Network for Studies on Pensions, Aging and Retirement

Netspar DISCUSSION PAPERS

Roel Beetsma and Alessandro Bucciol

The Consequences of Indexed Debt for Welfare and Funding Ratios in the Dutch Pension System

Discussion Paper 12/2010-090

The consequences of indexed debt for welfare and funding ratios in the Dutch pension system*

Roel M. W. J. Beetsma†
University of Amsterdam, Netspar, Mn Services, Tinbergen Institute, CEPR and CESifo

Alessandro Bucciol‡
University of Verona, University of Amsterdam and Netspar

This version: December 23, 2010

Abstract

Using a model of a two-pillar pension system, designed after and calibrated to the Dutch situation, we explore for the funding ratio of pension funds and the welfare of individuals the implications of replacing nominal debt in the pension fund's portfolio with indexed debt. We consider price-indexed, wage-indexed and longevity-indexed debt. The welfare consequences of indexed debt are only modest, while the same is the case for the effect on the volatility of the funding ratio when the full menu of shocks is present. The effect on the funding ratio is limited even in the presence of only the shocks to which the debt is indexed. In fact, a too large share of wage-indexed debt may destabilise the funding ratio.

Keywords: funded pensions, funding ratio, indexed debt, stochastic simulations.

JEL codes: H55, I38, C61

*Prepared for the Netspar book "The Three-Pillar Pension System Revisited: the Case of the Netherlands". Financial support from Netspar is gratefully acknowledged. We are grateful to Lans Bovenberg, Casper van Ewijk, Theo Nijman, Ed Westerhout and the participants of the conference "Macroeconomics of Pension Reform" (October 7-8, 2010, The Hague) for their many useful comments. The usual disclaimers apply.

†Contact information: University of Amsterdam, Dept. of General Economics, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands. Telephone: +31 (0)20 525 5280. Fax: +31 (0)20 525 4254. Email address: r.m.w.j.beetsma@uva.nl.

‡Contact information: University of Verona, Dept. of Economics, Viale Università 4, 37129 Verona, Italy. Telephone: +39 045 842 5448. Fax: +39 045 802 8529. Email address: alessandro.bucciol@univr.it.

1 Introduction

Dutch pension arrangements are rather unique in the world, in that they feature a second, funded pillar that is almost as large as the first, pay-as-you-go (PAYG) pillar in the Netherlands. The far majority of the employees now build up an occupational pension via a company or sectoral pension fund. It is expected that the relative size of the second pillar will rise further, because as a result of the rising life expectancy it will become more and more difficult to finance the PAYG part out of workers' contributions. Essential elements of the second pillar are "collectivity" and "solidarity" with a de jure or de facto obligation to participate. Contribution and accumulation rates are identical across the fund's participants. These elements help to keep the costs of the pension arrangement relatively low, while allowing for risk-sharing among the various groups of participants. For example, with indexation of benefits to wages, some of the wage risk is shifted onto the retired, while some of the financial market risk is shifted onto the workers when realised indexation is linked to the market performance of the fund's investment portfolio.

Within the Dutch second pillar workers build up nominal pension rights (expressed in euros) through their contributions to the pension funds. These rights promise a nominally-fixed benefit as of the retirement date. In the past, the pension rights were virtually always indexed for price or wage inflation. However, in recent years and in particular as a result of the economic and financial crisis, it has become clear that the financial health of the pension funds, as measured by the funds' buffers, is highly sensitive to a host of shocks. As a result, in many instances indexation has been abolished and the purchasing power of pensions has been eroded through inflation.

It has been frequently suggested that the vulnerability of the pension buffers and, hence, of the future pensions themselves to some of these shocks, such as inflation risk, can be reduced when the pension funds are able to invest in indexed debt. The nominal return on such debt moves up and down with the rate of inflation and, hence, to the extent that the pension fund invests in this type of debt, the value of its assets will be protected against inflation risk.

However, the Dutch government does not issue indexed debt,[1] while only a limited number of foreign governments issue such debt. None of this debt is indexed to the Dutch price level and, hence, pension funds have no other possibility than to hedge against inflation risk through derivative instruments. In this paper, and following Bucciol and Beetsma (2010), we develop an applied many-generation small open-economy OLG model with heterogeneous agents that is designed to capture the main features of the two pillars of the Dutch pension system. The model is calibrated to the Dutch economy and its unexpected aggregate shocks, which arise from demographic uncertainty (the size of newborn generations and survival probabilities), economic uncertainty (productivity growth and inflation) and financial uncertainty (bond and equity returns, and yield curve). Using stochastic simulations of the model, we explore for both the behaviour of the pension funding ratio, i.e. the ratio of the fund's assets and liabilities, and the welfare of various types of individuals the implications of substituting indexed debt for nominal debt in the portfolios of the funds. We consider price-indexed, wage-indexed and longevity-indexed debt. We limit the complexity of the framework by not explicitly modeling a government issuing the indexed debt. Hence, pension funds buy indexed debt on international markets and its returns are exogenously determined.

We find that with the full menu of shocks active the welfare effects of a switch from nominal to any of the aformentioned types of indexed debt are positive for a majority of individuals because consumption variability decreases. However, in most cases the effects are rather modest in size.[2]

[1] A discussion (in Dutch) on the desirability of the Dutch government issuing indexed debt is found in Beetsma and van Ewijk (2010).

[2] This contrasts with Koijen et al. (2010), who find that the welfare benefits of introducing indexed debt instruments can be substantial. However, their model substantially differs from ours. In particular, they have fewer

The effect on the volatility of the funding ratio when the full menu of shocks is present is rather small. The limited size of the effects may not be too surprising. To take the example of debt indexed to price inflation, the return on this type of debt still bears the risk associated with movements in the real interest rate (which, in turn, are linked to movements in the marginal productivity of capital), while the funding ratio continues to be subject to an even much wider range of risks, such as demographic risk, interest risk and stock market risk. Further, the policy rules regulating the funding ratio are identical across the various scenarios. This limits the consequences of a switch from one scenario to another for the volatility of the funding ratio and also for welfare. The implications for the funding ratio of investing in indexed debt are limited even in the presence of only the shock to which the debt is indexed, because pension funds invest only half of their value in debt instruments and at the annual frequency the actual indexation of pension rights to price or wage inflation often deviates substantially from exactly full indexation. In fact, a too large share of wage-indexed debt may destabilise the funding ratio.

The remainder of this paper is organised as follows. Section 2 describes the model. Section 3 discusses the model calibration and the rules for the adjustment of the policy parameters. Section 4 presents the results of the simulations and compares the cases of the various types of indexed debt with the case of nominal debt held by the pension fund. Finally, Section 5 concludes the paper. The appendix provides details on the estimation of the shock processes. Further details are available from an additional online appendix that can be downloaded at http://www1.fee.uva.nl/mint/beetsma.shtm.

2 The Benchmark Model

The model is an overlapping generations model with a number of D cohorts alive in any given period t. A period in our model corresponds to one year.

2.1 Cohorts and demography

Index $j = 1, ..., D$ indicates the age of the cohort, computed as the amount of time since entry into the labor force. Individuals face an exogenous age-dependent probability of dying in each period. The probability is stochastic and exhibits a downward trend, thereby causing the population on average to become older over time. Further, we assume that the cohort of newborn agents in period t is $1+n_t$ times larger than the cohort of newborn agents in period $t-1$, where n_t is also stochastic.

2.2 Skill groups and the income process

Each individual belongs to some skill group i, with $i = 1, ..., I$, and she remains in her skill group during her entire life. A higher value of i denotes a higher skill level. We assume that all the skill groups are of equal size. The skill level of a person determines her income, given her age and the macroeconomic circumstances. We allow for skill-induced differences in income, because below a certain income level individuals do not build up claims to a second-pillar pension and, hence, groups with low skills will hardly be affected by the policy differences in the second pillar. This may affect the aggregate welfare comparison between different policies. In addition, we want to capture the main elements of the Dutch system in order to be able to compare policies in a realistic setting.

Individuals work until the exogenous retirement age R and live for at most D years. During their working life ($j = 1, ..., R$), they receive a labour income $y_{i,j,t}$ given by:

sources of shocks.

$$y_{i,j,t} = e_i s_j z_t, \tag{1}$$

where e_i, $i = 1, ..., I$ is the efficiency index for skill group i, s_j, $j = 1, ..., R$ is a seniority index (income varies with age for a given skill level), and z_t is an exogenous process:

$$z_t = (1 + g_t)\, z_{t-1}, \tag{2}$$

where g_t is the exogenous *nominal* growth rate of the process and $z_0 = 1$. Hence, individuals in a given cohort in period t only differ in terms of their income, while all individuals in a given skill-group earn the same income per hour worked.

2.3 Social security

Social security is based on a two-pillar system that resembles the Dutch pension system. The first pillar is a PAYG DB program which pays a flat benefit to every retiree. It is organized by the government, which sets the contribution rate to ensure that this pillar is balanced on a period-by-period basis. Even though this pillar does not feature explicitly in our analysis, it provides an important part of retirees' income. In particular, because of the franchise for the second pillar (as explained below), for low-skilled individuals the first pillar is (virtually) the only source of income during retirement. Hence, we include it in our model to produce realistic effects of policy changes on income and welfare of the various groups in society. For example, because low-skilled groups receive hardly any income from the second pension pillar, they will also experience hardly any effect of policy changes in the second pillar. The second pillar consists of private pension funds that provide "defined benefit" nominal pensions that are usually indexed to some combination of price and productivity changes.

2.3.1 The first pillar of social security

Each period, an individual of working age pays a mandatory contribution $p^F_{i,j,t}$ to the first pillar of the pension system. This contribution depends on the size of his income $y_{i,j,t}$ relative to the thresholds $\delta^l y_t$ and $\delta^u y_t$:

$$p^F_{i,j,t} = \left\{ \begin{array}{ll} 0, & \text{if } y_{i,j,t} < \delta^l y_t \\ \theta^F_t \left(y_{i,j,t} - \delta^l y_t \right), & \text{if } y_{i,j,t} \in \left[\delta^l y_t, \delta^u y_t \right] \\ \theta^F_t \left(\delta^u y_t - \delta^l y_t \right), & \text{if } y_{i,j,t} > \delta^u y_t \end{array} \right\}, \quad j \leq R, \tag{3}$$

where δ^l, δ^u and θ^F_t are policy parameters and y_t is average income across all working individuals. Like those on low income (below $\delta^l y_t$), the retired pay no contributions, while for high-come workers the contribution is capped. In period t a retiree receives a flat benefit that is a fraction ρ^F of the average income in the economy:

$$b^F_t = \rho^F y_t. \tag{4}$$

Given the benefit formula in equation (4), each period the contribution rate θ^F_t is adjusted such that aggregate contributions into the first pillar equal aggregate first-pillar benefits.

