KU LEUVEN

CENTER FOR ECONOMIC STUDIES

DISCUSSION PAPER SERIES
DPS13.03

FEBRUARY 2013

The impact of an increase in the legal retirement age on the effective retirement age

Noelia BERNAL and Frederic VERMEULEN

Public Economics

Faculty of Economics
And Business

The impact of an increase in the legal retirement age on the effective retirement age*

Noelia Bernal† and Frederic Vermeulen‡

February 27, 2013

Abstract

We analyze the impact of an increase in the legal retirement age on the effective retirement age in the Netherlands. We do this by means of a dynamic programming model for the retirement behavior of singles. The model is applied to new administrative data that contain very accurate and detailed information on individual incomes and occupational pension entitlements. Our model is able to capture the main patterns observed in the data. We observe that as individuals get older their labor supply declines considerably and this varies by health status. We simulate a soon to be implemented pension reform which aims at gradually increasing the legal retirement age from 65 to 67. The simulation results show a rather small impact on the effective retirement age. Individuals postpone their retirement by only 3 months on average, while differences across individuals mainly depend on their health status.

Key words: normal retirement age, effective retirement age, dynamic programming.
JEL-classification: C61, D12, J26.

1 Introduction

Population ageing is one of the most important challenges that are posed to the OECD countries. The combination of an increasing life expectancy and a declining fertility implies a big burden on the long term sustainability of public pension programmes in many countries. As a response to this challenge, a growing number of countries raised, or are on the verge of raising, their official retirement age as part of a pension reform. Such a change was quite early adopted by the United States' government. One of the provisions in the Social Security Amendments of 1983, for example, gradually increased the age for collecting full social security benefits from 65 to 67 over a long period that began in 2000. Similar policies are being adopted in many member states of the European Union, with Denmark, Germany and the United Kingdom as important examples (see European Commission, 2012, for a detailed overview).

Still, it is well-known that in many OECD countries, there is a substantial gap between the official retirement age and the effective retirement age. Although there are notable exceptions (like Japan, Korea and Mexico), the average effective retirement age is lower than the official retirement age in most of the OECD countries. In countries like Austria, Belgium and Luxemburg, the average gap is not less than five years (OECD, 2009).

The aim of this paper is to analyze the impact of an increase in the official retirement age on the effective retirement age in the Netherlands. Like in many other OECD countries, there is an agreement between the Dutch government and the social partners to gradually increase the official retirement age from 65 to 67 by 2023. Obtaining insights into the efficacy of this reform is therefore of substantial policy relevance.

*We thank Tobias Klein, Robert Olieman, Ward Romp and seminar participants at the Central Bank of Peru, the University of Piura, the SVB and a Netspar Pension Day for useful comments.

†Netspar, CentER, Tilburg University. P.O. Box 90153, NL-5000 LE Tilburg, The Netherlands. E-mail: n.bernallobato@uvt.nl. Noelia Bernal gratefully acknowledges financial support from Netspar and and the Sociale Verzekeringsbank.

‡Department of Economics, University of Leuven. Naamsestraat 69, B-3000 Leuven, Belgium. E-mail: frederic.vermeulen@kuleuven.be. Frederic Vermeulen gratefully acknowledges financial support from The Netherlands Organisation for Scientific Research (NWO) through a VIDI grant and the Research Fund KU Leuven through the grant STRT/12/001.

A first feature of our analysis is that we make use of a dynamic programming (DP) model of retirement. DP models are based on recursive methods for solving sequential decision problems; in our case the decision when to retire. Such models have enabled economists to formulate and solve a variety of problems involving decisions over time and under uncertainty. As a result, they have now become an important and versatile tool in a host of areas, including labor economics, industrial organization, economic demography and marketing (see, e.g., Rust, 1994, 2006, and Adda and Cooper, 2003, for numerous references). Given the fact that the increase of the official retirement age will only be in full force by 2023, a DP model seems the most natural starting point for our analysis. One of the advantages of structural modeling is its potential to predict the impact of future or hypothetical policy changes.

A second feature of our study is that we will exclusively focus on older singles' retirement behavior. This choice has the big advantage that it allows one to circumvent well-known issues associated with the modeling of retirement decisions of individuals in couples. Most of these issues can be brought back to the question which model should be used to model such individuals' retirement decisions. Some authors (for example, Rust and Phelan, 1997, French, 2005, French and Jones, 2011, and Eckstein and Lifshitz, 2011) focus on the labor supply decisions of one of the spouses in a couple while they take the behavior of the other spouse as exogenously determined. They consider a unitary model, which assumes that a couple behaves as a single decision maker. One issue here is that there is quite some empirical evidence that this assumption is too strong (see Browning and Chiappori, 1998, Cherchye and Vermeulen, 2008, and Cherchye, De Rock and Vermeulen, 2009, for some recent examples). Other authors explicitly account of the fact that there are multiple decision makers in couples. Gustman and Steinmeier (2000, 2004), for example, assume a noncooperative approach, where each spouse maximizes own utility while taking into account the actions of the other spouse. Van der Klaauw and Wolpin (2008), on the other hand, assume that couples' preferences are captured by some weighted combination of both spouses' preferences. A collective model à la Chiappori (1988) is considered by Michaud and Vermeulen (2011), though they focus on a static retirement model. Given the above, it can be argued that our focus on the behavior of pure singles, where, given individual rationality, the unitary model applies to by definition, offers us a very clean setting to conduct our empirical analysis.

A third feature of our analysis is that the structural model that we build in this paper is applied to new administrative data from Statistics Netherlands (CBS). The main advantage of our data is that it includes very accurate and detailed information on individual incomes and pension entitlements. This level of accuracy will be most useful when we formulate the intertemporal individual budget constraints, with which rational preference maximizing individuals are confronted under uncertainty. As far as we know, an analysis that is based on Dutch administrative data on incomes and pension entitlements has not been conducted yet. A previous study for the Netherlands was made by Heyma (2004). He also used a DP model, but the latter was used on survey data that did not contain any information on pension entitlements.

The main aim of our empirical analysis is to investigate in what way the pension system, and more specifically the official retirement age, affects labor supply behavior of older single individuals in the presence of uncertainty about income, health status and life expectancy. Our model closely follows the models of retirement behavior proposed by Rust (1989), Rust and Phelan (1997) and Karlstrom, Palme and Svensson (2004). Like these authors, we focus on the binary choice between working and not working and assume that consumption equals net income in each period. Although this setting is restrictive, our data set does not allow us to go any further in this respect. It should also be noted though that the Netherlands have a well-developed occupational pension scheme (which we take into account in our analysis) that makes other savings relatively less important. One substantial difference between our application and those in Rust (1989), Rust and Phelan (1997) and Karlstrom, Palme and Svensson (2004) is that we solely focus on singles while they apply a unitary model to individuals who can also be married.

