Network for Studies on Pensions, Aging and Retirement

Netspar DISCUSSION PAPERS

Leon Bettendorf and Thijs Knaap

The Redistribution of Macroeconomic Risks by Dutch Institutions

Discussion Paper 11/2010-087

The Redistribution of Macroeconomic Risks by Dutch Institutions

Leon Bettendorf and Thijs Knaap (CPB)*

November 2010

Abstract

We study the effects of macroeconomic shocks on different cohorts in the Dutch economy. From a calibrated stochastic model of macroeconomic risks, we derive typical shocks to productivity, demography and asset returns. The effects of these shocks are then simulated using an overlapping-generations model that contains a detailed specification of taxes, premiums, and benefits under Dutch law. We look at both the direct impact of shocks on agents wealth and at the redistribution and insurance that are carried out by the government and the pension system. While both these entities generally act to insure shocks across cohorts, our results show that the insurance role of the government is much larger than that of the occupational pension funds. We further find that there is little cross-correlation between different risks, except in the case of rare disasters.

*l.j.h.bettendorf@cpb.nl Thanks are due to Lans Bovenberg, Peter Broer, Hans Fehr, Albert van der Horst, Bas ter Weel and Ed Westerhout for their detailed comments on an earlier version of this paper.

1 Introduction

Macroeconomic shocks influence the entire economy, but their effects may be different from person to person. Unexpected changes in a rate of return, for instance, have an immediate impact on those who are long or short in the asset. Then there may be secondary effects that are felt throughout the rest of the economy, when those initially affected change their behavior. But the final impact of any shock is not just the result of its propagation through the market economy. The result is modified by the influence of the government, who may wish to redistribute the effects, and by different parts of the pension system. For instance, a young worker without financial assets may still be affected by a financial shock through a change in his pension premium payments. Or a slowdown in the birth rate may affect the income that retirees receive from their pay-as-you-go pension.

This paper aims to chart the main macroeconomic risks for different cohorts in the Netherlands. We look at both the direct impact of different shocks and at the redistribution and insurance that are offered by the government and the pension system. To do this, we examine the effects of unexpected changes in productivity, demographics and financial returns in a calibrated overlapping-generations model of the country. We derive likely shocks from a calibrated stochastic model of macroeconomic risks—one scenario, for instance, is that productivity grows unexpectedly, and then converges back to its original path—and feed the time series of these changes into the OLG model. From the new equilibrium of that model, we can derive the effects on the balances of different generations: how has their labor income changed, the net transfers from the government and the pension funds, their financial wealth?

Our tools enable us to chart the change in wealth for different generations, as it results from the kinds of macroeconomic shocks that we can reasonably expect. It is possible to describe this as a form of generational accounting (see for instance Auerbach et al. 1999), which concerns fiscal sustainability and the intergenerational distribution of wealth by a country's institutions. For the Netherlands, a recent exercise is in Van der Horst et al. (2010). It finds that current arrangements are not sustainable and that there is a gap between the actual projected deficits and the maximum sustainable path that amounts to 4.5% of GDP.

In this paper however, we are not interested in levels of the accounts but rather in their dynamics. Instead of inquiring whether there are any sustainability problems, we assume that the sustainability problem has been solved by a one-time permanent decrease of government material consumption. This gives us a scenario in which government debt (as a percentage of GDP) ultimately stabilizes, which rules out a Ponzi-game. We study the effect of different macroeconomic shocks starting from this base path, and report changes in (lifetime) wealth relative to those in the initial scenario.

Using this approach, we find that both the government and the occupational pension funds act as insurers for most macroeconomic shocks. Insofar as we can summarize their actions across many different cohorts, the general rule seems to be that they transfer away from those who gain from the shocks, toward those who lose. When workers become more productive, for instance, higher taxes and pension premiums dampen their gains while the (initially unaffected) retirees see an increase in their benefits. It is to be noted, though, that the insurance role of the government is much larger than that of the pension funds: often, the size of the government's transfers is several time that of the pension funds. Among the shocks studied in this paper, there are some cases in which we note a distinct lack of insurance, or even perverse redistribution. This usually involves future, unborn, generations. We will indicate these cases in the text.

While we are not able to give an analysis of the welfare effects of different shocks, we can compute what would happen if each generation would suffer an equal (percentage) change in their consumption levels. This is a well-known benchmark case that, under assumptions, corresponds to optimal risk sharing (see Bohn 2006). We find that actual consumption changes are usually far away from this ideal, with most of the insurance taking place between the generations that are alive at the time of the shock.

Looking at the kinds of macroeconomic risks that are most relevant, we further find that there is little cross-correlation between different risk classes. That is, a shock in one area (*e.g.*, demography) has only small effects in other areas (*e.g.*, productivity). Only in the case of "rare disasters" do we find that several shocks happen at once. In that case, the insurance mechanism seems to protect current cohorts at the expense of future generations.

Our tools are not perfect, though, and we must compromise on several issues: the OLG model is deterministic with perfect foresight, and thus is ill-suited to analyze an environment in which shocks happen every period. The exclusion of government consumption and public health care from the utility function implies that welfare analysis is problematic. This makes it impossible to do normative analysis on the distribution of the shocks.

Aside from the obvious link with generational accounting, this work also relates to previous studies of the distribution and insurance of macroeconomic shocks. The work of Bohn (2006, 2009) is the most prominent example of this literature. Bohn uses a standard Diamond model of two overlapping generations to assess how different shocks should be shared among the young and the old. He derives the benchmark for risk sharing (an equal response of all cohorts that is proportional to their level of consumption) which represents the unique efficient way to distribute shocks if the degree of relative risk aversion is constant over life.

Bohn (2006) finds that market solutions leave the young overexposed to productivity risk. He then uses a stylized representation of US fiscal institutions to find that American fiscal policy actually magnifies the generational gap between working-age and retiree exposure to this risk. Demographic shocks (fertility and longevity) are shared through fiscal institutions, but to a degree that is too small compared to the benchmark.

Bohn's simpler framework allows an analysis of welfare, but only gives a coarse characterization of the risk-sharing institutions. Our analysis of the Dutch fiscal institutions and pension funds does *not* find that fiscal rules put extra productivity risk on the young. In contrast to the case studied by Bohn, government benefits to the old (pensions, and in-kind medical benefits) are indexed to wages and thus vary with productivity. This puts part of the risk associated with wages on cohorts that have already retired.

This result fits nicely with the idea that the government and pension funds should be sharing risks across generations, exposing retirees to wage risks and young workers to capital risks. Others have quantified the importance of this role in the case of pension funds: Cui et al. (2006), for instance, focus on the role of pension funds in insuring investment risk and inflation risk, and present welfare gains on the order of 1-4% from a defined benefit scheme. While our model does not allow explicit welfare comparisons, we do look at a much larger array of risks and include the government as a second party that provides insurance. In our results, it turns out that the government performs a much larger risk-sharing role than pension funds.

There are several reasons, however, to include pension funds as a special kind of redistribution authority in this paper. Their size is one of them: with assets in excess of Dutch GDP a large part of private wealth is stored in occupational pension funds, and it matters how these funds distribute it. Secondly, workers cannot opt out of the fund and premiums are usually not actuarially fair. The influence on behavior and wealth distribution is therefore non-negligible. Finally, with their large stake in the financial markets and the labor market, it can be argued that occupational pension funds are one of the most exposed parties when it comes to macroeconomic shocks. It is their role in redistribution that will be of special interest.

When a macroeconomic shock is applied to the baseline scenario, we will need to model how the government sector and the occupational pension funds will deal with the effects of the shock to ensure sustainability once again. For pension funds, it is not hard to find such a policy rule. With their constant exposure to shocks and as the subject of financial supervision, pension funds often explicitly publish their resolution mechanism for solvability problems. The rule in our model is discussed in section 2.2.

For the government, such explicit resolution mechanisms do not exist. Decisions are taken in the political arena and do not follow fixed rules. We assume that initial sustainability is achieved through a reduction of general spending, but this might well turn out different in practice. Furtheron in this study (section 4.2), we inspect several mechanisms by which to restore sustainability, each with a different effect on the redistribution characteristics of the government.

The next section discusses the tools of our analysis. It looks at the model for macroeconomic shocks, the OLG model of the Dutch economy and the base path. We map the effects of simple, isolated, shocks in section 3 and look at mixed shocks, where several (correlated) risks occur together, in section 4. Section 5 concludes.

2 The shocks, the model, and the base path

If we are to gain insight into the initial impact and redistribution of typical macroeconomic shocks absorbed by the Dutch economy, we need to know more about the shocks, and we need a model of the economy that includes the redistributing agents. For this, we turn to existing literature. Macroeconomic shocks are the subject of a literature review by Broer (2010), which includes a VAR model of the four main types of shocks. Our model of the economy comes from Draper and Armstrong (2007).

In this section we discuss these background works and set the stage for our analysis. We start below with the study of macroeconomic risks. A description of the economic model that we use to assess the shocks' effects is in section 2.2. The base path of the model is discussed in section 2.3.

2.1 Macroeconomic risks

The defining characteristic of macroeconomic risks is that they cannot be diversified within the period in which they occur. For a *series* of shocks some diversification over time is possible, but any effort to do so is associated with discontinuity risk (See, for instance, Gollier 2008). Because of their exposure to macroeconomic shocks, pension funds and their supervisors share an interest in the characteristics of these risks.