Note that under this system an individual earning a low income pays no contributions but still receives the same benefit as an individual with a high income.

2.3.2 The second pillar of social security

Each period, an individual of working age also pays a mandatory contribution $p^S_{i,j,t}$ to the second pillar if his income exceeds the franchise income level λy_t, where parameter λ denotes the franchise as a share of average income. Specifically,

$$p^S_{i,j,t} = \theta^S_t \max\left[0, y_{i,j,t} - \lambda y_t\right], \quad j \leq R, \tag{5}$$

where θ^S_t is a policy parameter. We assume that θ^S_t is capped at a maximum value of $\theta^{S,\max} > 0$. The Dutch pension contract generally imposes a cap on the contribution rate, but this cap may differ across funds. We take account of the presence of such a cap by imposing a maximum $\theta^{S,\max}$ on θ^S_t.

A cohort entering retirement at age $R+1$ receives a benefit linked to its entire wage history. Period t benefits of an individual in skill group i of cohort j are given by:

$$b^S_{i,j,t} = M_{i,j,t}, \qquad j \geq R+1, \tag{6}$$

where $M_{i,j,t}$ is the "stock of nominal pension rights" in euros accumulated by the end of period t. It is the second-pillar pension that a retiree receives each period from this year and on, as long as this number is not revised. Variable $M_{i,j,t}$ is a stock variable in the sense that a retiree's annual benefit increases for each additional year of work that she has provided. Precisely, $M_{i,j,t}$ evolves as follows:

$$M_{i,j,t} = \left\{ \begin{array}{ll} (1+m_t) \left\{ \begin{array}{c} (1+\kappa_t \pi_t + \iota_t (g_t - \pi_t)) M_{i,j-1,t-1} \\ +\mu \max\{0, y_{i,j,t} - \lambda y_t\} \end{array} \right\}, & j \leq R \\ (1+m_t)(1+\kappa_t \pi_t + \iota_t (g_t - \pi_t)) M_{i,j-1,t-1}, & j > R \end{array} \right\}, \tag{7}$$

where parameter μ denotes the annual accrual rate of nominal rights as a share of income above the franchise level. The productivity indexation parameter ι_t and the price indexation parameter κ_t capture the degree of indexation of earlier accumulated nominal rights $M_{i,j-1,t-1}$ to real income growth, $g_t - \pi_t$, and inflation, π_t, respectively. Indexation aims at following total wage growth, in which case $\iota_t = \kappa_t = 1$. However, the actual degree of indexation may depend on the financial position of the pension fund. Further, $m_t < 0$ captures a proportional reduction in nominal rights that is applied when the pension buffer is so low that the other instruments (the indexation rates and the contribution rate) are insufficient to restore the buffer within the allowed restoration period, while $m_t > 0$ when earlier cuts in nominal rights are undone. Each individual enters the labour market with zero nominal claims ($M_{i,0,t-j} = 0$ for any i and t). Notice that, in contrast to the first-pillar pension benefit, the second-pillar benefit depends on both the cohort and skill level of the individual.

As we shall describe below, the pension fund's instrument setting strongly depends on the so-called nominal funding ratio F_t, which is the ratio between the fund's assets, A_t, and its liabilities, L_t:

$$F_t = \frac{A_t}{L_t}. \tag{8}$$

At the end of period t the pension fund's assets are the sum of the second-pillar contributions from workers in period t *minus* the second-pillar benefits paid to the retirees in period t *plus* the pension fund's assets at the end of period $t-1$ grossed up by their return in the financial markets:

$$A_t = \left(\sum_{j=1}^{R} \frac{N_{j,t}}{I} \sum_{i=1}^{I} p^S_{i,j,t} - \sum_{j=R+1}^{D} \frac{N_{j,t}}{I} \sum_{i=1}^{I} b^S_{i,j,t} \right) + \left(1 + r^f_t\right) A_{t-1}, \tag{9}$$

where

$$1+r_t^f=(1-z^e)\left(1+r_t^{lb}\right)+z^e\left(1+r_t^e\right), \tag{10}$$

where r_t^f is the nominal rate of return on the pension fund's asset portfolio with a constant share z^e invested in equities and the remainder in long-term nominal bonds. Later we shall also consider variants in which the pension fund invests in indexed bonds rather than nominal bonds.

We take the standard perspective of a small open economy with perfect capital mobility. All the asset returns, that is the returns on long-term bonds (r_t^{lb}) and equities (r_t^e) are exogenously determined on the international financial markets. Given that we do not model a domestic government sector, we can think of the bonds being issued by foreign governments with foreign tax payers liable for the repayment of the debt.[3] To avoid complicating the model further, we also assume that the pension fund's portfolio composition z^e is exogenous. Actual data for Dutch pension funds show a stable composition over the years, which may point to pension funds aiming stable targets for the various asset categories.

We assume that the long-term bonds held by the pension fund always have a 10-year maturity. This implies that at the end of each year bonds of 9-year maturity are sold to purchase new 10-year bonds. In more detail, the fund's annual portfolio rebalancing operation works as follows. In year $t-1$, say, the pension fund buys 10-year zero-coupon bonds for an amount of B_{t-1}. Denoting the return on 10-year bonds by $r_{10,t-1}^b$, the value at maturity of the bonds is

$$P_{t+9}=B_{t-1}\left(1+r_{10,t-1}^b\right)^{10}, \tag{11}$$

hence, the present value B_{t-1} of the bond holdings in year $t-1$ is:

$$B_{t-1}=\frac{P_{t+9}}{\left(1+r_{10,t-1}^b\right)^{10}}.$$

In year t, only 9 years of maturity are left, and the bond return is $r_{9,t}^b$. The present value B_t is then

$$B_t=\frac{P_{t+9}}{\left(1+r_{9,t}^b\right)^{9}}.$$

Combining with (11) we obtain the following expression:

$$B_t=B_{t-1}\frac{\left(1+r_{10,t-1}^b\right)^{10}}{\left(1+r_{9,t}^b\right)^{9}}=B_{t-1}\left(1+r_t^{lb}\right).$$

The fund's liabilities are the sum of the present values of current and future rights *already accumulated* by the cohorts currently alive:

$$L_t=\sum_{j=1}^{D}\frac{N_{j,t}}{I}\sum_{i=1}^{I}L_{i,j,t}, \tag{12}$$

where $L_{i,j,t}$ is the liability to the cohort of age j and skill level i, which is computed by discounting the projected future nominal benefits resulting from the current stock of nominal rights against a term structure of annual nominal interest rates $\{r_{k,t}\}_{k=1}^{D}$:

[3] Introduction of a domestic government sector would be beyond the scope of the present paper, as it would introduce many complications. While with perfect capital mobility nominal asset returns would still be exogenous, we would have to introduce taxes that are needed to finance government expenditures and the repayment of debt. Tax policies could be used to affect redistribution and risk-sharing among generations. These issues cannot be addressed in the present paper.

$$L_{i,j,t} = \left\{ \begin{array}{ll} E_t \left[\sum_{l=R+1-j}^{D-j} \left(\prod_{k=1}^{l} \psi_{j+k,t-j+1} \right) \frac{1}{(1+r_{l,t})^l} M_{i,j,t} \right], & \text{if } j \leq R, \\ E_t \left[\sum_{l=0}^{D-j} \left(\prod_{k=1}^{l} \psi_{j+k,t-j+1} \right) \frac{1}{(1+r_{l,t})^l} M_{i,j,t} \right], & \text{if } j > R, \end{array} \right\}. \tag{13}$$

When $j \leq R$, we discount all future benefits to the current year t, but of course they will only be paid out once individuals have retired. Importantly, notice that the computation of the liabilities excludes the effects of possible *future* indexation. Hence, the liabilities are less than the total pension wealth accumulation by current generations. This is also the reason why pension funds that aim at maintaining the purchasing power of the accumulated rights need to maintain a funding ratio that is substantially above 100%.

2.4 Individuals

Period utility of individuals is given by

$$u(\widetilde{c}_{i,j,t}) = \frac{\widetilde{c}_{i,j,t}^{1-\gamma}}{1-\gamma},$$

where γ is the coefficient of relative risk aversion for the entire term in square brackets, and $\widetilde{c}_{i,j,t}$ is *real* consumption,

$$\widetilde{c}_{i,j,t} = \frac{c_{i,j,t}}{\prod_{s=1}^{t} (1+\pi_s)},$$

where $c_{i,j,t}$ is nominal consumption (the price level at the end of period 0 is unity). Further, $\widetilde{y}_{i,j,t}$ is labour or pension income net of contributions:

$$\widetilde{y}_{i,j,t} = \left\{ \begin{array}{ll} y_{i,j,t} - p^F_{i,j,t} - p^S_{i,j,t} & \text{if } j \leq R \\ b^F_t + b^S_{i,j,t} & \text{if } j \geq R+1 \end{array} \right\}.$$

To keep matters as simple as possible and enhance intuition, we assume that individuals take no decisions, implying that they simply consume all of their disposable income $\widetilde{y}_{i,j,t}$. Hence, their savings are zero. Since they enter the world with zero assets, their assets $a_{i,j,t}$ are zero throughout their life.

The individual's value function is

$$V_{i,j,t} = E_t \left[\sum_{l=0}^{D-j} u(\widetilde{c}_{i,j+l,t+l}) \frac{\beta^l}{\psi_{j,t-j+1}} \left(\prod_{k=-j+1}^{l} \psi_{j+k,t-j+1} \right) \right].$$

2.5 The shocks

We assume that there are only aggregate, hence no individual-specific shocks. Seven types of aggregate exogenous shocks hit the economy. Specifically, we allow for demographic shocks (to the growth rate of the newborns cohort and to the survival probabilities), inflation rate shocks, nominal income shocks (which, together with the inflation shock, produce a shock to the productivity growth rate) and financial market shocks (to bond returns, equity returns and the yield curve). The shocks are collected in the vector $\omega_t = \left[\epsilon^n_t, \epsilon^\psi_t, \epsilon^g_t, \epsilon^\pi_t, \epsilon^b_t, \epsilon^e_t, \epsilon^b_{2,t}, ..., \epsilon^b_{D,t} \right]'$ with elements

- ϵ^n_t: shock to the newborn cohort growth rate n_t,

- ϵ_t^{ψ}: a vector of shocks to the set of survival probabilities $\left\{\psi_{j,t-j+1}\right\}_{j=1}^{D}$,
- ϵ_t^{g}: shock to the nominal income growth rate g_t,
- ϵ_t^{π}: shock to the inflation rate π_t,
- ϵ_t^{b}: shock to the one-year nominal bond return r_t^b,
- ϵ_t^{e}: shock to the nominal equity return r_t^e,
- $\epsilon_{k,t}^{b}, k = 2, ..., D$: shock to the yield curve at maturity $k > 1$, $r_{k,t}^b$.

