feldstein (eds.), *Risk Aspects of Investment-Based Social Security Reform*, chapter 4, University of Chicago Press, pp. 113–152.
- Samuelson, P. A. (1969), “Lifetime Portfolio Selection by Dynamic Programming,” *The Review of Economics and Statistics*, **51**, 239–246.
- Sinn, H.-W. (2000), “Why a Funded Pension System is Useful and Why It is Not Useful,” *International Tax and Public Finance*, **7**, 389–410.

Teulings, C. N. and C. G. de Vries (2006), “Generational Accounting, Solidarity and Pension Losses,” *De Economist*, **154**(1), 63–83.

Towers Watson (2011), *Global pension asset study 2011*.

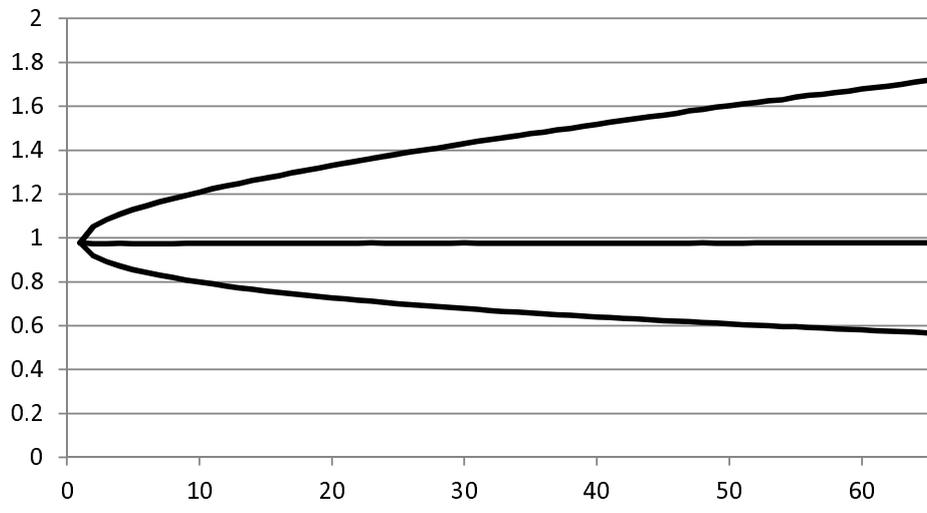


Figure 1: Consumption quantiles (5%, 50%, 95%) of the optimal individual scheme as a function of economic age

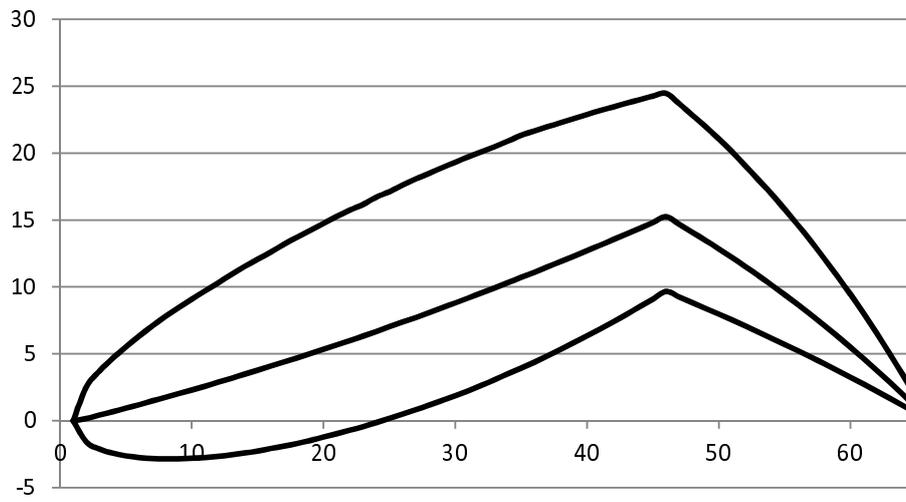


Figure 2: Initial financial wealth quantiles (5%, 50%, 95%) of the optimal individual scheme as a function of economic birth year