Our main source of information on the size and likelihood of different macroeconomic shocks is Broer (2010). In this literature survey, four types of shocks are discerned: demographic, financial, shocks to productivity and rare (but influential) disasters. In each of these areas, a large literature is available from which estimates of the underlying process can be obtained. Importantly, the four types of shocks all have effects on the same economy and it is natural that a shock in one area should have effects in another. Indeed, Broer finds that there are many spillovers, for instance from demography to productivity (as workers age, their productivity and their ability to innovate decline) and from productivity to financial markets.

It is these spillovers that are potentially very important for parties that are exposed to multiple risks. Pension funds are vulnerable to low financial returns, but also to higher wages. Their worries would be intensified if the two tend to go together. In most cases however, Broer finds that these spillovers are not very large. The one exception is the case of "rare disasters", a field of analysis made popular by Barro (2006). With a small but positive probability, economies can experience both a large loss in productivity and severly negative returns on risky assets. Recovery from these disasters can take a long time, and economies may not reattain the old path.

Broer summarizes his findings in a VAR model of the four main processes. Most of it is taken from the literature: it combines the financial model of Campbell and Viceira (2005) with the mortality model of Hári et al. (2008) and the disaster characteristics of Barro (2006). Fertility and productivity are estimated using a

Kalman filter and cross-correlation coefficients are taken from the literature. While the large number of demographic variables leads to a high-dimensional model (the demographic process uses 100 cohorts), most of the connections between demography and the rest of the economy go through the dependency ratio. This somewhat reduces the size of the ‘core’ model. Nonetheless, to get some insight into the joint distribution of all shocks a simulation study is necessary.

Due to the (auto-) correlation between the different variables, a single impulse in one of them can set off a response across many different dimensions. We will refer to an impulse in one variable, the resulting dynamics and the associated movements in other variables as a “mixed shock”. It is the kind of event that we see as typical and likely to occur in the future. Our analysis starts with “simple shocks,” in which we look at the dynamics of a single variable only, and the response of other variables is turned off. For instance, a simple shock studies the dynamic pattern of productivity after an impulse without changing the rates of return on the financial markets. The distribution of simple shocks is the subject of section 3.

2.2 Gamma

We simulate the shocks with the deterministic general equilibrium model Gamma. A detailed description of Gamma can be found in Draper and Armstrong (2007). The model describes a small open economy that is calibrated for the Netherlands. It considers a homogenous, tradable, good and a single asset. The price of the good and the interest rate are determined on world markets. The model distinguishes five sectors: households, firms, the government, a pension fund and the foreign sector.

Households are disaggregated into 99 different cohorts. Every generation is represented by one representative household. A household chooses consumption of goods and leisure by maximizing lifetime utility, subject to the lifetime budget constraint. The death rate increases with age. Households insure against the uncertain lifetime income by buying annuities (Yaari 1965). An individual starts working at the age of 20 years. Labor supply depends on net wages and pension accrual. The net wage is determined by the gross wage minus taxes and pension premiums. The gross wage is age-specific due to an exogenous productivity profile. Pension accrual is calculated as the discounted value of the increase of the benefits from the funded pension scheme. The pension system distorts labor supply incentives only if pension premiums are not actuarially fair. Labor supply does not depend on wealth. We do not consider bequests, implying that an individual is born without any financial wealth.

The representative firm uses labor and capital to produce the tradable good. The rate of labor saving technical progress is fixed at 1.7% on the base path. The demand for inputs is derived from maximizing the value of the firm. Since investment is subject to adjustment costs, the capital stock responds with delay to changes in employment.

The government sector is modelled in detail and captures a generational accounting framework, whereby publicly provided benefits and taxes are attributed to the different age cohorts. Government revenues include taxes on labor income, pensions, capital income and consumption. In reality, progressive taxes (including social premiums) are levied on labor income, transfers and pensions. In the model the progressive system is approximated by a proportional tax, with the specification of deductible pension premiums and of a fixed tax credit. With the modeling of constant marginal income taxes, risks are less shared than with existing institutions. Capital income is not progressively taxed but a flat tax is imposed on the value of financial wealth (above a threshold). In the model capital taxes are proportional to financial wealth. Public expenditures consist of consumption, several transfer schemes, social insurance schemes, education, health care spending and interest. Age profiles are specified for transfers in cash and in kind, including the public PAYG pensions. The public pension benefit is paid starting at the age of 65 and is independent of the labor history and income. The benefit is linked to the aggregate net wage. Since a single retired person gets a higher benefit than a person who still lives with a partner, the average public pension rises with age.

The long-term projections for public spending and taxes are made under the assumption that age-specific expenditures are indexed to wages; non age-related public expenditures are indexed to GDP and tax rates are held constant (see van der Horst et al. (2010)). The intertemporal budget constraint of the government is closed by a permanent adjustment of the material government consumption. Since material consumption is distributed equally over all living individuals, this implies that the net government benefit changes by the same amount for every individual in every year. Achieving sustainable public finances on the basepath requires a reduction of government consumption by 4.5% as from 2015.

Supplementary occupational pensions are provided by pension funds. These are based on a funded system. Pension premiums are compulsory for individual workers and tax deductible. Pension benefits are taxed upon realization and indexed to the average wage earned over working life. The replacement ratio depends on the labor history of an individual. Pension rights and benefits are indexed to a combination of wages and prices. Benefits are paid from the age of 65, irrespective when the person retires from the labor market. The assets of the pension fund should cover in each year the projected obligations, or the funding ratio, defined as the ratio between asset holdings and nominal liabilities (i.e. without indexation), should be larger than one. In case of an initial deficit, the pension fund temporarily raises the premium rate and restricts indexation. We limit the change in the (catching-up) premium to 1% per year. When the funding ratio is smaller than 1.35, less than full indexation is applied. In the long run, the premium rate converges to the actuarial fair level and the funding rate equals 1.45 (allowing for full indexation).

In the small open economy, excess supply of goods can be exported at the given world market price. Similarly, a savings surplus is invested on the world capital market at a fixed interest rate.

The model is calibrated on Dutch national accounts of 2008. The demographic projections are provided by Statistics Netherlands. The calibration is fully discussed in Draper et al. (2010). Some of the key parameters are in table 1.

Wage elasticity labour supply	0.14
Intertemporal substitution elasticity	0.50
Substitution capital - labour	0.50
Real rate of return	0.03
Technological progress	0.017

Table 1: Key parameters of the Gamma OLG model.

2.3 Net benefit on the base path

Generational accounting exercises usually take place in a world in which there is only (micro-) uncertainty surrounding each person's time of death. A Yaari (1965)-arrangement provides perfect insurance against this uncertainty. There is no other risk factor; in the financial sector, only a single asset exists whose rate of return has a predictable path. Agents thus are able to optimize their consumption in each current and future period, bound only by the intertemporal budget constraint

$$\mathrm{PV}(\{c_s\}_{s=t,\dots,t+T}) = \mathrm{PV}(\{w_s \cdot L_s^{\mathrm{s}}\}_{s=t,\dots,t+T}) + \\ \mathrm{PV}(\{b_s^G\}_{s=t,\dots,t+T}) + \mathrm{PV}(\{b_s^P\}_{s=t,\dots,t+T}) + W_t.$$

Here, the present value operator is defined as

$$\mathrm{PV}(\{x_s\}_{s=t,\dots,t+T}) = \sum_{s=t}^{t+T} \frac{x_s \Lambda_{s,t}}{\prod_{k=t}^{s}(1+r_k)}. \quad (1)$$

The (known) interest rate is r, and the cumulative probability of survival from period t to period s is $\Lambda_{s,t}$. The intertemporal budget constraint says that the present value of consumption c must equal the present value of gross labor income (with observed labor supply, call this human wealth), plus the present value of net government transfers b^G (which may be negative in years in which the agent is a net payer), plus the present value of future pension benefits b^P (where b^P is negative in the years that premiums are paid), plus current financial wealth W.[1] The discounting with the probability of survival indicates the presence of the Yaari (1965)-type reverse life insurance scheme.

Figure 1 maps out the elements of the right-hand side of this equation on the base path. On the horizontal axis is the year of birth, and plotted are the values

[1]Notice that the net government transfers (b^G) at the right-hand side include public consumption, in particular health care expenditures. Therefore, the left-hand side (c) should be interpreted as total consumption. The agents can only choose the level of private consumption. Marginal utility of private consumption is independent of the level of public consumption.

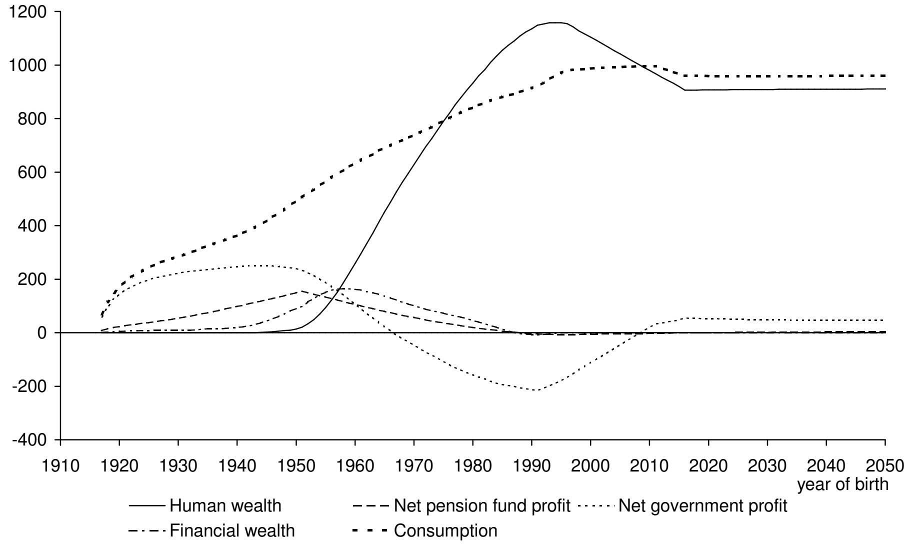

Figure 1: Discounted, expected, lifetime net benefit (thousands of euros, 2016). Income comes from four sources: labor income, pension fund and government net benefit and private wealth. These four sources add up to expected consumption.

of each of the four components in 2016. For those generations that are yet unborn in 2016, the value of their lifetime resources at the time of birth is discounted back to the year 2016 using the rate of technological progress. This has the attractive feature that all unborn generations have the same amount of human wealth ("expressed in 2016 euros") so that we can easily check whether the economy has arrived at a steady state. Discounting with the rate of interest would bring all these lines down to zero asymptotically.