All these shocks affect the size of the funding ratio (equation (8)), whereas only the demographic shocks affect the first pillar of the pension system. As a consequence, the key parameters of the pension system have to be adjusted to restore the balance in the first pillar and to maintain sustainability of the second pillar. Below we give a brief description of each shock process.

The demographic shocks are independent of each other and of all other shocks (at all leads and lags). The growth rate n_t of the newborns cohort depends on deterministic and random components:

$$n_t = n + \epsilon_t^n, \tag{14}$$

where n is the mean and ϵ_t^n the innovation at time t, which follows an AR(1) process with parameter φ:

$$\epsilon_t^n = \varphi \epsilon_{t-1}^n + \eta_t^n, \quad \eta_t^n \tilde{} N\left(0, \sigma_n^2\right). \tag{15}$$

The survival probabilities evolve according to a Lee and Carter (1992) model:

$$\ln\left(1 - \psi_{j,t-j+1}\right) = \ln\left(1 - \psi_{j,t-j}\right) + \tau_j\left(\chi + \epsilon_{t-j+1}^{\psi}\right), \quad \epsilon_{t-j+1}^{\psi} \tilde{} N\left(0, \sigma_{\psi}^2\right), \quad j = 1, ..., D, \tag{16}$$

with τ_j an age-dependent coefficient, χ a constant growth factor (to describe the historical trend increase in survival probabilities) and ϵ_{t-j+1}^{ψ} an innovation at time $t - j + 1$ that follows an i.i.d. process with variance σ_{ψ}^2.

We allow the shocks to the inflation rate, the nominal income growth rate, the one-year bond return and the equity return to be correlated with each other and over time. These variables feature the following multivariate process:

$$\begin{pmatrix} \pi_t \\ g_t \\ r_t^b \\ r_t^e \end{pmatrix} = \begin{pmatrix} \pi \\ g \\ r^b \\ r^e \end{pmatrix} + \begin{pmatrix} \epsilon_t^{\pi} \\ \epsilon_t^{g} \\ \epsilon_t^{b} \\ \epsilon_t^{e} \end{pmatrix}, \tag{17}$$

with annual means $\left(\pi, g, r^b, r^e\right)'$, and innovations $\left(\epsilon_t^{\pi}, \epsilon_t^{g}, \epsilon_t^{b}, \epsilon_t^{e}\right)'$ following a VAR(1) process,

$$\begin{pmatrix} \epsilon_t^{\pi} \\ \epsilon_t^{g} \\ \epsilon_t^{b} \\ \epsilon_t^{e} \end{pmatrix} = \mathbf{B} \begin{pmatrix} \epsilon_{t-1}^{\pi} \\ \epsilon_{t-1}^{g} \\ \epsilon_{t-1}^{b} \\ \epsilon_{t-1}^{e} \end{pmatrix} + \begin{pmatrix} \eta_t^{\pi} \\ \eta_t^{g} \\ \eta_t^{b} \\ \eta_t^{e} \end{pmatrix}, \quad \text{with} \quad \begin{pmatrix} \eta_t^{\pi} \\ \eta_t^{g} \\ \eta_t^{b} \\ \eta_t^{e} \end{pmatrix} \tilde{} N\left(\underset{4x1}{\mathbf{0}}, \underset{4x4}{\widetilde{\Sigma}}\right). \tag{18}$$

For the estimation of (17) and (18) we use inflation and income data for the Netherlands (source is OECD, 2009) and 1-year bond and equity data for the U.S. (sources are, respectively, Federal Reserve, 2009, and Datastream, 2009).

We finally turn to the term structure of annual nominal interest rates (the yield curve). We set the interest rate at one-year maturity equal to the one-year bond interest rate arising from the above multivariate process. To describe the remaining components of the yield curve, we focus on the rates in excess of the bond interest rate at maturity 1, $\widetilde{r}^b_{k,t}$. Following the prevailing literature (see, e.g., Evans and Marshall, 1998; Dai and Singleton, 2000), we model the excess interest rates as a vector autoregressive distributed lag (VADL) process with lag 1:

$$\begin{pmatrix} \widetilde{r}^b_{2,t} \\ \widetilde{r}^b_{3,t} \\ \vdots \\ \widetilde{r}^b_{D,t} \end{pmatrix} = \mathbf{\Gamma}_0 + \mathbf{\Gamma}_1 \begin{pmatrix} \widetilde{r}^b_{2,t-1} \\ \widetilde{r}^b_{3,t-1} \\ \vdots \\ \widetilde{r}^b_{D,t-1} \end{pmatrix} + \mathbf{\Gamma}_2 \begin{pmatrix} \pi_{t-1} \\ g_{t-1} \\ r^b_{t-1} \\ r^e_{t-1} \end{pmatrix} + \begin{pmatrix} \epsilon^b_{2,t} \\ \epsilon^b_{3,t} \\ \vdots \\ \epsilon^b_{D,t} \end{pmatrix}, \tag{19}$$

with

$$\begin{pmatrix} \epsilon^b_{2,t} & \epsilon^b_{3,t} & \dots & \epsilon^b_{D,t} \end{pmatrix}' \,\tilde{}\, N(\mathbf{0}, \Sigma). \tag{20}$$

Each period t, the excess interest rate at maturity k, $\widetilde{r}^b_{k,t}, k \geqslant 2$, is a linear combination of deterministic and random components. The deterministic part is a function of several variables at time $t-1$: the excess interest rates at all maturities $k \geqslant 2$ and the four macro and financial variables whose shocks follow the VAR(1) process (18). The random part is given by the innovations $\epsilon^b_{k,t}$, which may be correlated across maturities.

Actual yields at any maturity $k \geq 1$ are then built as the sum of the VADL(1) realisations and the realisation to the one-year bond interest rate:

$$\begin{pmatrix} r^b_{1,t} \\ r^b_{2,t} \\ r^b_{3,t} \\ \vdots \\ r^b_{D,t} \end{pmatrix} = \begin{pmatrix} 0 \\ \widetilde{r}^b_{2,t} \\ \widetilde{r}^b_{3,t} \\ \vdots \\ \widetilde{r}^b_{D,t} \end{pmatrix} + \begin{pmatrix} r^b_t \\ r^b_t \\ r^b_t \\ \vdots \\ r^b_t \end{pmatrix}.$$

The average yield curve $\{r^b_k\}_{k=1}^D$ is given by:

$$\begin{pmatrix} r^b_2 \\ r^b_3 \\ \vdots \\ r^b_D \end{pmatrix} = \begin{pmatrix} r^b \\ r^b \\ \vdots \\ r^b \end{pmatrix} + (I - \mathbf{\Gamma}_1)^{-1} \left(\mathbf{\Gamma}_0 + \mathbf{\Gamma}_2 \begin{pmatrix} \pi \\ g \\ r^b \\ r^e \end{pmatrix} \right), \tag{21}$$

where we have applied $E\left[\tilde{r}^b_{k,t}\right] = E\left[\tilde{r}^b_{k,t-1}\right]$ to (19) because of stationarity.

We estimate (19) using U.S. data (source is Federal Reserve, 2009).

2.6 Welfare comparisons between policies

We are interested in welfare comparisons based on simulations with one or more shocks. Generally speaking, welfare differences measured in this way are the sum of a "redistribution effect" of a shift from the benchmark scenario A to the alternative scenario B in the absence of shocks and a "risk-sharing effect" that captures the difference between the total welfare effect of the shift from the benchmark to the alternative under shocks and the redistribution effect. In the comparisons that we will consider below, only the risk-sharing effect will be present.

Welfare comparisons are made at the aggregate level by reporting the share of individuals alive at $t = 1$ in favour of the alternative. At the individual level (i.e., based on the combination of skill and age) welfare comparisons takes place at the start of period $t = 1$ of our simulations (see below) for those cohorts alive at that moment and at the start of the first year of life for the cohorts that are born later. Individual welfare is measured as the certain "consumption equivalent change" (CEC), as is standard in the literature of life-cycle models (see, e.g., Cocco et al., 2005, and Krueger and Kubler, 2006). It is defined as the change in certainty-equivalent consumption over the remainder of the individual's life under scenario B relative to scenario A. That is, if $V_{i,j,t}(S)$ is the welfare for skill group i of the generation aged j in year t under scenario $S \in \{A, B\}$, we compute[4]

$$CEC_{i,j,t} = \left(\frac{V_{i,j,t}(B)}{V_{i,j,t}(A)}\right)^{\frac{1}{1-\gamma}} - 1.$$

We consider $CEC_{i,j,1}$ for the alive generations and $CEC_{i,1,t}$ for the unborn generations with $t > 1$.

3 Calibration, simulation details and the policy rule

3.1 Benchmark calibration

We assume that the economically active life of an agent starts at age 25. Individuals work for $R = 40$ years until they reach age 65. They live for at most $D = 75$ years, until age 100. The discount factor β is set to 0.96, as is common practice in the macroeconomic literature (see, e.g., Imrohoroglu, 1989, or Krebs, 2007). The coefficient of relative risk aversion γ is set to 3. While there is substantial uncertainty about the size of this parameter, this assumption accords quite well with the assumed risk aversion in much of the macroeconomic literature (see, e.g., Imrohoroglu et al., 2003) as well as estimates at the individual level (for example Beetsma and Schotman, 2001). Further, we assume a rather large number of $I = 10$ different skill groups in order to be able to also capture the consequences of the lowest skill groups not participating in the second pillar. The efficiency index $\{e_i\}_{i=1}^{I}$ is given by the income deciles in the Netherlands for the year 2000 taken by the World Income Inequality Database (WIID, version 2.0c, May 2008). We normalise the index to have an average value of 1. The seniority index $\{s_j\}_{j=1}^{I}$ uses the average of Hansen's (1993) estimation of median wage rates by age group. We take the average between males and females and interpolate the data using the spline method.