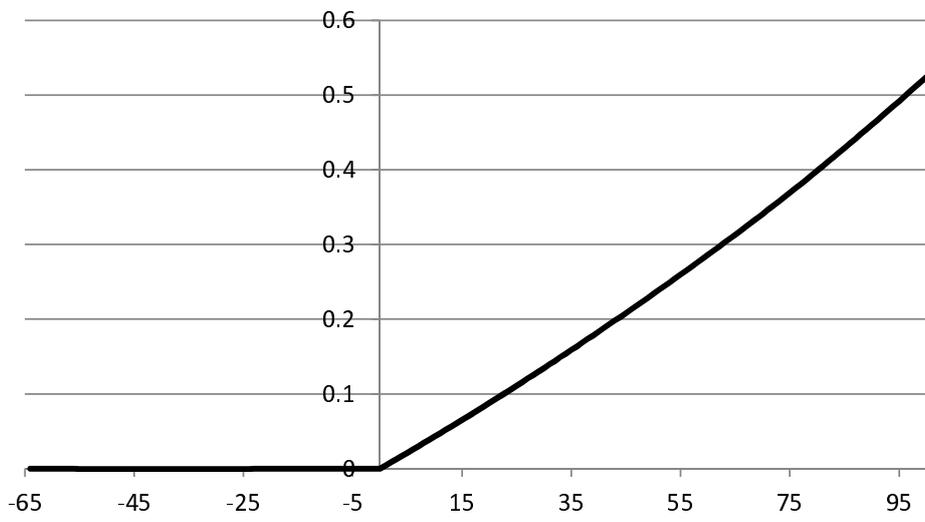


Figure 3: Welfare gains by birth year of the first-best scheme, compared to the optimal individual scheme

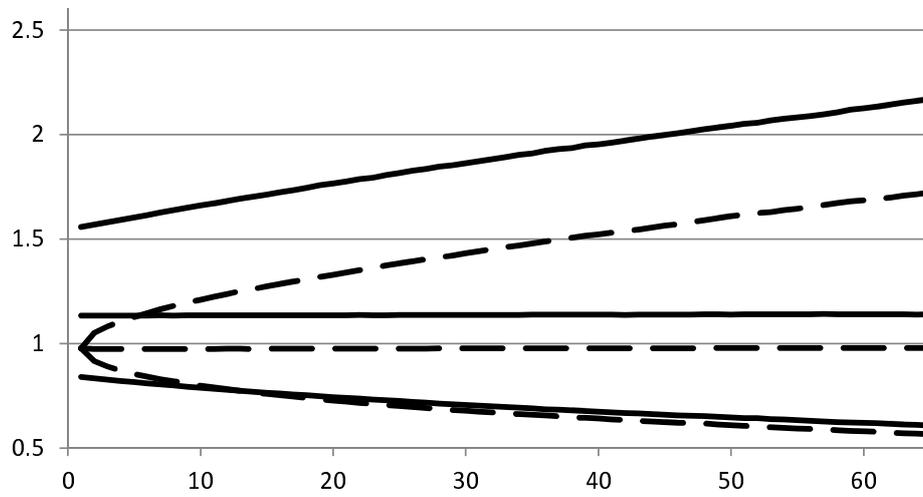


Figure 4: Consumption quantiles by age (5%, 50% and 95%) for the generation entering 20 years after the reform in the first-best scheme (—) and in the optimal individual scheme (- -)

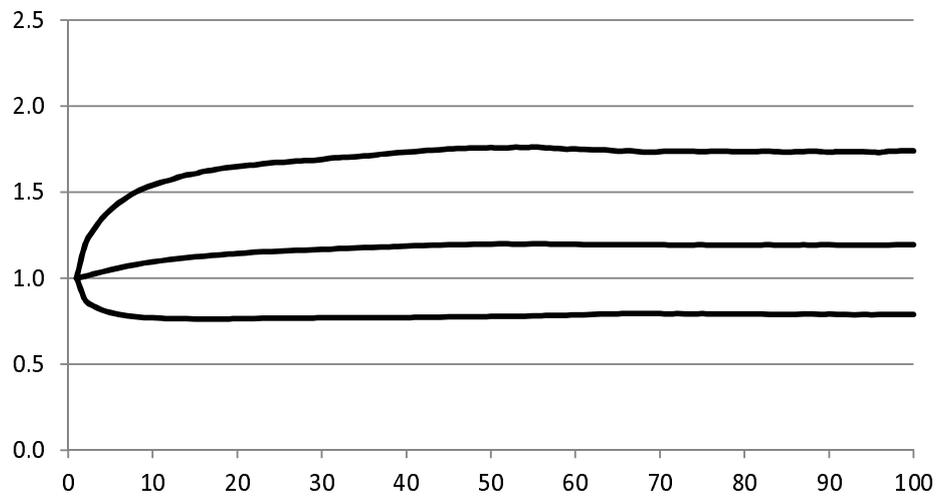


Figure 5: Distribution as represented by the 5%, 50% and 95% quantiles of the funding ratio over time