Figure 1 shows how the importance of the different wealth components differs between ages as the remaining part of the life cycle varies. Discounted labor income is zero for the retired and falls throughout the lifetime for the active part of the population. Those who have not yet entered the labor market see their human wealth increase at the rate of interest as their entry comes near. Expected pension fund benefits increase over the working life as pension rights are being accumulated; premiums paid in the past do not show up in this number. For those who have entered retirement, the value of the pension "asset" drops as the expected remaining lifetime goes down. For the unborn, the actuarially fair pension system with which they have not yet interacted does not represent a net gain or loss.

Remaining net benefits from the government are positive for the retired and the older workers: they expect to receive more than the amount they have left to pay (through taxes on labor, capital etcetera). Note that the expected value of net government benefits declines very slowly for those over 65. Most benefits at this stage consist of health care expenditures, which increase with age and are paid for with public funds. The amounts in the graph are per person and conditional on staying alive. It is not true, in general, that the 90-plus-cohorts claim a large part of the government budget.

Young people who have already consumed government services (such as education) in their youth have a negative expected remaining benefit. For the unborn, the expected benefit from the government is positive as they are given ownership of part of the government assets (which, at this time, still exceed the government's liabilities). Financial wealth, finally, is accumulated throughout the working life and decumulated in retirement. As required by the budget constraint, the sum of these four components adds up to the expected value of total lifetime consumption, which completes the graph.

Using graphs that derive from figure 1 we will chart the initial effect of macroeconomic shocks and the subsequent redistribution of this effect by the government and the pension funds. We do this for each wealth component, plotting the change expressed as a percentage of baseline consumption. The four wealth components add up to the net effect that a cohort experiences, which is its (percentage) change in lifetime consumption. One might be tempted to ask how much redistribution is in fact optimal, but it is clear that we cannot answer this question without a social welfare function to weigh the net benefits of different generations. We do not, in fact, make statements about welfare but restrict ourselves to statements about income rather than full consumption.

For comparisons, however, we may use the benchmark case of Bohn (2006). In this case, where agents make optimal saving and portfolio decisions over their lifetime, the relative effect of any shock on yearly consumption should be equal for all generations.[2] If part of the role of government and pension institutions is to help agents in achieving this optimal response, then we can measure the degree to which they succeed by judging whether the distribution of shocks over generations is close to uniform. If that is not the case, we can ask ourselves whether the government and pension institutions at least change the distribution towards uniformity, rather than away from it.

3 Initial impact, insurance and redistribution

To understand the response to shocks by agents, pension funds and the government, we initially reduce the shocks' complexity. With many variables moving at once, it becomes harder to pinpoint cause and effect in the OLG model and we therefore start our analysis with four "simple shocks." In these cases, we isolate the sub-VAR model for productivity (or demography, or financial returns) and look at a typical impulse-response pattern in just that submodel, leaving out correlations with other variables. We can more easily analyze the effect of these "simple shocks" in the OLG model. Descriptions of the effects of these shocks should give a feel for

[2]Note that this result is derived under the assumption of homothetic utility. If there is habit formation, for instance, it does not hold.

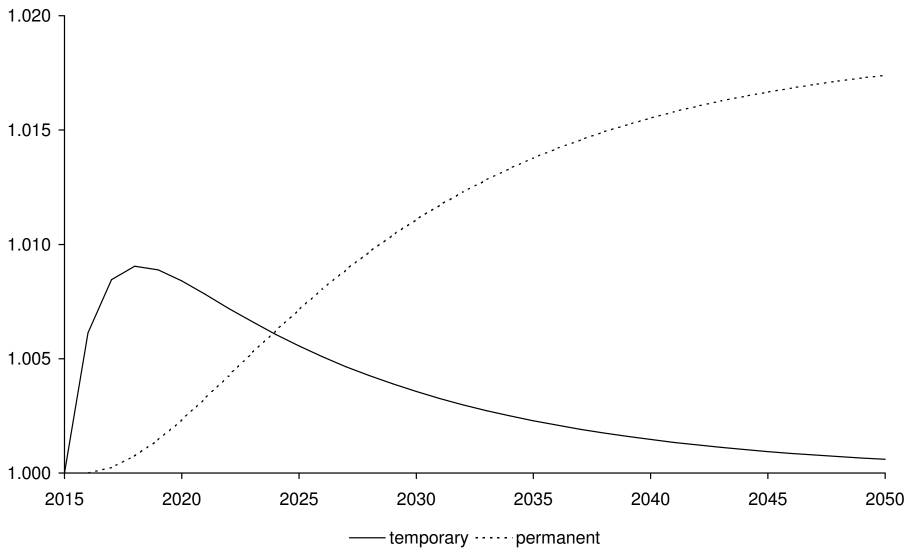

Figure 2: Two shocks in the rate of productivity, relative to the base path, administered in the year 2016. With the first shock, the level of productivity eventually drifts back to the base path. The second shock occurs to the latent process that has a unit root. Spillover effects cause the Dutch level of productivity to converge to the new, higher, path where it is permanently 1.87 percent higher than before.

the model, without overly going into detail. All shocks discussed below arrive unexpectedly in the year 2016. If the shock also implies changes after this year, these are foreseen by the agents after they learn of the initial shock in 2016.

3.1 A shock in productivity

If productivity is hit by a shock, does it stay permanently altered? This question (which was asked about GNP and answered in the affirmative by Nelson and Plosser 1982) is discussed at length in Broer (2010, section 2.1.1). At present, the debate on whether GDP has a unit root still continues. We sidestep the controversy by looking at two types of productivity shocks, one of which is permanent and one of which is transitory. Productivity is modelled as a (latent) unit root process with which actual productivity is cointegrated. We apply a one-standard-deviation shock to the productivity process to get a transitory effect, and a one-standard-deviation shock to the latent process to get a permanently higher productivity (in that case, 1.87 percent). The graphs of the two productivity paths are in figure 2.

The initial impact of a shock on productivity is on wages and profits. Since we look at simple shocks in this section, the effect of profits on financial returns has been turned off and the only initial impact is on wages. For the temporary shock, this means that labor income is higher than on the base path for the cohorts that are on the labor market while the shock lasts. For the chosen size and date of the shock, figure 3 shows that generations born between 1945 and 2030 are affected directly. Those with the highest participation rate and age-specific productivity see the largest effect: these are the generations born around 1970 with a gain of around 100 euros per year. This corresponds to 0.4% of lifetime consumption and is the result of a shock in which productivity increases just shy of one percent, relative to the base path, for a number of years (see figure 2).

The figure also shows how this initial effect is redistributed by the government and the pension funds. The government immediately takes part of the wage increase away through income taxes, and again some through the consumption tax. This accounts for the decline in net government benefits for the cohorts that see their wage income increase. A second effect of the increase in wages is that wage-indexed transfers go up. This benefits all living cohorts, including the elderly who receive first-pillar pensions. This group also sees the value of the health care it gets increase. For the retired, this means that the net benefit from the government goes up. In all, the sustainability of the government budget worsens and public expenditures are decreased somewhat as a result. This shows up as a negative effect on net government benefits for the unborn.

The pension sector also increases benefits to existing retirees, as a result of wage indexing. The decline in solvability caused by this is reversed by increasing premiums. This amounts to a transfer from the working-age to the retired, albeit a relatively small one as the dashed line in figure 3 shows.

What figure 3 clearly shows is that both the government and the pension sector actively redistribute the effects of the productivity shock away from the working-age who are the initial recipients, towards the initially unaffected retirees. It is the government sector that does the bulk of this redistribution, by disbursing higher pensions and paying for the more expensive health care. The pension system redistributes to slightly younger cohorts, but to a much smaller degree. The total effect of the temporary shock in productivity can be seen from the dashed line in figure 3 that shows the percentage change in consumption. It shows that on an individual level, the main recipients of the shock are the elderly, who feel its effects only through redistribution. A temporary increase in productivity actually makes the unborn generations worse off, as they have to contribute to finance the government's worsened solvability. This is one of the cases, announced in the introduction, where redistributions seems to be perverse, as temporary gain in capacity leads to losses for a number of cohorts.