The social security parameters mimic the institutional framework in the Netherlands. For the first pillar, the Dutch Tax Office ("Belastingdienst") reports for 2008 a maximum income assessable for contributions of $3,850.40$ euros per month. We therefore set our upper income threshold for

[4] Certainty-equivalent consumption c_S under scenario S, i.e. the constant consumption level at all future dates of life and under all states of the world under this scenario, follows from

$$V_{i,j,t}(S) = k_0 u(c_S) = k_0 \frac{c_S^{1-\gamma}}{1-\gamma},$$

where k_0 is some constant involving the individual's discount factor, survival probabilities and inflation rates. Hence,

$$\begin{aligned} CEC_{i,j,t} &= \frac{c_B}{c_A} - 1 = \left[\frac{(1-\gamma)V_{i,j,t}(B)/k_0}{(1-\gamma)V_{i,j,t}(A)/k_0}\right]^{\frac{1}{1-\gamma}} - 1 \\ &= \left(\frac{V_{i,j,t}(B)}{V_{i,j,t}(A)}\right)^{\frac{1}{1-\gamma}} - 1 = u^{-1}\left(\frac{1}{1-\gamma}\frac{V_{i,j,t}(B)}{V_{i,j,t}(A)}\right) - 1. \end{aligned}$$

contributions to $\delta^u = 1.10$, which is roughly equal to $3,850.40 * 12/42,403$, where EUR $42,403$ is our imputation for the economy's average income for 2008.[5] The lower income threshold is set to $\delta^l = 0.4685$, so as to generate an initial contribution rate of $\theta_1^F = 12.77\%$, identical to the initial second-pillar contribution rate, $\theta_1^F = \theta_1^S$, which is calculated assuming that aggregate contributions at time 1 coincide with aggregate benefits in the absence of shocks. This value of θ_1^S is close to the actual value in the Netherlands. In our simulations we will cap θ_t^S at $\theta^{S,\max} = 25\%$. We finally set the benefit scale factor at $\rho^F = 0.2435$.

For the second social security pillar, we set $z^e = 0.50$ for any level of the funding ratio F_t. Our choice roughly corresponds to the balance sheet average for Dutch pension funds over the past 10 years (source: DNB, 2009). Because the bond and equity investments in the pension fund's portfolio generally have different realised returns, at the end of each period t the portfolio is reshuffled such that the system enters the next period $t+1$ again with the original portfolio weight $z^e = 0.50$. We set the pension accrual rate μ at 2% and the franchise parameter λ to 0.381.[6] The choices of ρ^F and λ are meant to generate realistic replacement rates that on average are equal to 30.40% for the first pillar and 37.60% for the second pillar. The first-pillar replacement rate is higher for less skilled groups and ranges on average between 12.06% and 63.33% across the skill groups, while the second-pillar replacement rate is higher for more skilled groups and ranges on average between 3.78% and 56.64%. Overall, the total replacement rate from the two pillars is higher for more skilled groups and ranges on average between 67.11% and 68.70%.

Given the initial value of the second-pillar contribution rate $\theta_1^S = 12.77\%$, we choose initial assets A_0 so as to generate an initial funding ratio F_1 of 1.25 in the absence of shocks.[7]

The deterministic component of the growth rate of the newborn cohort, $n = 0.2063\%$, is the average annual growth rate based on the estimation of an order-one moving-average model for the annual number of births in the Netherlands over the period $1906 - 2005$ (source is the Human Mortality Database, 2009). Our calibration of survival probabilities is based on the estimation of a Lee and Carter (1992) model using Dutch period survival probabilities.[8] The combination of survival probabilities and birth rates determines the size of each cohort. The starting value of the old-age dependency ratio (i.e., the ratio of retirees over workers) is 20.99%, in line with the OECD (2009) figures for the Netherlands in 2005.

Crucial is the calibration of average price inflation, average nominal income growth and the average bond and equity returns. The calibrated averages are reported in the final four lines of Table 1. We loosely follow the literature in this regard (see, e.g., Brennan and Xia, 2002, and van Ewijk et al., 2006) and set the average inflation rate at $\pi = 2\%$, the average nominal income growth rate at $g = 3\%$ (which corresponds to an average real productivity growth of 1% per year),

[5] In Eurostat the most recent number on average income in the Netherlands refers to year 2005. The same source also provides the minimum income until year 2008. Exploiting the correlation between average and minimum income, we run an OLS regression of average income on minimum income. As a result, we predict the average income for year 2008 to be EUR $42,403$.

[6] The maximum accrual rate that is fiscally facilitated in the Netherlands is 2.25% for pension arrangements linked to average lifetime wages and 2% for pension arrangements linked to final wages.

[7] Initial assets A_0 are 1.4731 times aggregate income in the economy. This is on the high side compared to the actual Dutch situation. However, in our model every worker participates in the pension fund, while in the Netherlands this is only part (though a majority) of the employed. Moreover, a substantial fraction of the workers has its pension arranged through insurance companies, while the self-employed do not participate in pension funds either (they have the possibility to build up their own pension through an insurance company, but the financial reserves of insurance companies are not considered part of the pension buffers).

[8] With these probabilities, the average population age is initially set to 48.21 years and the remaining life expectancy is 33.54 years, as opposed to 33.23 years for a 48-year old in 2005 according to the actual data (see Human Mortality Database, 2009).

the average one-year bond interest rate at $r^b = 3\%$, and the average equity return at $r^e = 5.5\%$.[9]

Table 1. Benchmark calibration of the parameters

Symbol	Description	Calibration
General setting		
D	Number of cohorts (= maximum death age -25)	75
R	Number of working cohorts (= retirement age -25)	40
β	Discount factor	0.96
γ	Relative risk aversion parameter	3
$\{e_i\}_{i=1}^{I}$	Efficiency index	WIID (2008)
$\{s_j\}_{j=1}^{I}$	Seniority index	Hansen (1993)
First pillar pension parameters		
$\left\{\delta^l, \delta^u\right\}$	Income thresholds in the contribution formula	$\{0.469, 1.10\}$
ρ^F	Benefit scale factor	0.2435
Second pillar pension parameters		
z^e	Equity share in fund portfolio	0.5
$\left\{K^S, K^L\right\}$	Restoration periods	$\{5, 15\}$
μ	Second-pillar pension accrual rate	0.02
λ	Franchise share	0.381
F_1	Initial funding ratio	1.25
$\theta^{S,\max}$	Upper bound on contribution rate	0.25
Annual averages of the random variables		
π	Inflation rate	2%
g	Nominal income growth rate	3%
r^b	One-year nominal bond return	3%
r^e	Equity return	5.5%

3.2 Simulation details

We simulate $Q = 1,000$ times a sequence of vectors of unexpected shocks over $2D - 1 + 250 = 399$ years, drawn from the joint distribution of all the shocks. Our welfare calculation is based on the economy as of the D^{th} year in the simulation. Hence, we track only the welfare of the cohorts that are alive in that year, implying that those that die earlier are ignored, and we track the welfare of cohorts born later, the latest one dying in the final period of the simulation. In other words, the total number of years of one simulation run equals the time distance between the birth of the oldest cohort that we track and the complete extinction of the last unborn cohort that we track. At each moment there are D overlapping generations. For convenience, we relabel the D^{th} year in the simulation as $t = 1$. The purpose of simulating the first $D - 1$ years is to simply generate a distribution of the assets held by each cohort at the end of $t = 0$.

In each simulation run, we assume that the ageing process stops after $t = 40$, by setting the trends in newborn growth rates and in survival probabilities to zero after $t = 40$, while the shocks to both processes and thus the demographic uncertainty still remain. Hence, also mortality rates at any given age no longer fall. For two reasons we stop the ageing process after 40 years . First,

[9] With these assumptions, the funding ratio, which is initially set at $F_1 = 1.25$, is equal to 1.16 after 75 years in the absence of shocks and policy intervention.

it is hard to imagine that mortality rates continue falling for many more decades at the same rate as they did in the past. In particular, important common mortal diseases have already been eradicated, while it will become more and more difficult to treat remaining lethal diseases. Second, some important ageing studies, such as those by the Economic Policy Committee and European Commission (2006) and the United Nations (2009), only project ageing (and its associated costs) up to 2050 (hence 40 years from now), because the uncertainty in the projections becomes too large over larger horizons.

To allow for the cleanest possible comparison among the various policies, during each simulation run we use the same shock series for all policies, while, moreover, during the initialisation phase of each simulation run no policy responses occur. Hence, the situation at the start of $t = 1$ (before choices are made) is identical in each run under the various policies. Specifically, at the start of the initialisation phase all policy parameter values are set at their $t = 1$ levels (complete price indexation, zero productivity indexation and a constant contribution rate) and they remain unaffected during this phase. Further, at the start of the initialisation phase the pension rights of all individuals are set to zero. During the initialisation phase individuals accumulate pension rights according to (7), with zero indexation to productivity growth and full indexation to price inflation ($\iota_t = 0$ and $\pi_t = 1$), under the assumption that there are no shocks and income evolves according to (1) and (2), with g_t thus being constant. After the initialisation phase, at the end of $t = 0$, the process z_t is rescaled to unity and the nominal pension rights of all the individuals are rescaled by the same factor. Using (12) and (13), we can then compute total pension liabilities at the end of $t = 0$. Because welfare depends on the size of the buffer after the initialisation period in the simulation run, we reset the stock of pension fund assets such that the buffer at the end of $t = 0$ equals 1.25. In other words, the assets and liabilities of the pension fund at the end of $t = 0$ are identical across all policy variants.

3.3 The policy rule

The baseline and the alternatives to be considered below are subject to various sources of suboptimality. First, instead of optimising their decisions, individuals consume their disposable income. Second, we consider rules for the adjustment of the policy instruments that are not necessarily optimal, but that are intended to capture the main features of the policies imposed by the Dutch pension supervisor and followed by Dutch pension funds.

The government automatically adjusts the contribution rate $\theta_t^F \in [0, 1]$ to maintain a balanced first pension pillar. On average, this contribution rate increases over the years along with the ageing of the population.

More policy options are available to affect the funding ratio of the second pillar. There are three key parameters, whose period $t + 1$ values are determined on the basis of the funding ratio F_t in period t: the contribution rate $\theta_{t+1}^S \in \left[0, \theta^{S,\max}\right]$, the two indexation parameters $\{\kappa_{t+1} \in [0, 1], \iota_{t+1} \geq 0\}$ and, as a last resort, a reduction in the nominal pension rights ($m_{t+1} < 0$). The Board of the pension fund selects the contribution rate and the indexation parameters, but can only reduce nominal rights under special circumstances, as described below.

Policymakers start with a benchmark parameter combination $\left\{\theta_1^S, \kappa_1, \iota_1\right\}$ and a funding ratio equal to $\xi^m = 1.25$. We set $\kappa_1 = 1$ (complete price indexation) and $\iota_1 = 0$ (zero productivity indexation). We define two threshold values for the funding ratio, ξ^l and ξ^u, with $\xi^l < \xi^m < \xi^u$ and $\xi^l > 1$. In particular, we set $\xi^l = 1.05$ and $\xi^u = 1.50$. All policies are identical when the funding ratio F_t is above ξ^m. In that case, after restoring possible earlier cuts in nominal rights, the fund's Board sets the contribution rate at its initial level θ_1^S, price indexation to $\kappa_{t+1} = 1$ and productivity indexation to $\iota_{t+1} = \frac{F_t - \xi^m}{\xi^u - \xi^m}$. Hence, productivity indexation increases linearly in F_t

and becomes complete at ξ^u; it continues to increase at the same rate as F_t rises above ξ^u. This way the funding ratio is stabilised from above.