The rest of the paper is organized as follows. Section 2 provides a brief description of the pension system in the Netherlands and describes the administrative data that we use. Section 3 presents our DP model and gives details on the empirical specification. The estimation results are presented in Section 4. In Section 5, we provide simulation results with respect to the planned increase in the official retirement age. More specifically, we will analyze the impact of this policy change on the effective retirement age. Section 6 concludes.

2 Background and data

2.1 The pension system in the Netherlands

The Dutch pension system combines three pillars. As Ewisk (2005) describes, the first pillar consists of pensions that are pay-as-you-go financially based (with the Dutch abbreviation AOW). This state benefit is an old-age pension provided to all residents in the Netherlands at the age of 65 and it is linked to the minimum wage. No distinction is made between men, women, employees, self-employed or immigrants. There is no means-test to check eligibility, which implies that other forms of income have no effect on the level of the AOW benefit. It depends on household composition though. Single individuals receive a different benefit than cohabiting or married individuals. All residents between the ages of 15 and 65 are insured for the AOW. The entitlement accrual equals 2 percent for every insured year, which leads to a 100 percent entitlement for individuals who live in the Netherlands for 50 years without gaps. A gap occurs when a person resides outside the Netherlands.

The second pillar comprises the old-age occupational pensions (OP), which are privately organized by employers and employees. These occupational pensions are mandatory, funded and defined benefit for the large majority of workers. There are two ways to define the level of expected benefits. In the final pay scheme, the pension is based on an annual replacement rate of 1.75 percent of the final salary, whereas, in the average pay scheme, the pension is calculated based on a replacement rate of 2 percent of the average career salary. These accrual rates imply that, by each year individuals decide to continue working, they earn those percentages of their respective salary. It is expected that if the working period is between 35 to 40 years, the total pension benefit will be around 70 percent of the salary, including first pillar benefits. Therefore, the occupational pension scheme is considered supplementary to the AOW pension.

The third pillar comprises voluntary pensions. They are intended to complement or improve the occupational and AOW benefits and they must be organized by an insurance provider. This third pillar consists of defined contribution pensions and part of these receive a favorable tax treatment.

In this paper, we focus on the first and second pillars. Therefore, we will define an individual's retirement income as the sum of the state and occupational pensions.

2.2 Data

The data set used in our study is drawn from a combination of three administrative data sets from Statistics Netherlands (CBS). The first data set is called "Pension Entitlements" and it is available for the years 2005, 2006 and 2007. This data set contains new and valuable information on entitlements and attainable pensions from the second pillar. It is provided from pension funds, which ensures a high level of accuracy. Moreover, it also contains information about the franchise, which is an important variable for the computation of the pension income.

The second data set comes from an income panel study (called IPO). It contains detailed information about individual and household incomes and it is provided by the Dutch national tax administration. Again, this ensures a high level of accuracy. This data is basically a panel survey because it is based on a randomly drawn set of "key persons" who are followed over time. Information about all the household members in the key persons' households is also recorded. To match this data set with the pension data, we use the waves 2005, 2006 and 2007.

A third data set that we use is drawn from a population registration (called GBA). This dataset is used to exclude individuals who migrated and to take into account the fact that some individuals we focus on died over the period covered.

By merging these three data sets (using individual specific identifiers), we build a panel that contains the main information we need for our DP model. In addition to demographic characteristics such as age and gender, we have detailed information on incomes and pension entitlements. A drawback of our data set is that it does not contain any information about the individuals' educational level or on detailed categories of their health status. In the case of health status though, we can construct a binary variable, which indicates whether an individual is in good or bad health. The latter is then defined by receiving income from disability benefits.

As mentioned in the introduction, we focus on the retirement decisions of male and female singles, which allows us to consider a standard unitary model. Singles are defined as individuals who are not married or cohabiting, and who are separated, widowed or divorced. We further focus

Table 1: Socio-economic status in 2007

	Employment / Retirement Status				
Age in 2007	Employed	Disabled	Early retired	Retired	Total
60	107	1	15	-	123
61	99	-	25	-	124
62	48	-	20	-	68
63	33	1	11	-	45
64	13	-	5	-	18
65	19	-	-	9	28
66	1	-	-	15	16
67	3	-	-	9	12
68	2	-	-	-	2
69	1	-	-	-	1
70	2	-	-	2	4
71	-	-	-	1	1
72	-	-	-	1	1
Total	328	2	76	37	443

on individuals who are between 58 and 70 years old in the year 2005. These thresholds are chosen because we want to focus on the retirement behavior of older workers. Younger individuals might not be thinking yet about their retirement. Similarly, few individuals above 70 are still working and most of them are retired. We further focus on employees with an observed income (thus excluding the self-employed and individuals who receive assistance). We follow these employees until 2007 and record those who transit to completely disabled, early retired and normal retired status. We do not consider unemployed individuals because the latter are affected by unemployment benefit rules instead of pension rules. Finally, we keep individuals with a defined benefit plan since there are only few individuals with a defined contribution plan.

Our final sample has 1,329 observations of 443 individuals. Table 1 shows the socioeconomic status of these individuals in the year 2007. Note that these individuals will either remain employed or they will retire through complete disability, early retirement or normal retirement options. It is clear from the table that the share of people out of the labor force is increasing over time. Obviously, the pathway through early retirement is initially most important. This status is taken over by normal retirement for the oldest individuals in the sample. Table 2 gives some summary statistics on our sample.

Table 2: Descriptive statistics of initial state variables

	2005	2006	2007
Age (years)	59.98	60.98	61.98
	(2.179)	(2.179)	(2.179)
Male	0.542	0542	0.542
	(0.499)	(0.499)	(0.499)
Native	0.885	0.885	0.885
	(0.320)	(0.320)	(0.320)
Good health	0.955	0.950	0.944
	(0.208)	(0.217)	(0.231)
Employed	1.00	0.844	0.736
	(0.000)	(0.363)	(0.441)
Income (euros)	42,521	44,126	45,876
	(13,522)	(14,043)	(14,703)
Pension entitlement (euros)	28,842	29,433	30,055
	(6,440)	(6,630)	(6,833)
Sample	443	443	443

Note: Standard deviations in parentheses.