It seems odd that a temporary higher productivity should lead to a loss for future generations. This is the result of the current institutional setup: the wage increase leads to higher government expenditures, which are not reclaimed by the increase in the tax intake. In principle, however, it should be possible to use the windfall gain of higher productivity to make everybody better off. The dashed, bold, horizontal line in figure 3 shows by how much each generation could have increased their consumption if the shock had been distributed evenly. We compute it by taking the present value of all consumption changes, using the interest rate and mortality after the shock, and dividing evenly over all generations. Relative to

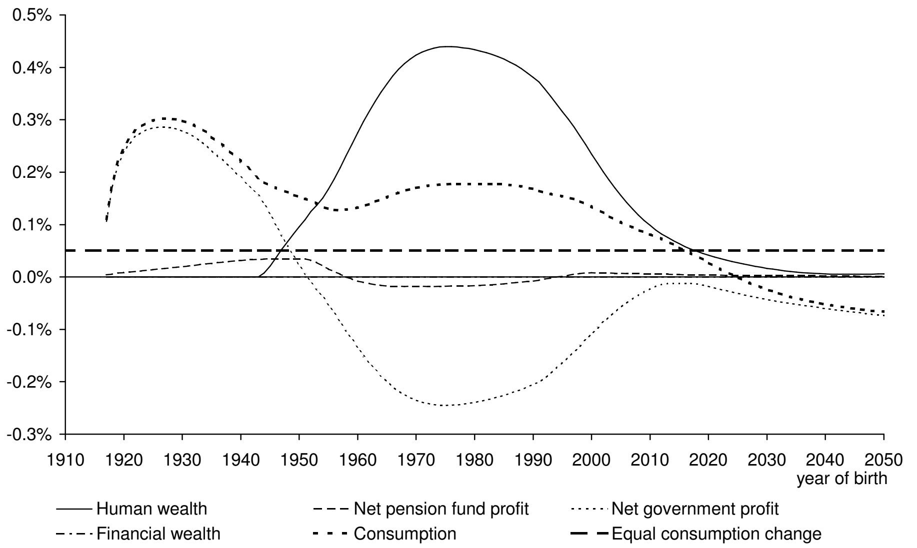

Figure 3: Discounted, expected, lifetime net benefit change from a temporary shock on productivity, expressed as a percentage of consumption on the base path. The change is split into human wealth, pension funds, and net government profit. Financial wealth is not affected. If each generation responded with the same percentage change in consumption, that level would be at the line *Equal consumption change*.

this benchmark of 0.05%, we see that all generations alive at the time of the shock get more than their share.

The finding that the government redistributes productivity risk to the retired is different from the main result in Bohn (2006), who finds that governments provide too much certainty to the retired vis-à-vis the working cohorts. The reason is that value of government transfers to the old in our model is uncertain, rather than certain, and correlates with the shock in productivity. This makes the retirees party to developments in the labor market, which they are not in the Bohn model. Pension funds play the same role on a smaller scale.

In case of a permanent shock to productivity, the change happens gradually and amounts to a 1.75 percent higher productivity level after 35 years (see figure 2). This increases the labor income of all cohorts that are active in the years after the shock, but the effect is skewed towards younger generations. Figure 4 shows that those born around the time of the shock can increase their consumption with about 1.75% (this is about 400 euros yearly). In this case too, the only direct effect of the increase in productivity is a higher wage rate. Part of this extra income is again redistributed by the government and the pension system, using the same mechanisms as with the temporary shock. The redistribution is, however, much less pronounced. With the temporary shock in productivity, transfers to retirees increase as the wage rate goes up. As the shock fades out, though, transfers decline again. This leads to a number of cohorts that have to pay higher taxes during their working age, but do not get extra benefits in retirement.

In contrast with this, with the permanent shock the working-age cohorts that see their income- and consumption taxes increase also get higher benefits in retirement, as the effects of the shock do not wear off. This is most easily seen by the gradual increase of the change in consumption in figure 4, which has a time profile not unlike that of the shock in productivity. In contrast to the temporary shock, the net government benefit to the unborn generations is not worse after a permanent shock. Note further that the same generations gain and lose from the redistribution of the government and the pension sector, and that the size of the redistribution is again much larger for the government.

Looking at the benchmark of equal distribution of the shock's effects, we see that this time the generations alive in 2016 get *less* than their share. Redistribution gives some cohorts of retirees an increase of 0.5%, but the uniform value of the shock is three times as much.

3.2 A wealth shock

We next turn to a simple financial shock, in which asset values decline unexpectedly while other exogenous variables stay on the base path. Again, this oversimplifies an actual shock by neglecting the correlations with productivity and future returns. We will turn to realistic shocks later and use this simple shock to understand the workings of the model.

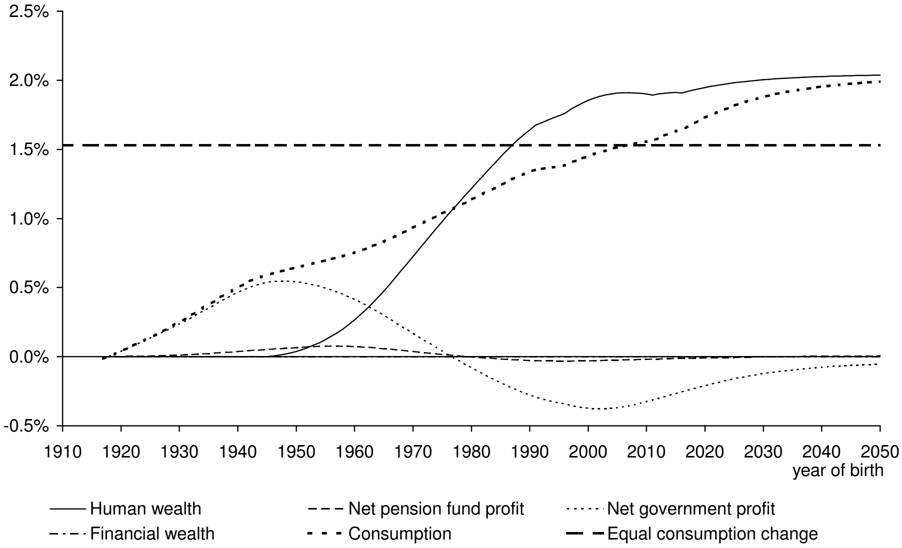

Figure 4: Discounted, expected, lifetime net benefit change from an enduring shock on productivity, expressed in percentage change in consumption. The change is split into human wealth, pension funds, government and financial wealth.

The OLG model works with the assumption of perfect foresight, and as such only has one financial asset. We aggregate the different real-world assets into this measure by assuming a fixed portfolio which hold 25% shares, 25% commercial bonds and 50% in the safe asset. In a disaster, the value of shares can drop anywhere between zero and 70 percent. Taking the expected value and using the portfolio weight of shares, we model a value decline in the asset of the OLG model as a one-time drop of 8.75%. This shock occurs without affecting the rate of interest, which stays constant at the steady state level both before and after the shock.

Figure 5 shows the effects of the drop on the cohorts' wealth components. The affected financial assets are held, in the model, by two parties: individual agents and pension funds, who both see their wealth shrink. For individual agents, the size of the loss is proportional to the amount of assets held, which can be seen in figure 1. The loss is felt heaviest by those generations near retirement when the shock occurs, and tapers off toward younger and older agents. This pattern can be observed in figure 5. Note that assets are slightly negative for the youngest generation at the time of the shock; they profit as the value of their debt falls. The effect through private wealth holdings is the largest of the four mechanisms.

Pension funds respond to their deteriorated solvability by decreasing benefit indexation and raising premiums. This primarily affects the same cohorts that were also the largest holders of private wealth, although the effects are a bit more spread out towards the older retirees and include those who do not hold assets at the time of the drop in value. Note that the pension system seems to exacerbate the initial effects of the shock, rather than redistributing them. Without a pension system,

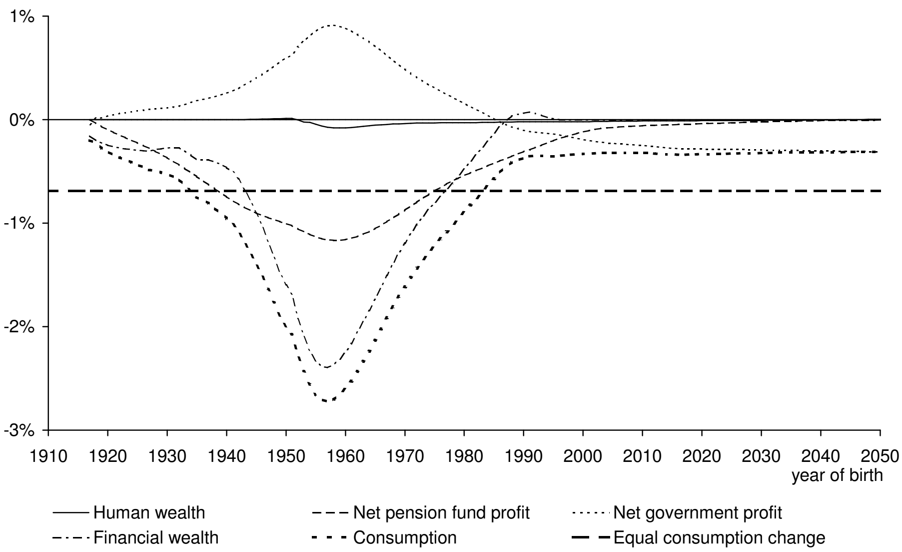

Figure 5: Discounted, expected, lifetime net benefit change from a one-time decrease in asset values of 8.75%, expressed in percentage consumption change. The change is split into human wealth, pension funds and government net benefit and financial wealth.

however, it is likely that agents would have held even larger personal balances. The much more spread-out profile of the change in pension system benefits indicates that some intergenerational risk sharing is indeed being provided. This result is in line with that of Bonenkamp et al. (2009), who compute the effects per cohort (through the pension system) of the decline in wealth that was the result of the 2008 drop in equity prices.

The government does redistribute in the case of this shock. Taxes on the affected cohorts drop, as do benefit levels for all generations when the government restores solvability after this loss in revenue. Because the latter also includes the unborn, this amounts to a transfer from those too young to suffer from the financial shock to those who did. In size, the redistribution of the government approximately cancels the effect of the pension loss for the most affected cohorts. Younger generations are given a share of the burden by both the government and the pension system. Compared to the benchmark of equal consumption changes, however, we see that the share of future generations under the current rules is too small.