As mandated by the Dutch Pension Law, when the funding ratio falls below ξ^m, but remains above ξ^l, a long-term restoration plan is started, while when it falls below ξ^l a short-term restoration plan is started. The latter situation is termed "underfunding". The long-term restoration plan requires a restoration of the funding ratio to at least ξ^m in at most $K^l = 15$ years (ignoring possible future shocks), while the short-term restoration plan requires its restoration to at least ξ^l in at most $K^s = 5$ years (ignoring possible future shocks). In the case of both a short-term or a long-term restoration plan, productivity and price indexation are always reduced first. If the adjustment is insufficient, the other instrument is also adjusted. Conform Dutch Law, when there is underfunding ($F_t < \xi^l$) and the adjustments in the indexation parameters and the contribution rate are jointly insufficient, nominal rights are scaled back by whatever amount is necessary to eliminate the underfunding within the allowed restoration period. In the case of a long-term restoration plan, nominal rights remain untouched.

4 Simulation results and comparisons

Financial instruments potentially help to protect against some specific risks. In principle, including such instruments in the pension fund's asset portfolio may reduce the overall volatility of the funding ratio and, therefore, limit the need for adjustments of the policy parameters. In this section, while keeping everything else unchanged, we compare the benchmark of a pension fund portfolio with long-term nominal debt with alternatives in which this nominal debt is replaced by, respectively:

- *price-indexed bonds*, of which the yields depend on the actual inflation rate plus a fixed spread, $r_t^{lb} = \alpha^p + (\pi_t - \pi)$.
- *wage-indexed bonds*, of which the yields depend on the actual nominal wage growth rate plus a fixed spread, $r_t^{lb} = \alpha^w + (g_t - g)$.
- *longevity-indexed bonds*, of which the yields depend on the difference between the actual and expected remaining life expectancies (which we denote by, respectively, T_t and $\widetilde{T}_t$) of the population, $r_t^{lb} = \left(1 + \alpha^l\right)^{1-\left(T_t - \widetilde{T}_t\right)}$. This way of modelling the return on longevity-indexed debt follows that in Blake and Burrows (2001).

We choose the parameters α^p, α^w, and α^l so as to generate an average return on indexed bonds equal to the return on long-term nominal bonds, r^{lb}. This way our analysis is purely devoted to comparing the consequences of a reduction in the mismatch between the funds assets and liabilities, while keeping the expected path for the value of its assets unaltered. Hence, on purpose we ignore the potential difference in risk premia paid on the various types of debt instruments. Notice that the respective instruments provide high (nominal) returns when there are upward shocks in price, wage growth and life expectancy. As mentioned earlier, all asset returns are exogenous.

Below we first investigate separately the use of each type of indexed debt in the presence of only the shock that it is supposed to hedge against. That is, we focus on price-indexed bonds when there are price shocks only, wage-indexed bonds when there are wage shocks only, and longevity-indexed bonds when there are demographic shocks only (shocks to survival probabilities and the growth rate of the newborn cohorts). Finally, we explore the role of all types of indexed debt in the presence of the full menu of shocks.

4.1 Price-indexed debt and inflation shocks only

We first investigate the use of price-indexed debt in the presence of inflation shocks only. In such an environment price-indexed debt should be beneficial. After all, price-indexed debt is often promoted as an instrument to reduce the inflation risk of pensions. Figure 1 shows the median funding ratio under the benchmark and with price indexed debt. During the first ten years it rises, because income through contributions and asset returns dominates the outlays in the form of pension benefits. Due to the ageing process after ten years the fund's outlays start to dominate the fund's income and the funding ratio starts to decline. Its decline is halted at the 125% level due to the restoration plans imposed by the supervisor. The median funding ratio remains essentially unaffected by the inclusion of price-indexed debt in the pension fund's portfolio. Comparing the volatility of the funding ratio under the two scenarios, we see that its median coefficient of variation is indeed quite a lot (one-fifth to a quarter) lower with price-indexed debt.

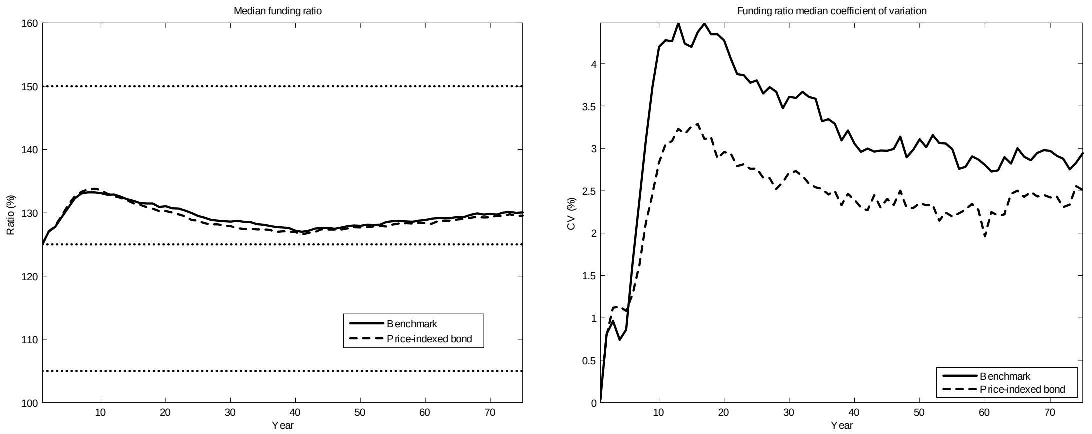

a. Median funding ratio **b.** Median coefficient of variation

Figure 1. Fund properties: price shocks only

Table 2 reports statistics on the funding ratio under the baseline and with price-indexed bonds. The volatilities of both assets and liabilities, as measured by the median coefficient of variation, are highest under price-indexed debt while, as expected, the volatility of the funding ratio is lowest. These facts are reconciled by the observation that the increases in volatility under price-indexed debt result from the density functions of both assets and liabilities becoming more skewed to the right, implying a larger frequency of relatively small values of assets and liabilities and, hence, less volatility in the ratio of assets over liabilities. The second panel of Table 2 reports the frequencies with which the funding ratio is below certain thresholds.

We observe that the likelihood of the funding ratio falling below $\xi^l = 105\%$ is zero, which is the result of the fact that inflation shocks alone are a relatively minor source of volatility in the funding ratio. The likelihood that the funding ratio falls below ξ^m is slightly lower under price-indexed debt than in the baseline case. In the third panel of the table we dissect the frequency of the funding ratio being below ξ^m into cases in which only the indexation rate needs to be adjusted, cases in which both the indexation rate and the contribution rate need to be adjusted and this adjustment is enough and cases in which the indexation and contribution rates are both adjusted, but this is insufficient for a (long-term) restoration plan. The cases are reported as frequencies of all the simulation observations. The fourth panel of the table reports average (over all observations) values of the policy parameters. The averages for the contribution rate and the rate of indexation to price inflation under price-indexed debt are rather similar to the benchmark, while the average rate of indexation to productivity improvements is somewhat lower. The volatilities of all instruments are smaller under price-indexed debt than under the alternative. However, the differences are small, except in the case of indexation to productivity.

Table 2. Policy comparison: price shocks only

%	**Non-indexed debt**	**Price-indexed debt**
Funding ratio volatility		
Median coefficient of variation	3.140	2.400
Median coeff. of var., assets	2.103	4.081
Median coeff. of var., liabilities	3.613	5.132
Assets-liabilities correlation	99.796	99.815
Autocorrelation, funding ratio	84.118	81.046
Probability of a funding ratio below a given threshold		
Below ξ^l	0	0
Below ξ^m	20.199	18.740
Below ξ^u	99.859	99.939
Joint probability of change in the policy parameters and funding ratio below ξ^m		
Only indexation rate	8.586	8.100
Index. and contr. rates enough	5.867	5.760
Index. and contr. rates not enough	5.756	4.880
Average policy parameters (%, standard deviation in parenthesis)		
κ_t	82.819	83.991
	(35.836)	(35.316)
ι_t	17.558	13.925
	(20.480)	(16.582)
θ_t^S	14.262	14.136
	(3.436)	(3.288)
$\kappa_t \pi_t + \iota_t (g_t - \pi_t)$	2.134	2.147
	(1.358)	(1.416)
$\theta_t^S \tilde{z}_t$	14.264	14.136
	(3.436)	(3.288)
Welfare comparison relative to benchmark		
% better off of those alive at $t = 1$	-	48.751

Note: autocorrelation measures the correlation of the ratio in two consecutive years.

Figure 2 depicts the cohort-specific welfare consequences under price-indexed bonds relative to the benchmark of non-indexed bonds. We are interested in separating the risk-sharing effects of switching to indexed bonds from the pure redistribution effects that may arise in the absence of shocks. However, the redistribution effects are zero, because in the absence of shocks the benchmark and alternative cases are identical. After all, in the absence of shocks, the returns on indexed debt are identical to those on long-term nominal debt. Therefore, Figure 2 displays the pure risk sharing effects. We see that the welfare effects of using price-indexed debt are on average rather small for each of the cohort-skill combinations depicted. However, these effects are unevenly distributed. Older cohorts alive at $t = 1$ experience a welfare loss, while the younger ones benefit from a switch to indexed debt in the pension fund's portfolio. The share of the population alive at $t = 1$ benefitting from the switch is slightly less than half (see final panel Table 2). We see that the cohort-wise pattern of the welfare effects is similar for the various skill levels. The effects are more muted (in both directions of the horizontal axis) for the lower-skilled individuals, because for these individuals the second pillar is a relatively small source of income when they are old.