3 The model

3.1 A dynamic programming formulation

We formulate our retirement model as one where individuals are faced with a sequential decision problem in a discrete finite horizon setting. We use the standard assumption that individuals are expected discounted utility maximizers. In each time period, the decision to retire is modeled as a binary choice between working and retirement. Individuals are assumed to observe the available information in each period and, conditional on that, calculate expected discounted utilities for each of the two alternative employment states. Finally, they are assumed to choose the alternative that maximizes their expected discounted utility. The choice between working and retirement in the next period defines our control (decision) variable, whereas the available information about income, pension entitlements, current employment status and some demographic variables define the set of state variables. Since individuals do not know their future labor income nor their future health status, they have to make decisions under uncertainty. This uncertainty then affects their decisions, which is modeled through conditional probabilities. These probabilities represent the individual's beliefs and are used to calculate expected discounted utilities.

Since this type of problem generally does not have a tractable analytical solution, we follow the typical approach that is based on Bellman's optimality principle. This implies that we use a backward induction process to obtain an optimal decision given certain conditions on a controlled process. To implement backward induction, we start in the last time period and for each possible combination of the state and control variables we calculate expected discounted utilities and decision rules. We continue the backward induction recursively for previous periods until we reach the first time period. This results in a decision rule that contains an optimal retirement sequence given individual beliefs and constraints. At every time period, individuals take the decision to retire if this alternative brings the highest expected discounted utility for every possible continuation of the problem. This is made by comparing the value functions for each alternative state, which summarizes the future consequences of choosing each alternative accounting for the uncertainty we described before.

The individual's period preferences are represented by a random utility function $U_t(s_t, d_t, \theta_u)$, where s_t is a vector of state variables at year t, d_t denotes the control variable (if the individual decides to stop (continue) working in time period t, then the control variable takes the value 1 (0)), and θ_u the parameters to be estimated. Following Rust (1989) and Rust and Phelan (1997), we assume a partition of the state variables into two components: $s_t = (x_t, \varepsilon_t)$, where x_t is a vector of observed state variables and ε_t is a vector of state variables that is observed by the individual but not by the econometrician. The vector of observed variables x_t considers the individual's current employment status, labor income, retirement income, health status, origin and age at time t. For empirical tractability, we assume the following additive form for the period utility function:

$$U_t(s_t, d_t, \theta_u) = u_t(x_t, d_t, \theta_u) + \varepsilon_t(d_t),$$

where the alternative specific error term $\varepsilon_t(d_t)$ is assumed to be independent and identically distributed according to a type 1 extreme value distribution. As mentioned above, the individual has to choose in an environment with uncertainty about future incomes, future health and life expectancy. This uncertainty is modelled through conditional probabilities that are represented by a transition probability matrix $p_t(x_{t+1}|x_t, d_t, \theta_p)$, where θ_p denotes a vector of unknown parameters that characterize an individual's expectations (or beliefs) about those uncertain variables.

Following Rust and Phelan (1997), we formulate our problem in terms of a value function $V_t(s_t)$, which summarizes the future consequences of choosing each alternative (retirement or working) while accounting for the uncertainty an individual is faced with. This function represents the expected discounted utility of an individual who is in state s_t and follows an optimal decision from time t onwards until she reaches the final period T (set at the year where the individual is 70 years old):

$$V_t(s) = \max E\left[\sum_{t=1}^{T} \beta^t U_t(s_t, d_t, \theta_u)|s_t = s\right]. \tag{1}$$

This value function, as well as its associated decision rule, depends on the underlying primitives of the structural model $(U_t\,(.)\,, p_t\,(.))$, which, on their turn, depend on two sets of parameters. The first set of parameters is the vector $\theta = (\theta_u, \theta_p, \beta)$, which contains the preference (θ_u) and beliefs (θ_p) parameters and the discount factor (β). The second set of parameters is the vector τ that contains the rules of the pension system (such as the normal retirement age, the accumulation of pension entitlements and the level of benefits). The details of how these rules influence the computation of pensions are given later.

By our assumption on the partition of the state variables, $s_t = (x_t, \varepsilon_t)$, and integrating out the unobserved state variables, we can derive a conditional choice probability $P_t(d|x, \theta, \tau)$ that provides the basis for estimating the unknown model parameters. More specifically, through the assumption of a type 1 extreme value distribution of $\varepsilon_t(d_t)$, a multinomial logit representation of the conditional choice probability can be derived:

$$P_t(d = d'|x, \theta, \tau) = \frac{\exp(v_t(x_t, d = d', \theta, \tau))}{\sum_{d \in D(x_t)} \exp(v_t(x_t, d, \theta, \tau))},$$

where $D(x_t)$ is the choice set in state x_t and v_t is the expected value function defined recursively by:

$$v_t(x_t, d_t, \theta, \tau) = u_t(x_t, d_t, \theta_u) + q_{t+1}\beta \sum_{\delta \in \Delta} \left\{ \log \sum_{d_{t+1} \in D(x_{t+1})} \exp(v_{t+1}(x_{t+1}, d_{t+1}, \theta, \tau)) \right\} p_t(x_{t+1}|x_t, d_t, \theta_p, \tau),$$

where q_{t+1} is the individual's survival probability from period t to $t+1$, and Δ is the set of possible transitions.

The expected value function $v_t(.)$ is related to the value function $V_t(.)$ defined in equation (1) by the following identity:

$$V_t(x_t, \varepsilon_t) = \max_{d \in D(x_t)} \left[v_t(x_t, d_t, \theta, \tau) + \varepsilon_t(d_t) \right].$$

The next step consists of defining the likelihood function to estimate the model parameters. Given panel data on observed state and control variables, x_t^i, d_t^i (where i is an index to refer to an individual in the data, $i = 1, ..., I$), we can estimate the model parameters by looking for the value of θ such that the following likelihood function is maximized:

$$L(\theta) = L(\theta_u, \theta_p, \beta) = \prod_{i=1}^{I} \prod_{t=1}^{T} P_t(d_t^i|x_t^i, \theta, \tau) p_t(x_t^i|x_{t-1}^i, d_{t-1}^i, \theta_p, \tau).$$

To estimate the model, we follow the two-stage estimation procedure that was proposed by Rust (1987). In a first stage, the beliefs parameters θ_p are estimated by using a partial likelihood function involving only products of the conditional probabilities $p_t(.)$. In a second stage, the estimates of θ_p are used to solve the backward recursion numerically which allows one to estimate the remaining parameters (β, θ_u) using a partial likelihood function with only products of the choice probabilities $P_t(.)$. Although this procedure is not as efficient as full maximum likelihood, Rust (1989) and Rust and Phelan (1997) argue that the efficiency loss is not too big and the computational burden is considerably reduced. Karlstrom, Palme and Svensson (2004) use a similar estimation procedure in their study.