Human wealth, finally, is virtually unaffected by this isolated financial shock since we have purposefully eliminated any real effects on, for instance, productivity. The small drop for the working generations is an indirect effect of higher (distortive) pension premiums.

An aspect of this shock that keeps things simple is that the change in valuation does not involve changes in the rate of interest, which is used to discount the future streams of revenue and consumption. Such a change is the subject of the next

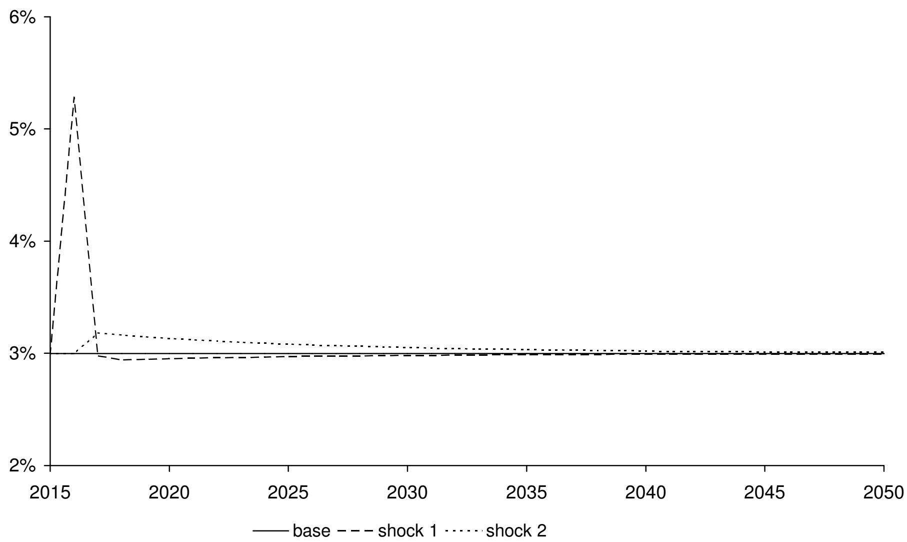

Figure 6: Two shocks in the (real) interest rate, administered in the year 2016. A spike, followed by a period of sub-normal rates and a long drift upwards.

section.

3.3 A shock in interest rates

We model financial shocks using the 5-asset VAR model of Campbell and Viceira (2005). In this model, there are five independent sources of error which hit the return on one asset initially, and then propagate through the values of all five assets. To apply this rich menu of assets in a one-asset OLG model, we work with a fixed portfolio that consists of three of the five assets. Observing how the shocks to different assets lead to a return-profile of this portfolio, we find that most tend to generate a similar pattern in returns. When we restrict ourselves to shocks in the real rate of interest, we find that we can reduce the number of independent shocks to two. Keeping inflation constant at 2%, the two shocks in figure 6 cover the range of possible dynamics: one is a spike in interest rates that quickly reverses and leads to a slightly lower rate for several years, the other is a prolonged increase with a very slow return to the mean value.

The effects of a shock in interest rates in our framework of generational accounting work through two channels. The first channel is similar to the other shocks above: the change in interest rates sets off a reaction in economic activity, usually because of a change in the desired capital-labor ratio. This changes behavior and income, depending on the ownership of these factors. The second channel is the computation of net benefits itself. Note that in equation (1) the interest rate is used to discount future receipts and outlays. A change in the interest rate thus directly affects the present value of future transactions. To some extent, this seems artificial: is someone with a certain, future, income really better off if the interest rate falls, even if prices stay the same? The answer is of course yes, as the interest rate itself is a price. This person could borrow against his future income and consume today; with a lower rate of interest, the amount that he could borrow would increase.

An unexpected change in the rate of interest also affects the holders of fixed income securities through a capital-gains channel; we disregard this effect by assuming that all domestic stocks and bonds are owned by foreigners. This means that a temporary change in the rate of interest has no effect on the *financial* wealth of the agents in this model at the time of the shock.

We illustrate using the first shock in interest rates, the 'spike' in figure 6. The changes in net benefit can be found in figure 7. The most important change in the economy is a (temporary) drop in wages of about 2% in 2017, when the capital shock adjusts (with a lag of one year) to the new interest rate. The increase in the interest rate leads to a decrease in investments, which has a negative effect on wages as the marginal product of labor falls. This directly affects the value of human wealth for the working cohorts at the time of the shock. The, quantitatively more important, indirect effect comes about through the interest rate itself: the present value of future wage income falls due to a higher rate of discount. As the graph shows, young workers lose over 2% of their consumption (about 200 euros, annualized). The wage recovers after the shock, and marginally exceeds its base path for several decades. Cohorts that are not yet born when the interest rate exceeds its normal level thus see the value of their human wealth increase slightly. That the unborn cohort gains while others lose is inefficient compared to the benchmark, but the redistribution is in the right direction; the government puts part of the costs of compensating current workers on future generations.

With the lower wages income taxes fall as well, which softens the blow for working-age cohorts but deteriorates the government balance. However, the value of government benefits also goes down, both due to the lower wages of those who provide benefits in kind and due to the wage-indexed nature of many benefits. The higher interest costs on the debt tip the balance and make the government slightly worse off. As a result it has to decrease outlays across the board.

Pensioners see no direct effect of the change in wages, but their position vis-à-vis the government declines as their benefits become cheaper, and are more heavily discounted.[3] The working-age cohorts pay less income-related taxes and see the value of their future benefits go up slightly; those outside the labor market are worse off as they suffer the government's actions to restore solvability.

Pension funds use the lower wage rate to scale down their (wage-indexed) benefits. Their value of their future obligations falls, which increases solvability and causes the funds to decrease premiums and to increase the pension rights of all

[3]The 99-year olds see a much smaller effect of this shock than other, slightly younger, retirees. Under the model's assumptions, they are only alive in 2016 and escape the drop in wage-related benefits (in particular, the in-kind health care benefits).

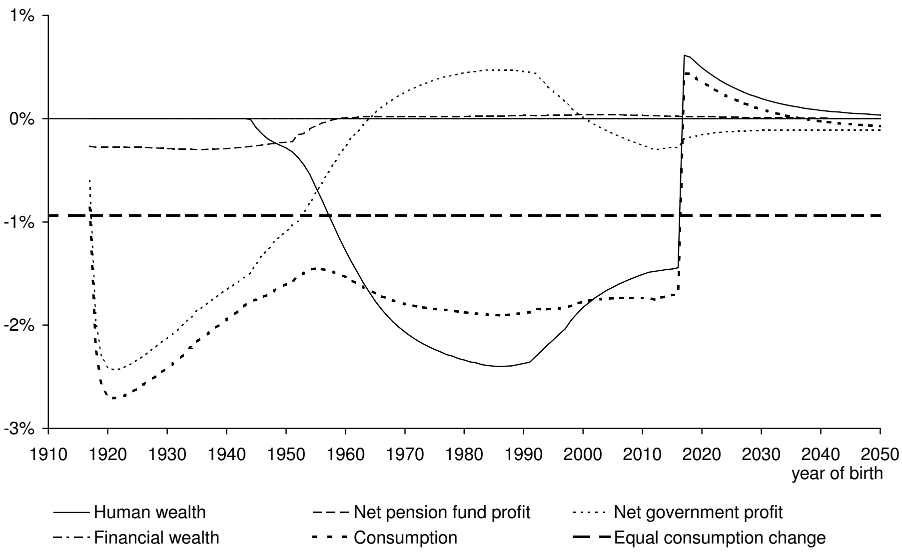

Figure 7: A spike in the interest rate in 2016 (see figure 6 above) results in these discounted, expected, net benefits changes. The changes are normalized to cohort consumption, and they compare the previous net benefits (discounted with a constant interest rate) to present net benefits (discounted with the post-shock interest rate).

participants, spreading a windfall gain over pensioners and the working-age alike. The net result is lower pensions for the retired, and slightly higher benefits for the working age.

Figure 7 shows that the spike in interest rates is bad for most cohorts. Those that rely on wage (-related) income or future cash flows are made worse off. Higher interest rates also cause a windfall loss for the owners of capital. In the present model however, the parties that bear the risk on capital are assumed to be foreign and so these negative effects do not show up in the analysis.

The alternative shock in interest rates is a modest increase (about 18 basis points maximum) but one that stays positive over a sustained period of several decades. Compound interest being what it is, the total effect can be quite impressive. The effects on net benefits are in figure 8. The main effect of this shock on the economy is that the wage rate drops between 0 and 1 percent for a long period of time.