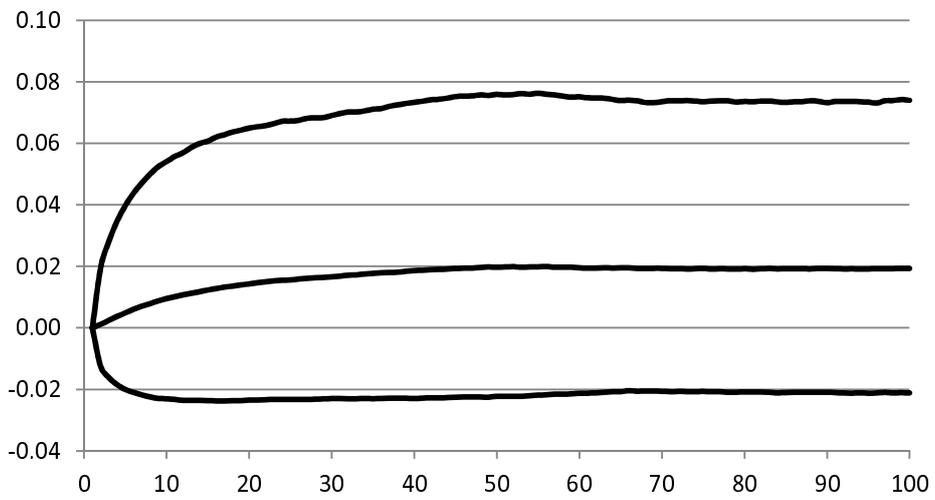


Figure 6: Distribution as represented by the 5%, 50% and 95% quantiles of the adjustment to pension rights over time

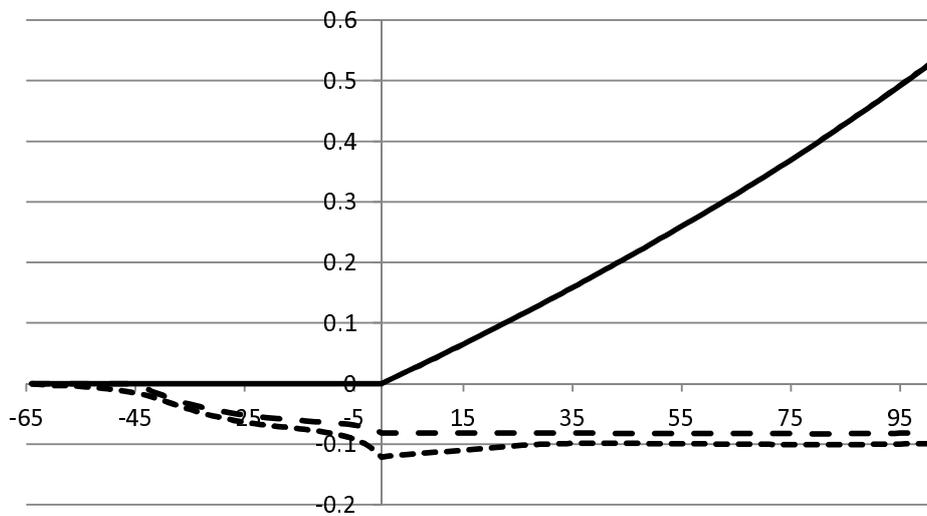


Figure 7: Welfare gains by birth year of the first-best scheme (—), the real-world collective scheme (dotted line) and the individual DC scheme (---), all compared to the optimal individual scheme