3.2 Individual preferences

The structural part of the individual's preferences at time t are assumed to be represented by a Cobb-Douglas utility function, which implies:

$$U_t(s_t, d_t, \theta_u) = [\alpha \ln c_t + (1 - \alpha) \ln l_t] + \varepsilon_t(d_t),$$

where c_t and l_t respectively denote consumption and leisure at time t. The parameter α is associated with the consumption share in the individual's full income at time t, while $(1 - \alpha)$ represents the share of leisure. In the empirical analysis, we assume that α depends on the individual's health status (denoted by h_t), which acts as a preference shifter.[1] In particular, we assume that $\alpha = \exp(\alpha_0 + \alpha_1 h_t)/(1 + \exp(\alpha_0 + \alpha_1 h_t))$, where α_0 and α_1 are to be estimated. This individual preference shifter is a state variable in the DP model. Coherency of the utility function is guaranteed since $\alpha \in]0, 1[$.

As stated before, there are no savings in our model. As such, an individual's consumption in period t equals her income in that period. This income contains both labor income and non-labor income, where the latter contains pension benefits for individuals who are claiming benefits. In the case of leisure, we define a set of two leisure options: one leisure value that corresponds to full time working and one leisure value that corresponds with retirement. Taking into account a working week of 40 hours and 46 working weeks per year and allowing for time to sleep and personal care, we obtain $1,040$ hours of leisure for the option of working and $2,880$ hours for the retirement option.

3.3 Control and state variables

The next step is to carefully define control and states variables in the DP model. The constructed panel allows us to formulate a model with a five-dimensional vector of observed state variables and a one-dimensional vector of control variables.

3.3.1 Control variable

d_t: This binary variable denotes the employment or retirement decision. It takes the value $d_t = 1$ when the individual retires. It means that she decides to enjoy income from only state and occupational pensions and $2,880$ hours of leisure. In contrast, the decision variable takes zero value, $d_t = 0$, when the individual continues working, which means that she chooses to receive labor income (and possibly an occupational pension at a later age) and enjoy $1,040$ hours of leisure. Similar to Rust and Phelan (1997) we assume that individuals have a perfect control over their future employment status, so $e_{t+1} = d_t$ with probability 1, where e_{t+1} denotes the employment status at period $t + 1$. Although individuals are uncertain about their future income, this assumption implies that an individual who decides to continue working is committed to that decision until the next period in which she makes a new decision.

3.3.2 State variables

c_t: Given the assumption of no savings, consumption is set equal to the individual total annual income, that is discretized into five intervals. The total income includes labor and non labor income net of contributions. Non labor income is defined as the sum of a disability benefit, a state pension, a survivor and an occupational pension income (where some incomes, of course, are equal to zero for some states). Before retirement, the data shows that the main income source comes from labor income (although non labor income from a disability benefit is quite important for many individuals). After retirement, the main source of income comes from state and occupational pensions. We deduct from the total income those contributions made to the state and occupational pension systems. This in order to include a most accurate measure of income in the utility function specification of our model. The cut points for the income intervals are constructed by taking the 20th, 40th, 60th and 80th percentile of the distribution of income. For each year, the specific values of the cut points vary according to the income profile estimation discussed in the next subsection.

g_t: This binary variable denotes the gender of an individual. It takes the value of 1 for males and 0 for females.

h_t: This binary variable captures an individual's health status in period t, which can be either good or bad. We consider individuals in bad health if they receive disability income according to the IPO dataset.

The unique decision (control) variable for individuals is whether to retire or to continue working. We assume that this decision is taken at the beginning of each year. This means that, at the beginning of each year, the individual observes the available information and, conditional on that, she calculates expected discounted utilities for the two alternatives and chooses that alternative which maximizes expected discounted utility. The available information (state variables) at the beginning of each year is the income in the previous period, the pension entitlements she has

[1] In principle, more preference shifters can be included in the analysis. However, given the particular data at hand, we faced a multicollinearity problem when we also included variables like age and gender.

built in the state and occupational pension system, the age and the current employment status. However, the individual is uncertain about future income, future entitlements and consequently future pensions. Therefore, we need to estimate the evolution of these variables and, as indicated before, we use these estimates (beliefs) in the second stage of the estimation procedure.

4 Estimation results

In what follows, we will first discuss the evolution of the state variables and give the (first stage) estimation results of the individuals' beliefs. Next, we will discuss the (second stage) estimation results of the preference parameters.

4.1 Evolution of state variables and estimation of beliefs

The transition matrix $p_t(.)$ represents the individual's beliefs about her future health, income and life expectancy. Similar to Rust (1989), Rust and Phelan (1997) and Karlstrom, Palme and Svensson (2004), we impose two main assumptions: the assumption of individual rational expectations and exclusion restrictions. The first assumption implies that beliefs about future health, income and life expectancy coincide with the population behavior of these variables. The second assumption implies that we can decompose the transition matrix as a product of conditional probabilities for each component and estimate them separately. The following equation shows the decomposition of $p_t(.)$:

$$p_t(x_{t+1}^i|x_t^i, d_t^i, \theta_p) = \pi_y(y_{t+1}^i|y_t^i, d_t^i, h_t^i, a_t^i) \times \pi_h(h_{t+1}^i|h_t^i, a_t^i) \times \pi_q(q_{t+1}^i|q_t^i, g_t^i, a_t^i),$$

where $\pi_y(.)$ is the conditional transition probability for income, $\pi_h(.)$ is the conditional transition probability for health status and $\pi_q(.)$ is the conditional survival probability. Each of these conditional probabilities can be estimated independently of each other in the first stage of the estimation procedure. In the second stage, we use these estimates to solve the DP model by recursion in a numerical way and estimate the remaining parameters (β, θ_u).

4.1.1 Labor income

We ran two regressions to estimate the labor income profiles, the fixed and the random effects model. Our preferred specification assumes that (log) income from employment is explained by demographic variables such as age, health, gender, origin and birthyear. As expected, we find that under both models age and health have a positive and significant effect on income (see Table 3). Older workers expect a 3.3 percent annual increase in their income whereas healthy workers expect an increase of respectively 19.5 and 27.9 percent relative to unhealthy workers all else equal. Note that the estimated coefficient of health status varies depending on the model chosen.

Recall that the fixed effects model allows us to control for omitted time-invariant variables such as education level. Since we lack educational information in all of the data sets that we use, this model might be more appropriate than the random effects model. We performed a Hausman test and we find that the fixed effects model is to be preferred over the random effects model.

Next, we predict the evolution of labor income by using the fixed effects estimates and the estimated unobserved effects. Following Rust (1989) and Rust and Phelan (1997) we do this by using the continuous (log) income variable and we obtain increasing income profiles explained by the significantly positive effect of age. Nevertheless, since we do not observe many individuals working after 65 years old, we stop the prediction at this age and assume that individuals who work beyond that age keep their last income. This assumption is reasonable if we consider that few employers and employees at this stage take actions to significantly improve productivity.