Figure 8 shows that pension funds reduce current pensions and transfer the gains from the increased solvability to the working-age cohorts. Human wealth again declines, due to both lower wages and more discounting. With the sustained higher rate of interest, unborn generations also lose from this shock, contrary to the earlier spike in interest rates. The interactions with the government are dominated by the change in wages. This makes indexed financial benefits and benefits in kind worth less, and so all generations are worse off in terms of the value of their

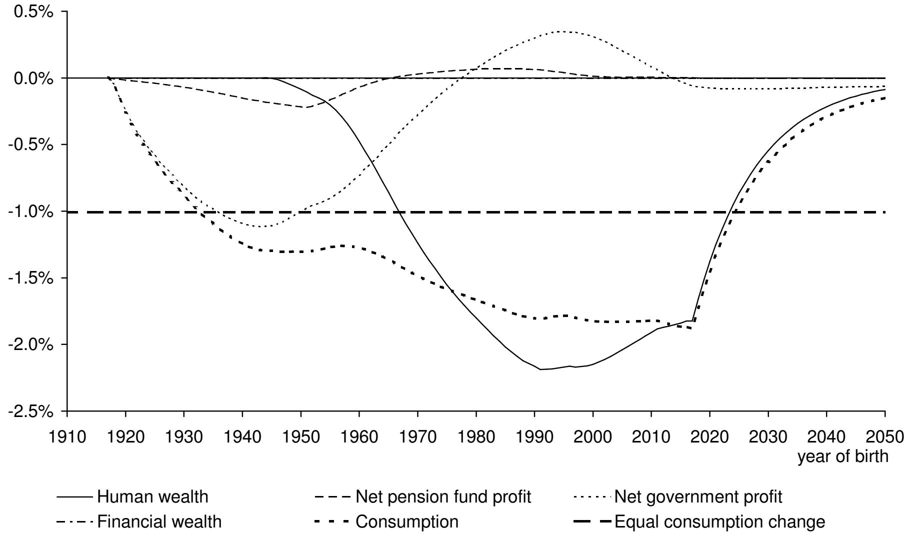

Figure 8: A general increase in interest rates starting in 2016 (see figure 6 above) results in these discounted, expected, net benefits changes. The changes are expressed as a percentage of consumption, and they compare the previous net benefits (discounted with a constant interest rate) to present net benefits (discounted with the post-shock interest rate).

benefits. Taxes, however, go down for those whose labor income is diminished and cause those born after 1977 to see an increase in expected government benefits. The overall effect on government solvability is again negative, leading to a lower expected benefit for the unborn.

Summarizing, we see that this long period of elevated interest leads to losses for all cohorts. The government plays a major redistributive role in this scenario, lowering net benefits to the retired generations and compensating the young for their loss in labor income. In the end, both shocks are bad for fiscal sustainability. This stands in contrast to a permanent increase in the interest rate that was analyzed in Van der Horst et al. (2010); it shows that temporary changes may have different sustainability effects than permanent changes. Pension funds engage in similar redistribution, but on a smaller scale. The unborn generations do not see much redistribution of the shock, and underpay compared to the benchmark of equal consumption change.

3.4 A shock in mortality

The mortality process in Broer (2010) uses the well-known framework of Lee and Carter (1992), which has a latent variable μ_t that influences the probability of death for all cohorts:

$$\log \lambda_{t,\tau} = \alpha_\tau + \beta_\tau \mu_t + \epsilon_{t,\tau}$$

where τ is cohort age and μ_t is a random walk with (downward) drift. $\lambda_{t,\tau}$ is the probability of dying during the age of τ in year t (not to be confused with the survival probability $\Lambda_{s,t}$ in formula 1). Both the constant α and the coefficient on the latent variable β are cohort-specific. A single shock to the latent process means that mortality rates across all cohorts move in the same direction (all β's have the same sign) and that the shock is permanent. The actual model by Hári et al. (2008) that is incorporated adds a twist, though, which is that errors to the latent process come from a moving average process and partly reverse themselves. How much of the error is reversed depends on the age of the cohort, with no reversion for the youngest versus two thirds for the very old. Also, the average size of the shock is larger for younger cohorts.

We study the effects of a shock in mortality whereby all cohorts experience a decrease in their mortality that is equal to the standard error of the underlying Lee-Carter process. Shocks are reversed inasfar as that is usual for the age-group. This means that two years after the shock, the probability of dying in one particular year goes down by about 10 percent for the middle-aged to about 2 percent for the elderly, relative to its previous value. The effect declines slightly over time with the standard drift of mortality rates but stays present in mortality rates for ever.

As with interest rates, the decrease in mortality does two things: firstly, it has an immediate impact on people's behavior, the balance sheet of the government and that of the pension system. Secondly, it affects the computation of generational accounts as the rate of mortality shows up in the present-value formula (1). Because mortality acts as a discount factor, lower mortality means that the value of future benefits goes up.

Figure 9 shows these principles in action. A decrease in the mortality rate increases the total amount of wages someone of working age can expect to receive, as the probability of dying before retirement goes down. We thus see a positive effect on human wealth that is proportional to the expected remaining time in the labor force. The marked increase around the year of the shock comes from the relatively high probability of dying in the year of birth..

Retirees, who get benefits conditional on not dying, see the value of their pension rights go up, while pension funds suffer a decrease in their funding ratio. Solvability is restored by charging higher premiums to the working generations, and by lowering the indexation of pension rights. The base premium rate is adjusted for the new expected lifetime and young people who enter the fund after the shock see no effect on their pension wealth.

Finally, net government benefits dominate the effects. A longer expected life means that expected net benefits go up for retirees. There is a second effect, however, which is a consequence of the way health is coupled with expected lifetime. As mortality goes down, the expected health of the old improves along with it. This does not decrease total health care costs over a person's lifetime, but it does move these costs further into the future. The increased discounting of these costs is seen as a decrease in the value of net government benefits. For most retirees, the net effect is positive but for the very old it turns into a loss. The total result

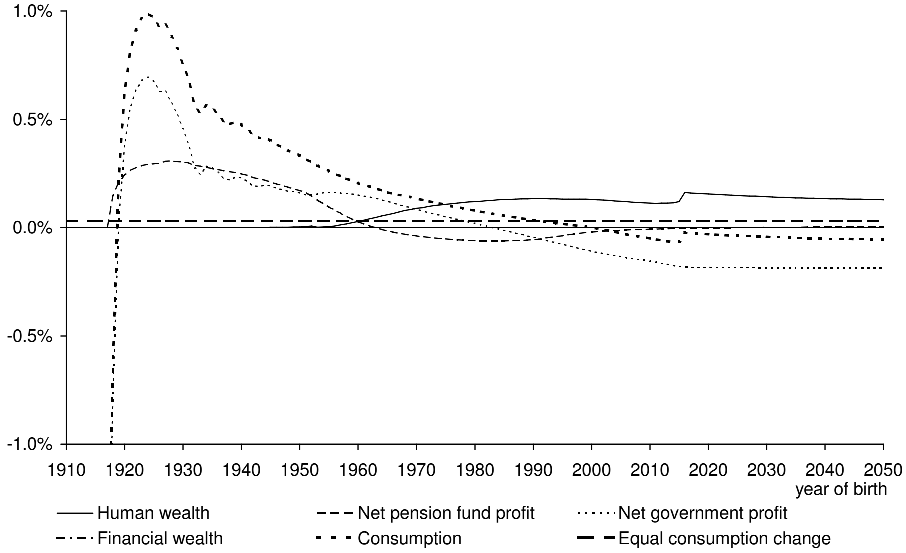

Figure 9: Discounted, expected, lifetime net benefit change from a shock in mortality. The change is split into human wealth, pension funds, and government.

of the shock on government sustainability also turns out to be negative, and future generations see about a 60 euro yearly decline in net benefits. This is a case of (mild) overshooting in the redistribution process, as the shock can in principle be used to make each generation better off.

4 Mixed shocks

The previous section showed how relatively simple shocks to one exogenous variable are absorbed and redistributed in the Dutch economy. While this is helpful in understanding the mechanisms in the model, we argued in section 2.1 that in reality, macroeconomic shocks affect several variables at the same time. In the current section, we turn our attention to these "mixed" shocks. While each originates in one part of the economy, they all spread through several domains. Our analysis of simple shocks will be of great help in understanding the effects of the mixed shocks.

In the current section, our aim is to look at shocks that are all equally likely. For almost all shocks, we start with a one-standard deviation impulse in one random variable. The only exception is the case of a rare disaster, explained below. Since the distributions of the shocks are symmetric, it is also true that the opposite shock is about as likely as the ones we discuss. The reader may find the effects of an increase in mortality, or in productivity, by negating the results in this section.

In this section, we also present the results of two different closure rules for the government. Up to now, any change in the solvability of government finances has

Variable	shock	dynamics
Productivity	-9.5%	permanent effect, arrived at over decades
Interest rate	-1.5%	immediate, recovery in 10 years
Wealth shock	-8.75%	immediate and permanent

Table 2: The parameters of the rare disaster

been countered by permanently adjusting government material consumption. In section 4.2 we close using different instruments after the 'rare disaster'-shock (see below) and inspect whether different generations shoulder the burden.

4.1 Wealth, productivity and interest: a rare disaster

Not all macroeconomic shocks follow a standard normal distribution. Barro (2006) notes that with a yearly probability between 1 and 2%, countries experience an economic disaster that results in a large drop in per-capita GDP, falling interest rates and large loss of value on financial markets. This combination of events is exactly the kind of correlation between macro risks that we are interested in, and we start with an analysis of the fallout of such a disaster. First though, a word of warning: as expected, *rare* disasters do not occur with the same frequency as the other shocks that we discuss in this paper. Broer (2010) puts the probability at 1.38% each year, which is an order of magnitude lower than the one-standard deviation shocks we discuss in other sections.

The disaster in this section combines the permanent fall in productivity of section 3.1 with the asset market drop of section 3.2. Interest rates fall by 150 basis points and stay depressed for 10 years. An overview of the parameters is in table 2. Note that the shock on productivity is given to the latent productivity variable, which makes it permanent but also introduces a lag of several years before the effect is fully felt.

The first thing to note about figure 10 is the scale of its vertical axis: the effects are much larger than those in previous diagrams. That there are many different things going on at the same time is evidenced by the behavior of the total wage sum: in the first seven years after the shock, this variable goes up as the low interest rate leads to investment that increases the marginal value of labor; after that, the productivity loss kicks in and the wage sum is smaller than on the base path. The (temporarily) higher wages and lower rate of discount explain why human wealth actually goes up after the disaster for the older cohorts on the labor market. Younger cohorts and the unborn see a decline that dominates all other effects for those born after 1985.