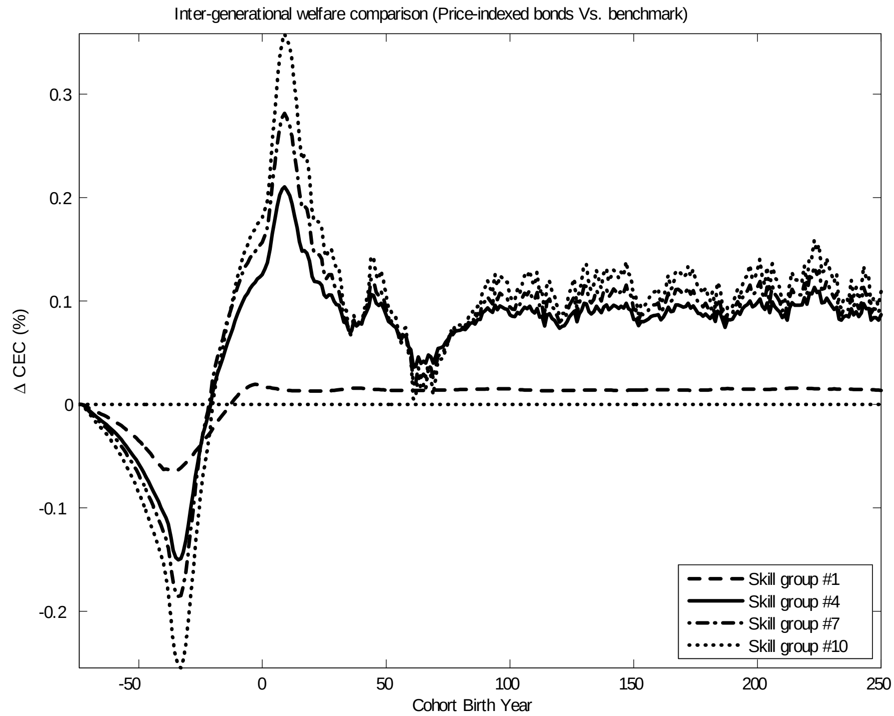

Figure 2. Welfare comparison: price-indexed bonds versus benchmark

To explore further the sources of these findings, it is useful to see how average lifetime consumption and its volatility are affected by the switch to price-indexed debt. Panel a of Figure 3 shows the (average over the simulations of the) consumption level over the remaining lifetime for those already alive at $t = 1$ and over the entire lifetime for those that are not yet born at $t = 1$. We see that average consumption is virtually identical under price-indexed debt and the benchmark. Panel b of Figure 3 presents the (average over the simulations of the) coefficient of variation of the consumption level over the first ten years of life for those not yet born at $t = 1$ and the first ten years of the remaining lifetime for those already alive at $t = 1$. Consumption variability is relatively high for the retired and the older workers. The reason is that the funding ratio starts at 125%, which implies a high chance that the older generations at $t = 1$ are confronted with restoration plans that involve reduced indexation of their pension rights and, hence, lower pension benefits. Table 2 reports the average pension contribution $\theta_t^S \tilde{z}_t$ (the contribution rate times its base) and its volatility.[10] Both are lower under price-indexed debt. As individuals consume their disposable income, changes in pension contributions translate directly into changes in consumption of workers. Hence, because pension contributions are more stable under price-indexed debt, consumption variability for those born after $t = 1$ is lower under this type of debt, which explains why those individuals benefit from the switch to price-indexed debt. The same pattern is observed if we use a window of five or twenty years. It is no longer observed when we calculate our consumption variability measure over the entire (remaining) lifetime. However, given that our measure neglects the effects

[10] Here, $\tilde{z}_t$ is "detrended income", given by $\tilde{z}_t = \left(1 + \epsilon_t^g\right) \tilde{z}_{t-1}$, where we recall that ϵ_t^g is the shock to the nominal income growth rate. Because nominal income is growing at a positive average rate, we consider the product of the contribution rate and detrended income as a more suitable indicator of the impact of the pension system on disposable income.

of discounting and the increasing chances of dying as one becomes older, calculated over the entire (remaining) lifetime the measure gives far too large weight to consumption uncertainty far into a person's life and forms an inadequate indicator of the welfare consequences of indexed debt. The disposable income of the retirees is directly affected by the "overall indexation" $\kappa_t \pi_t + \iota_t (g_t - \pi_t)$ of their pension rights. The volatility of overall indexation increases under a switch from non-indexed to price-indexed debt, which explains why welfare of the older generations is negatively affected by the switch. The source of this higher volatility is that an increase in inflation is more beneficial for the pension fund's assets when it invests in indexed debt, which allows for a higher indexation rate κ_t at given inflation rate. Vice versa, when inflation decreases. The higher correlation between κ_t and π_t raises the volatility of overall indexation.

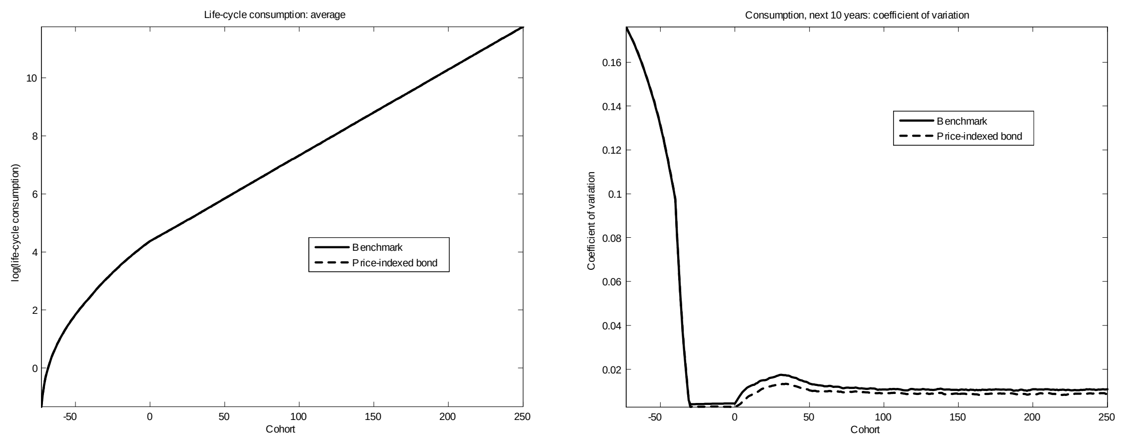

a. Average over lifetime **b.** Coefficient of variation over the next 10 years

Figure 3. Household consumption

4.2 Wage-indexed debt and wage shocks only

In this subsection we assume that the only source of shocks are the nominal wages. All the other shocks are set to zero. The consequences of introducing wage-indexed debt under these circumstances may be somewhat counterintuitive. In fact, we see that the replacement of nominal debt with wage-indexed debt in the pension fund's portfolio *de*stabilises the funding ratio. The intuition is as follows. Start from the situation with nominal debt. When a positive wage shock occurs, the pension fund's assets rise more than its liabilities. Its assets rise because the increase in the wage rate raises contributions. Liabilities also rise, but to a lesser extent, because the indexation of the nominal rights to wage shocks is generally only partial. Now, suppose that wage-indexed debt is substituted for the nominal debt. A positive wage shock now has an even larger positive effect on the value of the fund's assets, thereby raising the funding ratio even further. Of course, a more volatile funding ratio requires more frequent adjustment of the policy parameters. However, this offsets only part of the increase in the volatility. The welfare effects of the switch from nominal to wage-indexed debt are again generally small for the various individuals in the economy and will not be discussed any further. Figures for the behaviour of the funding ratio, welfare and household consumption are available in the companion online appendix.

4.3 Longevity-indexed debt and demographic shocks only

In this subsection, we assume that the only shocks are those to the number of newborns and to the survival probabilities, which affect life expectancy. An unexpected increase in life expectancy raises the fund's liabilities and, hence, has a negative effect on the funding ratio. In the benchmark case with nominal debt assets are unaffected. However, by substituting the longevity-indexed debt for nominal debt the mismatch between assets and liabilities is reduced, because the value of assets rises in case of a positive life-expectancy shock. Indeed, the volatility of the funding ratio falls with the use of longevity-indexed debt. The profiles of average lifetime consumption and consumption variability are very similar to those under price-indexed debt and are not discussed here. The uncertainty of life expectancy around its average is rather limited, implying that the welfare effects of a switch to longevity-indexed debt for the various individuals are generally small. Also in this case, figures for the funding ratio, welfare and household consumption are available in the companion online appendix.

The small size of the welfare effects may not be surprising given that the amount of uncertainty to be hedged away is only small. Table 3 below reports the average remaining life expectancy $\widetilde{T}_t$, as well as the 90% confidence bands around this average. While $\widetilde{T}_t$ rises quite substantially (more than 3 years) over a period of four decades, the confidence band is very narrow compared to this increase. The confidence bands on $\widetilde{T}_t$ are also translated into confidence bands for the return according to the formula $r_t^{lb} = \left(1 + \alpha^l\right)^{1-\left(T_t - \widetilde{T}_t\right)}$. In view of the historical movements in nominal interest rates these confidence bands are also narrow.

Table 3. Average remaining life expectancy

t	$\widetilde{T}_t$	T_t, **90% conf. int.**	r_t^{lb} **(%), 90% conf. int.**
1	27.3	(27.2, 27.3)	(2.82, 3.18)
20	28.9	(28.8, 28.9)	(2.73, 3.12)
40	30.5	(30.4, 30.5)	(2.75, 3.17)

4.4 Simulation results for full set of shocks

We now turn to the simulations of each type of indexed debt under the full set of shocks. The format in which the results are reported follows that in the previous subsections. Figure 4 shows the median funding ratio (panel a) and its volatility as captured by its median coefficient of variation (panel b) under the benchmark scenario and the three alternative cases. Again the median funding ratio tends to increase during the first ten years. However, the increase is much sharper than in the presence of a single shock only. The median funding ratio also remains higher over the median run and stabilizes it at around 140%. Although the averages of the shocks are zero, the funding ratio is a nonlinear function of those shocks and, hence the median with the full menu of shocks differs from that with a single shock only. In particular, because the funding ratio is regulated more strictly at its lower bound than at its upper bound negative shocks are more strongly offset than positive shocks, thereby on average pushing up the funding ratio upwards. While the median funding ratio exhibits rather similar patterns under the various scenarios, the figure shows that it tends to be somewhat less volatile under the alternatives to the benchmark of nominal debt.

Importantly, the volatility of the funding ratio is substantially higher than in the presence of a single shock only. The reason is that the various individual shocks feature low or zero correlation. In particular, shocks or groups of shocks do not offset each other. Hence, the presence of each of the individual shocks adds to the volatility of the funding ratio.