The discretization of the income distribution is an important concern in our model. To keep it simple and numerically feasible, we only consider five grids but it could easily be extended to an arbitrary number of grids, with the unavoidable increasing computation time though. For example, Rust and Phelan (1997) consider 25 intervals for the total family income, whereas Karlstrom, Palme and Svensson (2004) consider 400 intervals for the labor earnings. However, they mention that their estimation results were rather robust with smaller numbers of points. In our case, we compute the

Table 3: Estimation of labor income

	Fixed effects	Random effects
Age	0.033^{**}	0.033^{**}
	(0.003)	(0.003)
Male		0.397^{**}
		(0.022)
Native		0.111^{**}
		(0.031)
Health	0.195^{**}	0.279^{**}
	(0.038)	(0.032)
Birthyear		0.041^{**}
		(0.003)
Constant	8.313^{**}	-72.11^{**}
	(0.143)	(6.665)
σ_u	0.636	0.558
σ_e	0.188	0.188
ρ	0.920^{**}	0.898

Note: A double asterisk refers to significance at the 95 percent level. Standard errors in parentheses

cut points for the intervals by taking the 20th, 40th, 60th and 80th percentiles of the distribution and we repeat this process during the whole period of analysis.

We then estimate the labor income transition probability matrix conditional on previous income, health status and age category, $\pi_y(y_{t+1}|y_t, d_t = 0, h_t, a_t)$. We decide to condition on the last two variables because both have a positive and significant effect on the income profiles according to the fixed effects model. We also take this decision because we do not have enough observations per age if we additionally consider other variables such as origin or gender. Moreover, in order to have enough observations, we group individuals in five classes, where each class contains five age intervals starting from 58 years old until 70 years old (the last class includes individuals between 65 and 70 years old). We assume that individuals in each class share the same transition probabilities.[2]

Table 4 displays the income transition probability matrix. Each cell displays the probability of being in each quintile in period $t+1$ conditional on having been in a specific quintile, and having a specific health status and age category in period t. For instance, a 58 years old individual who has been observed in good health status and in the first quintile of the distribution has a 97 percent probability to stay in the same quintile in the next period, whereas a 60 years old individual has only a 94 percent probability to remain in the same quintile. In general, conditional probabilities tend to decrease with age for both health states.

4.1.2 Pension income

The next step is to estimate the pension entitlements and the pension income transition probability matrix, $\pi_y(y^i_{t+1}|y^i_t, d^i_t = 1, a^i_t)$. We model the AOW (state pension) entitlements in a deterministic way by assuming that in each year the individual resides in the Netherlands she accumulates 2 percent of the full benefit. This is based on the current regulation which says that the accumulation of AOW entitlements depends on the years an individual has lived in the Netherlands from his 15 birthday until he is 65 years old. The accrual rate (ar) is assumed to be 2 percent for each year in which there is insurance and 0 percent for the years the individual lives abroad (b^i_t):

$$AOWentitlement^i_t = (50 - b^i_t)ar.$$

The expected state pension, $AOWpension^i_t$, is obtained by simply multiplying the entitlement with the annual AOW benefit at year t, $AOWbenefit^j_t$. The annual AOW benefit for a single equals € 11,211 in 2005. This amount is multiplied by 100 percent if the individual has lived in the Netherlands for at least 50 years, whereas it is multiplied by a lower percentage if the individual has lived abroad. Since we do not have a variable that records years lived abroad, we are unable to accurately compute the expected state pension. Therefore, we use the origin of individuals as a

[2] We do not have enough observations for individuals in the last class, so we have assumed that they share the same transition probabilities as individuals in the previous categorie (60-64 years old).

Table 4: Transition probability matrices: Labor income

Good health; age 58-59	$\hat{\pi}_y(y_{t+1}\|y_t, d_t = 0, h_t = 1, a_t)$				
y_t	20^{th}	40^{th}	60^{th}	80^{th}	100^{th}
20^{th}	0.97	0.03	0.00	0.00	0.00
40^{th}	0.02	0.94	0.04	0.00	0.00
60^{th}	0.00	0.02	0.95	0.03	0.00
80^{th}	0.00	0.00	0.00	0.97	0.03
100^{th}	0.00	0.00	0.01	0.04	0.95
Good health; age 60-64	$\hat{\pi}_y(y_{t+1}\|y_t, d_t = 0, h_t = 1, a_t)$				
y_t	20^{th}	40^{th}	60^{th}	80^{th}	100^{th}
20^{th}	0.94	0.04	0.00	0.01	0.00
40^{th}	0.06	0.84	0.10	0.00	0.00
60^{th}	0.02	0.02	0.85	0.12	0.00
80^{th}	0.01	0.03	0.03	0.78	0.15
100^{th}	0.00	0.00	0.04	0.04	0.91
Bad health; age 58-59	$\hat{\pi}_y(y_{t+1}\|y_t, d_t = 0, h_t = 0, a_t)$				
y_t	20^{th}	40^{th}	60^{th}	80^{th}	100^{th}
20^{th}	1.00	0.00	0.00	0.00	0.00
40^{th}	0.13	0.87	0.00	0.00	0.00
60^{th}	0.00	0.17	0.83	0.00	0.00
80^{th}	0.00	0.00	0.25	0.50	0.25
100^{th}	0.00	0.00	0.01	0.04	0.95
Bad health; age 60-64	$\hat{\pi}_y(y_{t+1}\|y_t, d_t = 0, h_t = 0, a_t)$				
y_t	20^{th}	40^{th}	60^{th}	80^{th}	100^{th}
20^{th}	0.75	0.00	0.25	0.00	0.00
40^{th}	0.00	0.67	0.33	0.00	0.00
60^{th}	0.23	0.00	0.77	0.00	0.00
80^{th}	0.00	0.00	0.00	1.00	0.00
100^{th}	0.00	0.00	0.04	0.04	0.91

Table 5: Transition probability matrices: Pension income

Age 58-59	$\hat{\pi}_y(y_{t+1}\|y_t, d_t{=}\,1, a_t)$				
y_t	20^{th}	40^{th}	60^{th}	80^{th}	100^{th}
20^{th}	0.94	0.04	0.01	0.01	0.00
40^{th}	0.06	0.86	0.06	0.01	0.01
60^{th}	0.00	0.09	0.79	0.11	0.01
80^{th}	0.00	0.00	0.11	0.82	0.07
100^{th}	0.00	0.00	0.01	0.11	0.88
Age 60-64	$\hat{\pi}_y(y_{t+1}\|y_t, d_t{=}\,1, a_t)$				
y_t	20^{th}	40^{th}	60^{th}	80^{th}	100^{th}
20^{th}	0.84	0.05	0.05	0.02	0.05
40^{th}	0.03	0.77	0.09	0.00	0.11
60^{th}	0.00	0.06	0.70	0.11	0.13
80^{th}	0.01	0.00	0.06	0.79	0.14
100^{th}	0.01	0.01	0.00	0.04	0.95

proxy variable in our estimation. If the individual was born in the Netherlands, we assume that she retires receiving the full benefit of € 11,211. If the individual was born abroad, she is assumed to have gaps equivalent to 30 percent, so she only gets an incomplete state pension of € 7,848.