The wealth shock on financial markets is bad for holders of private wealth. Pension funds suffer the same wealth shock plus an increase in the value of their obligations due to the lower interest rates; the effect of the swerving wage rate can go both ways. It turns out that only the very old see an increase in their net benefits from the pension fund; the rest face a loss, whose distribution is similar to that in

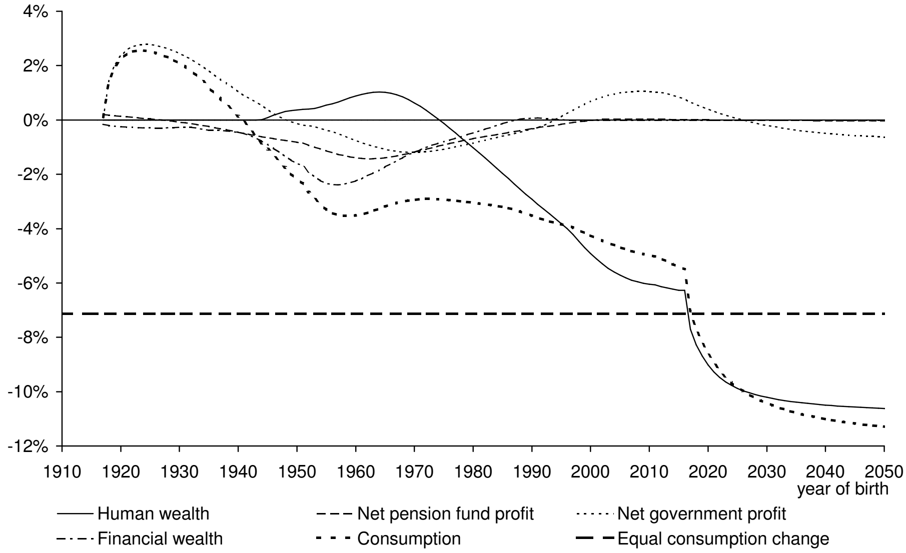

Figure 10: Discounted, expected, lifetime net benefit change from a rare disaster that has permanent effects on productivity. The changes are expressed as percentage of consumption and split into human wealth, pension funds, net government profit, and financial wealth.

figure 5: the pension fund spreads the loss over a larger number of cohorts than just the holders of financial wealth.

Those who expect positive net benefits from the government in the future see the present value of that promise increase as interest rates fall. This is true for retirees as well as the very young. When wages start dropping below base case levels, taxes on labor and consumption fall, but so does the value of benefits in kind and indexed pensions. Government interest outlays fall temporarily while the rate is below average. On balance, the effect on the solvability of the government is negative and net government profit falls for future generations. The change in net benefits switches sign several times: benefits get redistributed from current and unborn workers to retirees and children. Overall, the government appears to be exacerbating the effects of the disaster rather than insuring them. Due to the permanent loss of productivity, the unborn generations are the ones who most feel the direct impact of the disaster. This is the second case in which redistribution goes in the wrong direction, after the one noted in section 3.1.

4.2 Different government closure rules

The large effects of the rare disaster in the previous section make this shock well suited to test the effects of different government closure rules. Up to now, we have worked with the assumption that the goverment adjusts to changes in fiscal sustainability by changing its material consumption, forever, to plug the gap. While somewhat arbitrary, the advantage of this mechanism is that it does not distort prices and affects all generations equally. In practice, however, changing material consumption may not always be possible. In this case, fiscal sustainability may be brought about by changing one or more tax rates. This introduces (or, with positive shock, reduces) distortions in the economy that influence the distribution of the shock's effects.

We consider two alternative closure rules in this section: using the indirect tax on consumption and investment, and using the direct labor tax. We impose these rules on the base path as well as the after-shock solution, which means that the base path for each shock is different. This follows from the fact that creating a sustainable base path already requires the government to use its closure rule to fix the current fiscal insolvability.

To assess the effects of different closure rules on the division of the shock, consider the three lines in figure 11. They are the change in lifetime consumption, per cohort, after the 'rare disaster'-shock of section 4.1 above. The change in consumption summarizes the effects of the changes in all wealth components, not just net government profit. It is easy to see that other components may also be affected by the closure rule; for instance, changing the direct tax on labor income affects net pension benefits, whose premiums are tax deductible.

In the figure, note that using indirect taxes as an instrument leads to little changes in the way the shock is distributed over the different generations. Both workers and retirees regularly pay indirect taxes and are affected by their increase. Cohorts still in their working age appear to be affected slightly less than retirees. The similarity does not hold for the third instrument, direct taxes on wage income. Using this tax rate as an instrument of closure puts the burden of the shock squarely on the shoulders of the workers, the young, and the unborn generations. The reduction in human wealth is exacerbated by the reduction in labor supply that follows the lower net wages. Somewhat perversely, using direct taxes on labor as an instrument widens the gap between retirees, who profit from the shock, and younger generations.

Lastly, note that with all these three closure rules we do assume that the government immediately realizes the extent of the fiscal gap and takes action to change the relevant tax rate (or outlays) and put the budget back on a sustainable path. This is a distinct case from a non-forward looking rule, for instance if the government strives for year-to-year budget balance.

One reason to deviate from forward-looking rules is that they may cause large, politically infeasible, swings in the budget deficit. If we return to the situation in which government outlays are used as an instrument to guarantee long-run solvency, it follows that government debt becomes an endogenous variable. For all the shocks discussed in this paper, only three cause the debt ratio to deviate from its base path by more than 1.5 percent. These are the interest rate spike, the wealth shock and the rare disaster. The debt paths generated by the standard closing rule after these shocks are in figure 12. The wealth shock results in much higher government profit for the cohorts that live through it, financed by higher taxes in the

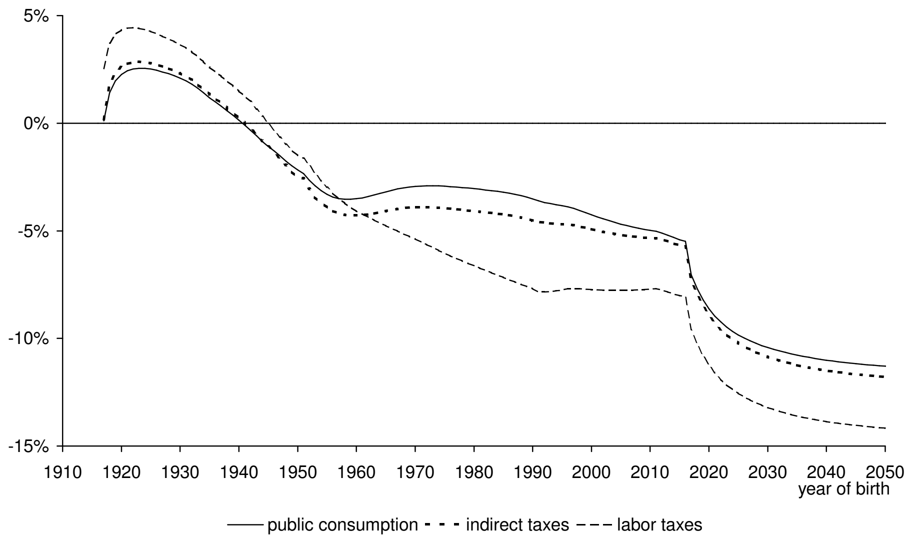

Figure 11: Discounted, expected, lifetime consumption change after a rare disaster that has permanent effects on productivity. The government closes using three different instruments: decreasing material consumption (as in figure 10), raising indirect taxes on consumption and investment or by raising direct taxes on labor income.

future. This translates into a a gradually increasing debt ratio, up to about 10 percentage points extra in 2075. This path does not seem to violate any constraints on the deficit. In contrast, the rare disaster leads to an immediate jump in the debt level by almost 25 percentage points in the year of the shock. Though not completely unthinkable (in 2008 the Dutch debt ratio jumped by 13 percentage points) this does indicate drastic action by the government. The interest rate shock, lastly, leads to temporary surpluses because of the decreased price of benefits in kind; this is probably not a problematic scenario for politicians.

4.3 Mortality and other factors

In this section, we apply the across-the-board decrease in mortality of section 3.4 to the full VAR model, taking into account cross-correlations between different macroeconomic risks. As with all demographic shocks, the path of causality to other risk factors is through the dependency ratio. In this scenario, the dependency ratio increases compared to the base path with about 0.6 percentage points in the long run.

The dependency ratio has a direct, negative effect on interest rates (Broer 2010, section 2.3.8) due to the fact that capital supply varies over the lifecycle. The idea is that older people are net capital owners and that a relative increase in their number drives down the rate of return. The effect is small: in the current scenario, the interest rate goes down by two basis points.

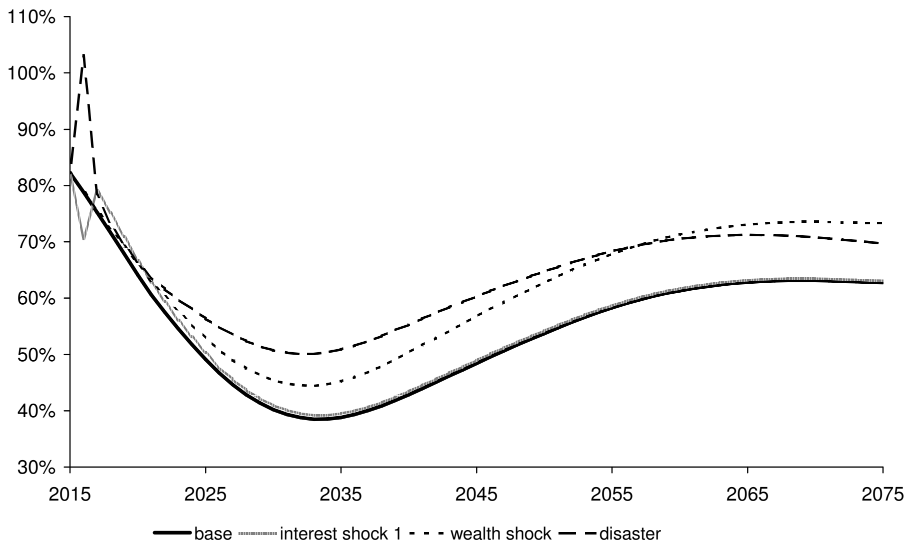

Figure 12: Government debt as a percentage of GDP, on the base path and in three scenarios with shocks: the wealth shock from section 3.2, interest shock 1 from figure 6, and the rare disaster of section 4.1.