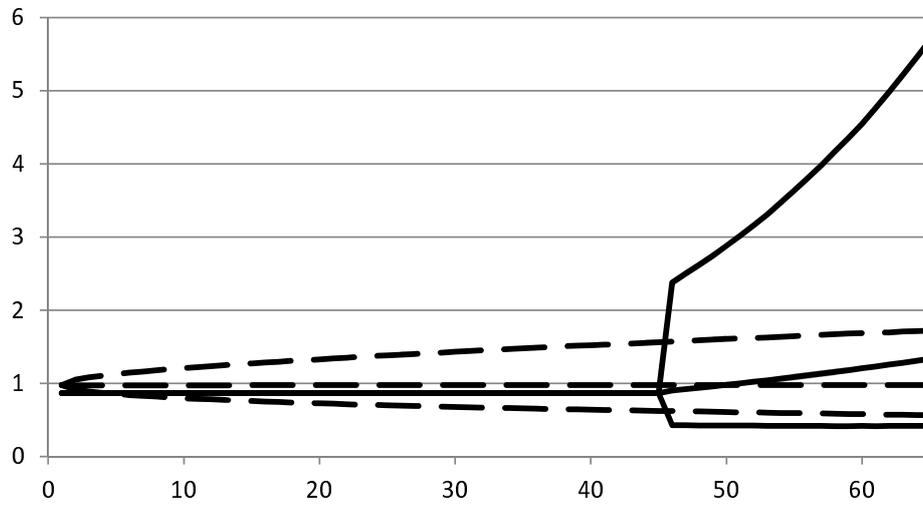


Figure 8: Consumption quantiles by age (5%, 50% and 95%) for the generation entering 20 years after the reform in the real-world collective scheme (-) and in the optimal individual scheme (- -)

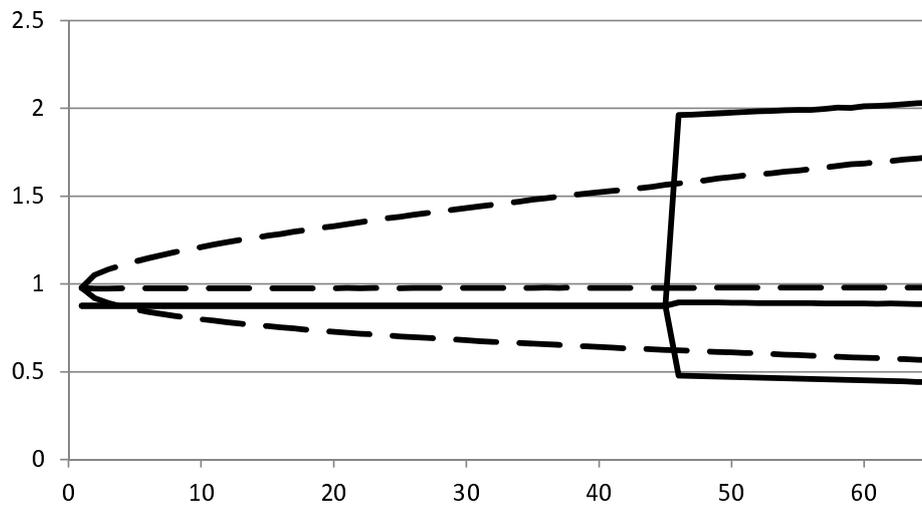


Figure 9: Consumption quantiles by age (5%, 50% and 95%) for the generation entering 20 years after the reform in the individual DC scheme (-) and in the optimal individual scheme (- -)

Table 1: Parameter values for the benchmark simulation

General	
r^f (in %)	2.0
$E(r)$ (in %)	6.0
σ_r (in %)	20.0
T^E	45
T	65
δ	0.041781
γ	5.0
w	1.0
Benchmark real-world contract	
ψ	2.0
λ_{pp}	0.0
λ_{π}	0.1

Table 2: Comparison of pension schemes

	FB	IDC	Real-world contract			
			BM	ALT 1	ALT 2	ALT 3
Assumptions						
λ_π			0.1	0.1	0.05	0.005
λ_{pp}			0.0	0.1	0.1	0.17
μ (%)			50	50	50	73
acc (%)			1.2	1.2	2.0	1.5
AEQV Results (%)						
Relative to BM	14.4	-6.5	-8.4	-6.1	-3.9	4.4
Relative to IDC	20.9		-1.9	0.4	2.6	10.9