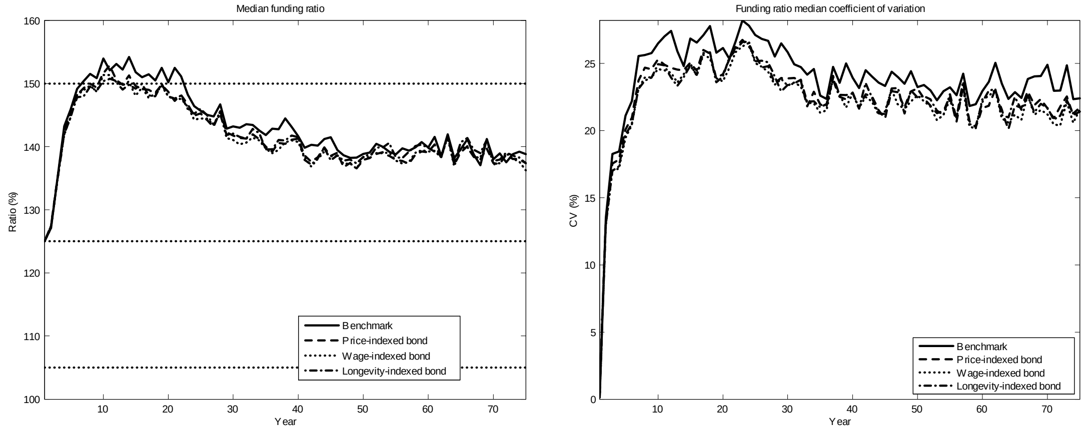

a. Median funding ratio **b.** Median coefficient of variation

Figure 4. Fund properties: full set of shocks

Table 4 reports statistics on the funding ratio under the benchmark and the three alternative cases. For the reason laid out above, the volatilities of both assets and liabilities, as measured by the median coefficient of variation are higher under the alternatives than under the benchmark, while the volatility of the funding ratio is lower under the alternatives. The second panel of Table 4 reports the frequencies with which the funding ratio is below certain thresholds. Not surprisingly, there is now a non-negligible chance that the funding ratio falls below 105%. This likelihood is somewhat lower, though, under the alternatives than under the benchmark. Again, the fourth panel of the table reports average values (over all observations) of the policy parameters. The averages for the contribution rate and the rate of indexation to inflation are very similar over the various cases, while the average rate of indexation to productivity is somewhat higher under the benchmark. Also the volatilities of the contribution rate and the rate of indexation to inflation are very similar in the various cases. However, the alternatives to the benchmark benefit from a more substantial reduction in the volatility of the rate of indexation to productivity.

The final line of Table 4 reports the aggregate welfare consequences of including indexed debt in the fund's portfolio, as measured by the percentage of those alive at $t = 1$ that strictly prefer the alternative scenario to the benchmark one. The measure shows that around 80% or more of these individuals prefer a scenario in which the pension fund invests in indexed debt.

Table 4. Policy comparison: full set of shocks

%	Benchmark	Price-ind.	Wage-ind.	Longevity-ind.
		Funding ratio volatility		
Median coefficient of variation	23.835	22.409	22.081	22.303
Median coeff. of var., assets	21.245	23.047	23.745	25.713
Median coeff. of var., liabilities	27.251	29.600	30.162	31.089
Assets-liabilities correlation	85.641	85.727	85.844	85.484
Autocorrelation, funding ratio	59.566	60.668	60.771	61.236
		Probability of a funding ratio below a given threshold		
Below ξ^l	16.664	16.195	15.943	15.828
Below ξ^m	34.088	34.273	34.359	34.071
Below ξ^u	56.431	57.817	58.061	57.712
		Joint probability of change in the policy parameters and funding ratio below ξ^m		
Only indexation rate	4.469	4.744	4.811	4.579
Index. and contr. rates enough	29.231	29.147	29.109	29.116
Reduction in nominal rights	0.388	0.383	0.439	0.376
		Average policy parameters (%, standard deviation in parenthesis)		
κ_t	49.034	49.561	49.719	49.498
	(47.447)	(47.939)	(47.939)	(48.025)
ι_t	99.332	91.580	90.569	93.329
	(208.054)	(191.010)	(188.206)	(195.408)
θ_t^S	19.141	19.082	19.061	19.057
	(5.914)	(5.923)	(5.925)	(5.931)
$\kappa_t \pi_t + \iota_t (g_t - \pi_t)$	2.598	2.507	2.496	2.522
	(3.996)	(3.748)	(3.709)	(3.810)
$\theta_t^S \tilde{z}_t$	19.340	19.328	19.314	19.347
	(6.986)	(7.126)	(7.126)	(7.224)
		Welfare comparison relative to benchmark		
% better off of those alive at $t = 1$	-	85.357	85.357	79.516

Note: Autocorrelation measures the correlation of the ratio in two consecutive years.

Figure 5a considers the specific case of price-indexed bonds and shows the welfare effects per cohort for some specific skill groups. The welfare consequences of moving from the benchmark situation to one in which the pension fund invests in price-indexed bonds are positive for most cohort-skill combinations. Even most retirees are now better off. While the average overall indexation rate is lower under indexed debt, this effect is dominated by the reduced volatility of the overall indexation rate. The benefit of the shift to indexed debt is smaller for lower-skilled groups, because for them second-pillar income is relatively smaller compared to first-pillar income. Some skill-cohort combinations are worse off. This happens only for relatively low-skilled groups. In the absence of shocks the lowest-skilled group would be indifferent about the type of debt the fund invests in, because the wage of this group would never exceed the franchise. However, in the presence of shocks it sometimes does and the lowest-skilled group suffers from the slightly increased uncertainty about the pension contributions which for this group dominates the reduced uncertainty about indexation leading to a more stable benefit once they are retired. We find that the welfare effect is quantitatively never large, with a maximum *CEC* effect of 0.8%. In most instances the effect is rather modest. Figure 5b shows the welfare effects per cohort for some specific skill groups for the case of wage-indexed bonds, while Figure 5c shows those effects for the case of longevity-indexed debt. The patterns of the welfare implications are very similar to those for a switch to price-indexed debt.

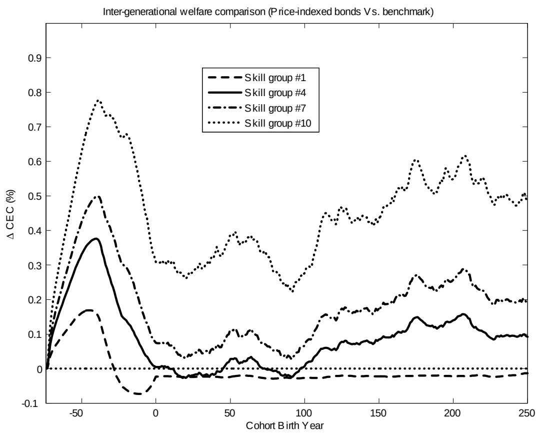

a. Price-indexed debt

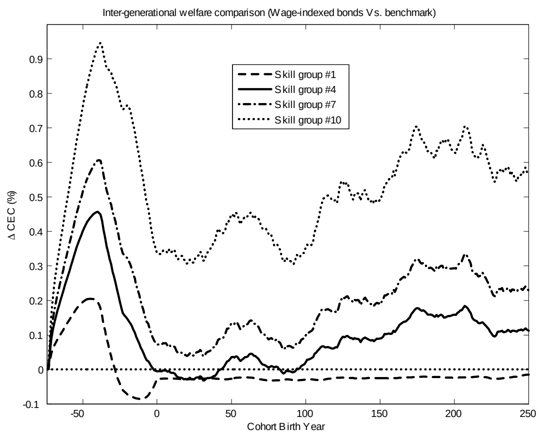

b. Wage-indexed debt

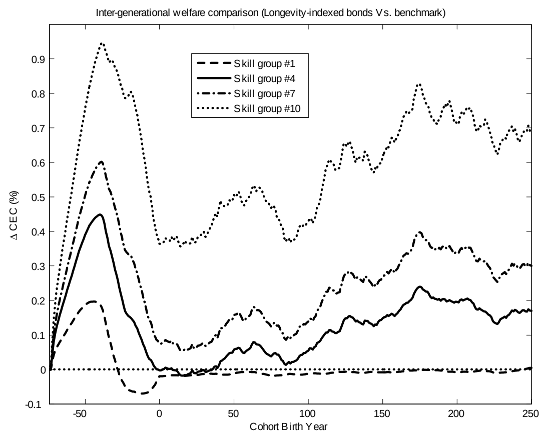

c. Longevity-indexed debt

Figure 5. Welfare comparison versus benchmark

5 Conclusions

In this paper we explored the implications of replacing nominal debt with indexed debt in pension fund portfolios for both the behaviour of the funds' funding ratios and the welfare of (groups of) individuals. We considered three types of indexed debt: price-indexed debt, wage-indexed debt and longevity-indexed debt. Our simulations were conducted using an overlapping-generations model of a small-open economy that was calibrated to the Dutch situation both in terms of pension arrangements and the shocks hitting the economy. To gain more insight into the effects of including indexed debt in the fund's portfolios, we first considered the different types of shocks (inflation shocks, wage shocks and demographic shocks) individually, before moving on to simulations in the presence of the full menu of shocks. The consequences of including indexed debt may appear somewhat counterintuitive to the proponents of investing in indexed debt. While price-indexed (longevity-indexed) debt reduces funding ratio volatility in the case of inflation (life expectancy) shocks only, investing in wage-indexed debt may actually destabilise the funding ratio by worsening the mismatch between assets and liabilities in the presence of wage shocks. Some robust conclusions emerge from the analysis. First, the volatility of the funding ratio and, hence, the frequency of changes in the policy parameters, is substantially higher in the presence of the full menu of shocks than in the presence of any single type of shock. This implies that the inclusion of indexed debt of some particular type can only have a limited effect on the funding ratio volatility when the full menu of shocks is active. Secondly, with all shocks active, the welfare effects of a switch from nominal to indexed debt are positive for a substantial majority of individuals, although they tend to be modest in size in most cases.

We have limited the complexity of our framework by not explicitly modeling a government issuing the indexed debt. An important avenue for further research would be to allow for such an extension and to consider individuals both as pension fund participants and payers of taxes to a government whose budget constraint is affected by inflation risk. In such an analysis we need to integrate the consequences of inflation risk for retirement benefits, pension contributions and tax benefits and we need to investigate how the overall effect of inflation risk on individuals' resources can be potentially mitigated if the government issues indexed debt that is (partially) held by pension funds.

References

[1] **Beetsma, R.,** and **C. van Ewijk** (2010), Over de Wenselijkheid van de Uitgifte van Geïndexeerde Schuld door de Nederlandse Overheid, *Netspar NEA Paper*, No. 30.

[2] **Beetsma, R.,** and **P. Schotman** (2001), Measuring Risk Attitudes in a Natural Experiment: Data from the Television Game Show LINGO, *Economic Journal* 111(474), 821-848.

[3] **Blake, D.**, and **W. Burrows** (2001)**,** Survivor Bonds: Helping to Hedge Mortality Risk, *The Journal of Risk and Insurance* 68(2), 339-348.

[4] **Brennan, M.J.,** and **Y. Xia** (2002), Dynamic Asset Allocation under Inflation, *Journal of Finance*, 57(3), 1201-1238.