A second direct effect of the increased dependency ratio goes through labor productivity. Broer (2010, section 2.2.2) specifies an elasticity of -0.1 between log labor productivity and the dependency ratio. This is caused by relative over-payment of older workers due to implicit contracts and a reduced ability to adapt to innovations. The effect is transitory, though: after several years of slightly lower growth (on the order of 0.1 percentage points) the process reverses and the economy returns to the old growth path.

The two cross-effects cancel each other out when it comes to wages in the first decade. After that, the productivity decrease is reversed but the lower interest rate remains, and wages increase by about 0.11%. This can be seen by inspecting the effect of this combined shock on the human wealth of unborn workers, on the right side of figure 13. Note that the left side of this figure is not very different from figure 9, where the cross-correlations are not taken into account. The effect on the consumption of unborn generations changes sign, from negative (without cross-correlations) to positive when cross-correlations are taken into account. Also note that for most cohorts, the effect on consumption is quite close to the benchmark of equal consumption change.

At the same time though, we must conclude that the added effect from letting the other macro-risks vary with a shock in mortality is negligible for most generations, especially the older, living, cohorts. This finding repeats itself across all of the other mixed shocks, which is the reason why we do not report mixed shocks from other sources in this section.

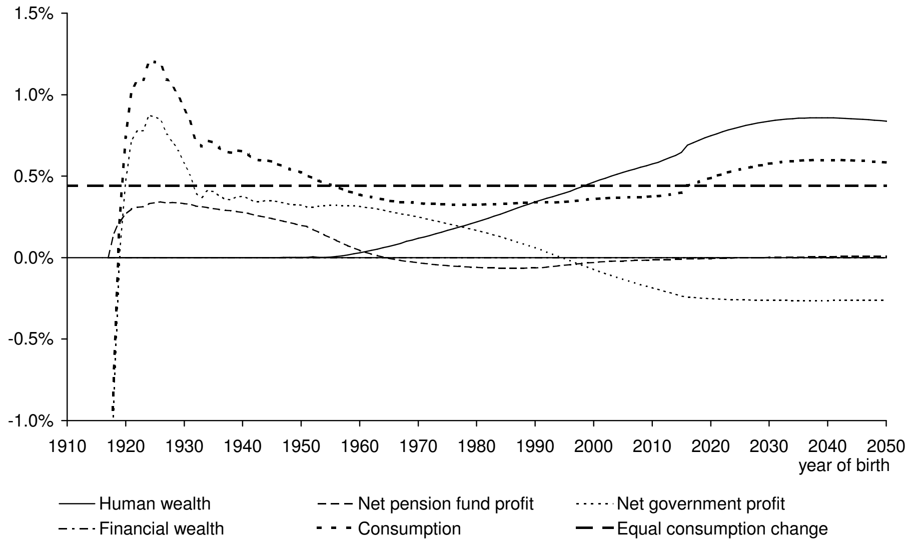

Figure 13: Discounted, expected, lifetime net benefit change from a mixed shock in mortality. The changes are expressed relative to consumption and split into human wealth, pension funds, net government profit, and financial wealth.

5 Conclusion

In this paper, we study the effects of macroeconomic shocks on different cohorts in the Dutch economy. We aim for realism rather than optimality: the shocks come from a calibrated model of relevant macroeconomic risks and the effects are computed using a large OLG model that contains detailed information on the way taxes and benefits are computed in the Dutch economy. This approach leads to compromises: the OLG model is deterministic, uses perfect foresight, and welfare analysis is problematic. We do, however, succeed in showing the impact of different kinds of shocks on the wealth components of different cohorts before and after redistribution.

We find that both the government and the pensions funds engage in the redistribution of shocks. The government does this by taxing the winners and transfering the money to those who lose from the shock. Pension funds also engage in redistribution but their role is more subtle: when financial assets lose value, the effect on private savers is much more concentrated (in terms of cohorts affected) than the effect on current and future members of the pension fund. For the relevant shocks it turns out that the government is by far the most important party in this redistribution. The changes it causes in the wealth of cohorts are several times as large as those caused by pension funds. This is an important lesson that can be taken away from this paper: while pension funds are often thought to play an important role in intergenerational risk sharing, the importance of their role is trumped by that of the government.

Note however that we take a very kind view of the way in which the government deals with sustainability problems, *i.e.* by spreading them over all possible cohorts. If the actual response is to fix problems by keeping balances close to zero, the amount of risk sharing decreases. We also neglect bequests and fix agents' portfolios, thus ruling out private risk-sharing arrangements which may overstate the role of the government.[4]

Finally, we find that only in the case of rare disasters do we see a relevant cross-correlation between different kinds of macroeconomic shocks. During these events, financial assets, discount rates and productivity growth are all affected at the same time. Outside of these disasters, shocks in one area can (in theory) have effects in another. We find that the size of these second-order effects does not warrant excessive worrying by researchers or policy makers. The effect of cross-correlations between shocks on generational wealth changes is quite small compared to the initial impact of the shock. It is notable, however, that rare disasters are one of the few cases where the government redistributes the wrong way, *i.e.* from cohorts that lose substantially towards cohorts that are much less affected.

References

Altonji, J., Hayashi, F. and Kotlikoff, L.: 1992, Is the extended family altruistically linked? Direct tests using micro data, *American Economic Review* **82**, 1177–1198.

Auerbach, A. J., Kotlikoff, L. J. and Leibfritz, W.: 1999, *Generational Accounting around the World*, University of Chicago Press, Chicago, IL.

Barro, R. J.: 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* **121**(3), 823–866.

Bohn, H.: 2006, Who bears what risk? An intergenerational perspective, *in* D. Blitzstein, O. S. Mitchell and S. P. Utkus (eds), *Restructuring retirement risks*, Oxford University Press, USA, chapter 2, pp. 10–36.

Bohn, H.: 2009, Intergenerational risk sharing and fiscal policy, *Journal of Monetary Economics* **56**, 805–816.

Bonenkamp, J., van Ewijk, C., ter Rele, H. and Westerhout, E.: 2009, Herstel dekkingsgraad pensioenfondsen vergt grote inkomensoffers, *Economisch Statistische Berichten* **94**(4556), 166–169.

Broer, D. P.: 2010, Macroeconomic risks and pension returns. CPB Memorandum 241, The Hague.
URL: *http://www.cpb.nl/nl/pub/cpbreeksen/memorandum/241/*

[4]Note however that Altonji et al. (1992) present empirical evidence which suggests that dynastic risk sharing is rather weak.

Campbell, J. Y. and Viceira, L. M.: 2005, The term structure of the risk-return tradeoff, *CEPR Discussion Paper 4914*, CEPR.
URL: *http://www.cepr.org/pubs/dps/DP4914.asp*

Cui, J., de Jong, F. and Ponds, E. H.: 2006, Intergenerational risk-sharing within funded pension schemes. Netspar Discussion Paper 2006-016, University of Tilburg.
URL: *http://arno.uvt.nl/show.cgi?fid=57905*

Draper, D. and Armstrong, A.: 2007, Gamma, a simulation model for ageing, pensions and public finances. CPB Document 147.
URL: *http://www.cpb.nl/nl/pub/cpbreeksen/document/147/*

Draper, N., Bettendorf, L., Bonenkamp, J., Nibbelink, A., Rosenbrand, R. and ter Rele, H.: 2010, Calibration of Gamma 2010. CPB Netherlands Bureau for Economic Policy Analysis, Memorandum 248.
URL: *http://www.cpb.nl/nl/pub/cpbreeksen/memorandum/248/*

Gollier, C.: 2008, Intergenerational risk-sharing and risk-taking of a pension fund, *Journal of Public Economics* **92**(5-6), 1463–1485.

Hári, N., de Waegenaere, A., Melenberg, B. and Nijman, T. E.: 2008, Estimating the term structure of mortality, *Insurance: Mathematics and Economics* **42**, 492–504.

Lee, R. D. and Carter, L. R.: 1992, Modeling and forecasting U.S. mortality, *Journal of the American Statistical Association* **87**, 659–671.

Nelson, C. R. and Plosser, C. I.: 1982, Trends and random walks in macroeconomic time series. some evidence and implications, *Journal of Monetary Economics* **10**(2), 139–162.

van der Horst, A., Bettendorf, L., Draper, N., van Ewijk, C., de Mooij, R. and ter Rele, H.: 2010, Vergrijzing Verdeeld: Toekomst van de Nederlandse Overheidsfinanciën. CPB Bijzondere Publicatie 86 (in Dutch).
URL: *http://www.cpb.nl/nl/pub/cpbreeksen/bijzonder/86/*

Yaari, M. E.: 1965, Uncertain lifetime, life insurance and the theory of the consumer, *Review of Economic Studies* **32**, 137–150.