In the case of the OP entitlements estimation, it is important to realize that the accumulation follows the income process because individuals accumulate a percentage of their income for each year they continue working. In our case, entitlements accumulated until 2007 are recorded in the "Pension entitlements" data set. Therefore, we use these as starting points of the optimization problem. For the future entitlements estimation, we use the estimated labor income profiles ($\hat{y}_t^i$) and apply the following formula:

$$OPentitlement_t^i = (\hat{y}_t^i - franchise_t) r [(1 + index)^\wedge(65 - a_t^i)].$$

The equation says that for a a_t^i years old employee, we assume that she will accumulate entitlements until she reaches the age of retirement and this process depends on the earnings net of franchise, accrual rate (r) and indexation. In our model, labor earnings are those estimated before (using the fixed effects model) and the franchise variable remains the same as the one observed in 2007. The accrual rate equals 2 percent and we do not assume any indexation. The expected annual OP pension is then computed by adding accumulated entitlements until 2007 with the future (accumulated) entitlements until retirement:

$$OPpension_t^i = Accumulatedentitlement_t^i + [\sum_{a_t^i}^{64} OPentitlement_t^i].$$

The total pension income is obtained by simply adding the AOW and OP annual pensions. Similarly as with labor income, we first estimate both pensions using the (log) continuous variables and then, we discretize total pension by taking the same quintiles as for labor income. Next, we estimate the transition probability matrix conditional on previous income and age category, $\pi_y(y_{t+1}^i | y_t^i, d_t^i = 1, a_t^i)$. We decide to condition only on age in order to have enough observations and reasonable probabilities. Similar as before, we group individuals in five age categories which are assumed to have the same transition probabilities.[3]

Table 5 displays the estimated transition probability matrix for pension income. Similar as for the transition for labor income, the conditional probabilities tend to decrease with age.

4.1.3 Health status

Being healthy or not affects the ability to work and enjoy leisure and this effect might be related with previous health status. To construct this variable we use the information recorded in the IPO dataset and we consider individuals in bad health as those who receive a disability benefit.[4]

[3] Individuals in the last categorie share the same transition probabilities as individuals in the previous categorie (60-64 years old).

[4] Note that individuals can receive a disability benefit while still working.

Table 6: Estimation of health status

	Static probit	Dynamic probit
Age	-0.045**	-0.041*
	(0.018)	(0.022)
Male	-0.222	-0.051
	(0.186)	(0.217)
Native	0.145	-0.116
	(0.250)	(0.321)
Lagged health		1.871**
		(0.536)
Initial health		4.637**
		(1.692)
Constant	7.300**	0.243
	(0.999)	(1.184)
σ_u	2.546	1.455
	(0.086)	(0.454)
ρ	0.866**	0.679**
	(0.008)	(0.136)

Note: Dependent variable: good health status. A double asterisk refers to significance at 95 percent level; an asterisk refers to significance at the 90 percent level. Standard errors in parentheses.

Individuals in good health then are those who do not receive this benefit. The advantage of this measure is that it is objective because individuals should be evaluated to determine their degree of disability. Only handicaps above a certain level lead to the right to receive a benefit which implies that individuals must be in a very bad health to claim this benefit. The disadvantage, however, might be that it underestimates the proportion of individuals in bad health, since it is possible to be in bad health and less able to work without receiving any disability income.

After constructing the health variable, we estimate the transition probabilities from the results of a dynamic probit model which accounts of unobserved heterogeneity and previous health condition. In this specification, we control for the initial condition problem by an approach suggested by Wooldridge (2002). We find that the lagged value of health has a stronger effect than other variables alone (which points to strong state dependence). This implies that the chance to continue being in good health depends strongly on having been in good health in the previous period. We also performed a static probit with age, gender and origin as explanatory variables. We found only a significantly estimated (negative) effect of age on the probability of being in good health and a considerable amount of unobserved heterogeneity. Therefore, we decided to stick to the dynamic probit model. Obviously, in our specification, unobserved heterogeneity is still present but less prominent than in the static case (67.9 percent versus 86.6 percent of the unexplained variation is captured by individual effects). Table 6 shows both the static and dynamic probit estimations.

Table 7 shows the estimated health transition probability matrix, $\pi_h(h^i_{t+1}|h^i_t, a^i_t)$. We only condition on previous health status and age because they are the only significantly estimated explanatory variables in our dynamic probit specification. In this case, we do not need to group individuals by age intervals because we obtain reasonable probabilities by age. Our results show that the estimated probability of remaining in good (bad) health as a function of previous good (bad) health status is pretty high, which reflects the strong state dependence discussed before. The negative effect of age on health status is only slightly observed in the probability of remaining in good health, whereas it is more clearly observed in the probability of remaining in bad health. For instance, a 58 years old individual has a 99.9 percent chance to remain in good health status if she has been observed in good health in the previous period, whereas she has 85.8 percent chance to continue having a bad health if she had a bad health status before. In contrast, for a 64 years old individual, the chance to remain in bad health status is much higher (99.5 percent) and the chance to continue having a good health status is practically unchanged.[5]

[5] For individuals older than 66, we assume that they have the same transition probabilities as individuals at this age.

Table 7: Transition probability matrices: Health status

Age=58	$\hat{\pi}_h(h_{t+1}\|h_t, a_t)$	
h_t	Good	Bad
Good	0.99	0.01
Bad	0.12	0.88
Age=60	$\hat{\pi}_h(h_{t+1}\|h_t, a_t)$	
h_t	Good	Bad
Good	0.99	0.01
Bad	0.11	0.89
Age=64	$\hat{\pi}_h(h_{t+1}\|h_t, a_t)$	
h_t	Good	Bad
Good	0.99	0.01
Bad	0.01	0.99

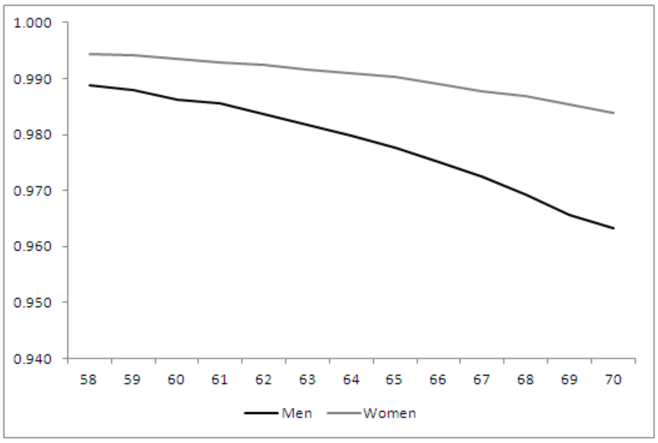

Figure 1: Survival probabilities by age and gender

4.1.4 Survival

The survival probabilities are specified exogenously by age and gender from the current mortality tables available at Statistics Netherlands. As expected, probabilities decline with age for both genders and they are higher for woman across all age cells (see Figure 1).