[5] **Bucciol, A.**, and **R. Beetsma** (2010), Inter- and Intra-generational Consequences of Pension Buffer Policy under Demographic, Financial and Economic Shocks, *CESifo Economic Studies* 56(3), 366-403.

[6] **Cocco, J.F., F.J. Gomes** and **P.J. Maenhout** (2005), Consumption and Portfolio Choice over the Life-Cycle, *Review of Financial Studies* 18(2), 491-533.

[7] **Dai, Q.,** and **K. Singleton** (2000), Specification Analysis of Affine Term Structure Models, *Journal of Finance* 55(5), 1943-1978.

[8] **Datastream** (2009), www.datastream.com .

[9] **DNB** (2009), www.statistics.dnb.nl .

[10] **Economic Policy Committee and European Commission** (2006), The Impact of Ageing on Public Expenditure: Projections for the EU25 Member States of Pensions, Health-Care, Long-Term Care, Education and Unemployment Transfers (2004-2050), *European Economy*, Special Reports, 1.

[11] **Evans, C.,** and **D. Marshall** (1998), Monetary Policy and the Term Structure of Nominal Interest Rates: Evidence and Theory, *Carnegie-Rochester Conference Series on Public Policy* 49(1), 53-111.

[12] **Ewijk, C. van, N. Draper, H. ter Rele** and **E. Westerhout** (2006), Ageing and the Sustainability of Dutch Public Finances, *CPB Special Publication*, 61.

[13] **Federal Reserve** (2009), www.federalreserve.gov/releases/H15.

[14] **Hansen, G.D.** (1993), The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked, *Journal of Applied Econometrics*, 8(1), 71-80.

[15] **Human Mortality Database** (2009), www.mortality.org.

[16] **Imrohoroglu, A.** (1989), Costs of Business Cycles with Indivisibilities and Liquidity Constraints, *Journal of Political Economy* 97(6), 1364-1383.

[17] **Imrohoroglu, A., S. Imrohoroglu** and **D.H. Joines** (2003), Time Inconsistent Preferences and Social Security, *Quarterly Journal of Economics* 118(2), 745-784.

[18] **Koijen, R.S.J., T.E. Nijman** and **Bas J.M. Werker** (2010), Optimal Annuity Risk Management, *Review of Finance*, forthcoming.

[19] **Krebs, T.** (2007), Job Displacement Risk and the Cost of Business Cycles, *American Economic Review* 97(3), 664-686.

[20] **Krueger, D.** and **F. Kubler** (2006), Pareto Improving Social Security Reform when Financial Markets are Incomplete!? *American Economic Review* 96(3), 737-755.

[21] **Lee, R.D.,** and **L.R. Carter** (1992), Modeling and Forecasting U.S. Mortality, *Journal of the American Statistical Association* 87(419), 659-671.

[22] **OECD** (2009), www.stats.oecd.org.

[23] **United Nations** (2009), www.un.org/esa/population.

A Appendix: estimates of the shock processes

A.1 Growth rate newborn cohort

We estimate the model (14) and (15) for the number of births in the Netherlands between 1906 and 2005 (source is Human Mortality Database, 2009). This yields $n = 0.0021$, $\varphi = -0.0624$ (standard error 0.1093) and $\sigma_n = 0.0492$ (standard error 0.0129). In our simulations we set $n = 0$ after year $t = 40$.

A.2 Survival probabilities

We collect from the Human Mortality Database (2009) Dutch period life tables from 1906 to 2005. We define by $\psi^p_{j,t}$ the population probability of surviving year t for individuals aged $j-1$ at the start of year t. It is measured by comparing at the end of period t the number of alive individuals of age $j-1$ with the number of alive individuals of age j. (Hence, probability $\psi^p_{j,t}$ should distinguished from the cohort survival probability $\psi_{j,t-j+1}$.) To distinguish the trend from fluctuations, we estimate with singular value decomposition the parameters of the Lee and Carter (1992) model:

$$\ln\left(1-\psi^p_{j,t}\right) = \alpha_j + \tau_j \chi_t + \eta^\psi_t,$$

subject to the suggested restrictions that $\sum_{t=1}^{T} \chi_t = 0$, where $t = 1,..,T$ is the sample period, and $\sum_{j=1}^{D} \tau_j = 1$, and where α_j and τ_j are age-varying parameters, χ_t is a time-varying vector and η^ψ_t is a random disturbance distributed as $N\left(0, \widetilde{\sigma}^2_\psi\right)$. With the proposed restrictions the estimate of α_j will be the average probability over the sample that someone dies at age j , when having survived up to age $j-1$. We assume that the mortality index χ_t evolves as a random walk with drift χ:

$$\chi_t = \chi_{t-1} + \chi + \epsilon^\psi_t,$$

with $\epsilon^\psi_t \tilde{} N\left(0, \sigma^2_\psi\right)$. We estimate $\hat{\chi} = -1.6733$ and $\hat{\sigma}_\psi = 0.0957$. This implies a trend fall in the probability of dying at any age j, conditional on having survived up to age $j-1$. In our simulations we set $\hat{\chi} = 0$ after year $t = 40$.

From the period life table estimates and the trend in the mortality index we calculate the cohort life tables as follows:

$$\begin{aligned} \ln\left(1-\psi_{j,t-j+1}\right) &= \hat{\alpha}_j + \hat{\tau}_j \left(\hat{\chi}_{t-j+1} + j\hat{\chi}\right) \\ &= \hat{\alpha}_j + \hat{\tau}_j \hat{\chi}_{t+1}, \end{aligned}$$

where $t-j+1$ is the year of birth of the cohort. In our model, the survival probabilities of the cohort born in year $t = 0$ are set equal to those of the actual cohort of individuals born in 1950. It is easy to see that the conditional survival probabilities of cohorts of age j are as follows linked over time:

$$\ln\left(1-\psi_{j,t-j+2}\right) = \ln\left(1-\psi_{j,t-j+1}\right) + \hat{\tau}_j \hat{\chi}.$$

The following table based on simulations provides an idea of the quantitative implications of the model estimates. We see that the confidence intervals are rather narrow, but that expected survival probabilities increase rather substantially during the first four decades of the simulations.

	Likelihood surviving to 65			Likelihood surviving to 80		
t	Expected	90% conf. int.	95% conf. int.	Expected	90% conf. int.	95% conf. int.
1	65.9	(65.7, 66.1)	(65.7, 66.2)	20.6	(20.4, 20.7)	(20.4, 20.7)
20	72.4	(72.1, 72.6)	(72.1, 72.6)	27.3	(27.0, 27.4)	(27.0, 27.5)
40	77.9	(77.6, 78.0)	(77.5, 78.0)	34.7	(34.3, 34.9)	(34.3, 34.9)

A.3 Economic and financial shocks

Estimation of the model (17) and (18) is at an annual frequency over the period 1976-2005. Given the limited number of annual observations relative to the number of parameters to be estimated, we decided to limit the VAR process to one lag only. However, re-estimating a VAR(1) for the residuals of the initial VAR(1) model yields insignificance of all coefficients, except the bond return which is still significant at the 10% level. The estimate of matrix **B** in (18) is reported in Table A.1., panel a. The annual innovations follow a multivariate normal distribution centered at 0 and with an estimated covariance matrix given in Table A.1., panel b. Panel a of Table A.1. shows that only in the specification of the equity return the Wald chi-squared test on the joint significance of the coefficients does not reject the null hypothesis of a purely random (white noise) process. Further, as one might have expected, the inflation process is highly persistent (our simulations imply a correlation of inflation with lagged inflation of 74%). Given that the coefficients of the other variables are insignificant, the inflation equation in the VAR model quite closely resembles the estimates of a univariate AR model for inflation. Estimated on a constant with one lag, the coefficient of the first lag is 0.69 (standard error is 0.10). With two lags, the coefficient estimates of the first and second lag are, respectively, 0.99 (standard error is 0.18) and -0.29 (standard error is 0.15). With three of more lags the second lag ceases to be significant at the 5% level. Also the wage growth process is persistent, although somewhat less so than inflation.

Table A.1. VAR(1) regression

a. Deterministic coefficient estimates (matrix **B** in (18))

Variable	**Inflation**	**Wage**	**Bond**	**Equity**
Inflation (-1)	0.7685***	0.5258***	0.0584	-0.3263
	(0.1789)	(0.1848)	(0.2668)	(2.4723)
Wage (-1)	-0.1757	0.0108	0.0222	-2.7298
	(0.1828)	(0.1888)	(0.2726)	(2.5258)
Bond (-1)	0.0670	0.0479	0.8700***	0.8933
	(0.0692)	(0.0714)	(0.1032)	(0.9560)
Equity (-1)	-0.0062	-0.0133	0.0152	-0.0123
	(0.0128)	(0.0132)	(0.0190)	(0.1764)
Wald chi-squared	58.5525	38.4297	98.2896	4.8642
p-value	0.0000	0.0000	0.0000	0.3015

*Note: standard deviations in parentheses. ***: significant at 1%. Wald chi-squared: test on the joint significance of the coefficients in each column, following a chi-squared distribution with 4 d.o.f.*

b. Residual covariances and correlations (%)

Variable	**Inflation**	**Wage**	**Bond**	**Equity**
Inflation	0.0107	*33.3503*	*34.8257*	*-26.7663*
Wage	0.0037	0.0114	*-25.9594*	*-6.7091*
Bond	0.0056	-0.0043	0.0238	*-11.9154*
Equity	-0.0396	-0.0102	-0.0263	2.0449

Note: correlations in italic.

The estimated standard deviation of the equity return (14.30%) is consistent with the historical values reported in the literature (e.g., see Cocco et al., 2005).

A.4 The bond yield curve

We use an annual time series of US yield returns at maturities 2, 3, 5, 7, 10, 20 and 30 (the only observed maturities – source is Federal Reserve, 2009) over the period 1976-2006. In the sample there are occasionally missing values for the yields at maturities 20 and 30, that we impute using a linear interpolation method. The shocks estimated in (19) are highly correlated, usually above 80% and never below 71%. The shocks tend to be more volatile at longer maturities, with a variance ranging from 0.0003 at maturity 2 to 0.0027 at maturity 30, but they remain small compared to the variance (0.0238) of shocks to the one-year bond returns estimated in the model (17) and (18). Having estimated (19) for the maturities that are available, we use a linear interpolation to obtain the interest rates at any discrete maturity between 1 and 30. Interest rates at maturities longer than 30 are set equal to the interest rate at maturity 30. The average bond yield curve exhibits a quadratically-looking profile as a function of maturity k. It increases monotonically up to $k = 30$, where it reaches an estimated interest rate of 4.26%.