We did not estimate the survival (mortality) transition probabilities for single individuals because we do not have enough observations to do it consistently, especially for the very old. Most probably, we would have had to use external data sources and extrapolation techniques in order to match mortality rates from Statistics Netherlands projections (see Rust and Phelan, 1997). For simplicity, we just use the current mortality tables conditional on age and gender and we assume that single individuals have the same survival probabilities as the entire population.

4.2 Estimation results with respect to the preference parameters

We next use the estimates of the conditional transition probabilities in the second stage of the estimation procedure to solve the model numerically and estimate the preference parameters (β, θ_u). We do this by incorporating the conditional probabilities in the expected value function $v_t(x_t, d_t, \theta, \tau)$ of the model. This means that, in each time period, individuals calculate their expected discounted utilities for the two alternatives (retire or work) using their probabilities of possible changes in health status, income and survival and choose the alternative that maximizes their expected discounted utility.

Table 8 shows the estimates of the parameters of the utility function.[6] Our baseline results

[6] As often happens in this literature, we set the discount factor equal to 0.97. Though the discount factor is in principle identifiable, there is not sufficient variation in the data to effectively estimate the discount factor.

Table 8: Model estimation result

Parameters	Estimate	Standard errors
α_0 (constant)	1.652	1.325
α_1 (goodhealth)	1.999	2.399
Loglikelihood	-1,277	-

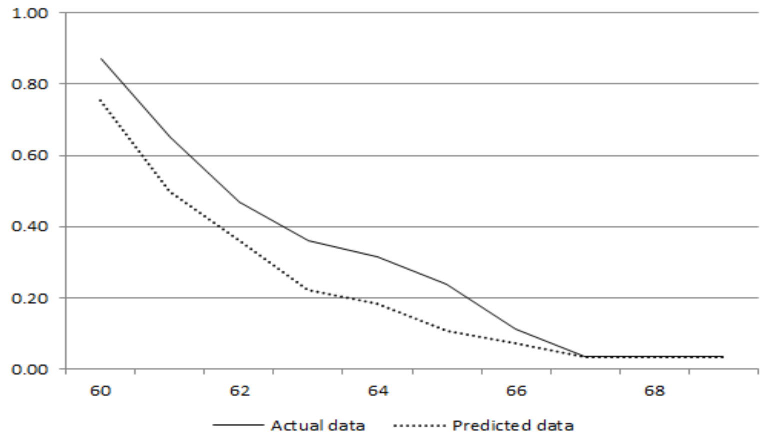

Figure 2: Cumulative distribution of labor force participation

indicate that individuals value consumption and leisure differently depending on whether they are in good or bad health.[7] Healthy individuals tend to value consumption more than less healthy ones (with estimated α-parameters that are equal to respectively 0.97 and 0.84). Therefore, individuals who are in bad health are willing to give up more consumption to gain more leisure than individuals who are in good health.

Figure 2 plots the observed and predicted cumulative probability functions of labor force participation for all individuals. We observe that our estimated parameters allow us to plot a predicted probability that captures the main pattern in the data. However, the model turns out to underpredict the probability to continue working for all cohorts.

Shifting attention to nonparticipation, the model predicts increasing nonparticipation rates as individuals get older, which is in line with the observed data (see Table 9). However, for younger cohorts (below age 64) the model overpredicts the retirement rates in comparison with the actual data, which means that for around 40 percent of individuals in these cohorts the optimal decision is to retire instead of working. For older cohorts (aged between 65 and 69) the model predicts a lower nonparticipation rate compared to the actual data.[8]

In the analysis by health status, we observe that healthy individuals tend to stay longer in the labor force than individuals who have a bad health status. Given their preferences, individuals in good health face a larger trade off between leisure and consumption, which makes them better off when they continue working. Individuals in bad health, on the contrary, turn out to be better off by

[7] Note that the standard errors are only indicative given the nonlinearity of the DP model. As demonstrated by Gregory and Veall (1985), the outcome of a Wald test depends on the particular parameterization of the hypothesis under study.

[8] Recall that individuals equal or above 70 years old are imposed to be retired because there is no enough information for the model to correctly predicts behavior

Table 9: Actual and predicted nonparticipation rates (percentages)

Age 2007	Actual	Predicted
60-64	21.2	38.6
65-69	55.9	45.8
70+	66.7	100.0

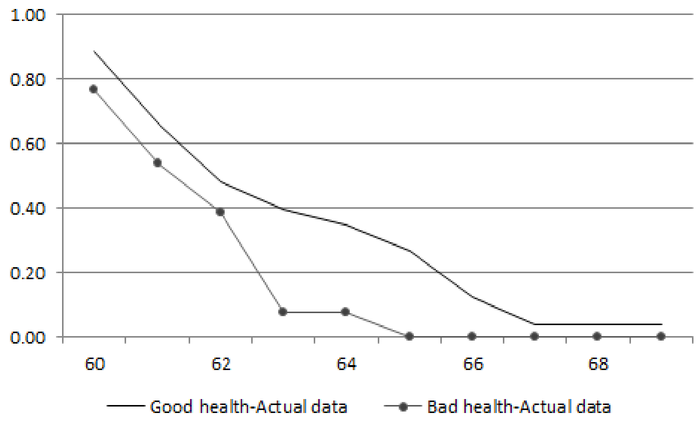

Figure 3: Cumulative distribution of labor force participation by health status: Actual

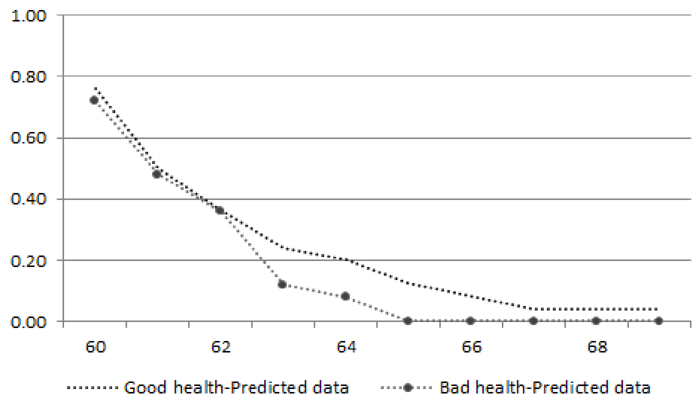

Figure 4: Cumulative distribution of labor force participation by health status: Predicted

retiring and enjoying more leisure given their health condition. This behavior is quite well predicted by the model (see Figures 3 and 4, and Table 10 for nonparticipation rates). Considering our definition of bad health, our results also indicate that those who are receiving disability benefits are less likely to participate in the labor force across all ages and they start to retire much earlier than those who do not receive this benefit. Heyma (2004) finds similar results in the sense that restricting eligibility for disability benefits hardly increases Dutch elderly labor force because individuals in bad health still retire early through the unemployment route.

5 Simulation

On July 12, 2012, a broad coalition of Dutch political parties passed a law that aims at the gradual increase of the legal retirement age from 65 to 67. In this section, we use the estimated model to

Table 10: Actual and predicted non employment rates by health status (percentages)

Age 2007	Good health		Bad health	
	Actual	Predicted	Actual	Predicted
60-64	19.2	34.6	52.2	100.0
65-69	56.1	43.9	50.0	100.0
70+	66.7	100.0	-	-

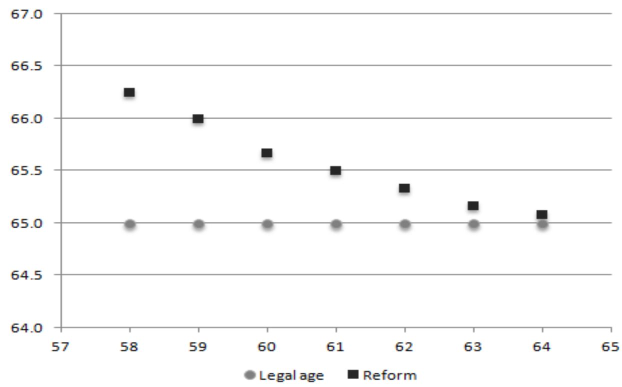

Figure 5: Legal retirement age in pre- and post-reform scenarios

simulate the effects of this soon to be implemented policy reform. The content of the reform has been extensively discussed in the Netherlands in the past few years.

The basic idea of the reform is that the retirement age increases more for younger individuals than for older ones. For individuals who are close to their retirement (specifically, those who are aged 61 or more) it is planned that they have to wait a few months after their 65th birthday before they can claim a complete public pension, whereas for younger individuals (aged 60 or less) it is planned that they have to wait more time; around one year after their 65th birthday. For instance, an individual who is 64 years old (and therefore close to retirement) has to wait only one more month, whereas an individual who is 59 years old (and therefore relatively far from retirement) will have to wait one full year.[9] Figure 5 shows the proposed legal retirement ages associated with the pension reform.

We simulate the reform by evaluating its effect on the labor force participation of the workers in the sample. To model the reform we assume that an individual's retirement behavior is solely affected by the economic incentives implied by the regulation of the legal retirement age. This means that the only change compared with our baseline model (discussed in the previous section) is that we increase the age at which individuals start to receive the complete public pension. We thus leave the level of the pension benefits unaffected. To be consistent, individuals also contribute to the public pension system until they are able to receive the pension. We also assume they continue accumulating entitlements in the second pillar. In our model, the impact of the reform takes place through the non-labor income in the budget constraint, which in turn affects the decision whether to work or retire.

Following the current reform rules, we find that there is little change in individual retirement behavior: individuals postpone their retirement by only three months on average. Table 11 shows that in absence of the reform (baseline scenario), the expected retirement age is 65.65, whereas in the simulation of the reform that gradually increases the retirement age, the expected retirement age increases to 65.88. This result slightly varies by gender. Men's optimal behavior is to continue working for two more months whereas women's is to work for three more months. In the baseline scenario, women expect to retire at 65.24 years old whereas with the reform they decide to retire at 65.47 years old on average, which leads to a difference in the expected retirement age of three months. In the case of men this difference is a bit smaller, 0.15 years or two months, which implies that the reform has slightly stronger effects on women's retirement behavior relative to men's.

Table 12 gives the expected retirement age by health status computed on the basis of the baseline and simulation results. According to the simulation scenario, individuals in good health decide to work four more months than in the baseline scenario. As we already explained above, the health status turns out to be a key determinant of the labor supply response, so by postponing the age of retirement, the alternative to retire becomes less attractive in comparison with the alternative to work. For individuals in bad health, the simulation results show that the alternative to retire continues to be very attractive, so they do not decide to postpone retirement with virtually no

[9] The retirement age is increased even more for individuals below 58 years old. However, these cohorts are not included in the sample.

Table 11: Expected retirement age by gender

Singles	Baseline (65)	Reform (Monthly increase)
Total	65.65	65.88
	(2.52)	(2.57)
Men	66.31	66.46
	(3.10)	(3.08)
Women	65.24	65.47
	(1.97)	(2.06)

Note: Standard deviations in parentheses.

Table 12: Expected retirement age by health status

Singles	Baseline (65)	Reform (Monthly increase)
Good health	66.27	66.57
	(1.93)	(1.86)
Bad health	61.00	61.00
	(1.22)	(1.22)

Note: Standard deviations in parentheses.

change in their expected retirement age as a consequence.

Figure 6 plots the cumulative probability function of labor force participation for the reform scenario. This gives us an idea about how labor supply might be affected by the implemented reform. If we focus on the dotted line (reform scenario), then we observe that increasing the legal retirement age on a monthly basis leads to a slightly upward shift of the probability function between 64 and 67 years old, which means that the reform has a small positive impact on labor force participation. By relatively reducing the value of the alternative to retire in comparison with the alternative to work, the current reform creates incentives for individuals to choose to work a little more.

This type of policy change has typically been studied from the point of view of its labor supply consequences. According to Gruber and Wise (2005) an increase in eligibility ages naturally increases labor force participation, especially of older employees. So, it is not surprising to observe that individuals keep on working for longer time, which is also the case in our simulation of the planned pension reform. Other authors have found similar results as well. In their exercise for Sweden, Karlstrom, Palme and Svensson (2004) find also a upward shift in the cumulative distribution function of labor supply meaning that more individuals decide to continue working instead of retiring. In their model, however, they include both the effect of delaying the age of retirement and the change in the pension benefit. Van der Klaauw and Wolpin (2008) also find that annual hours

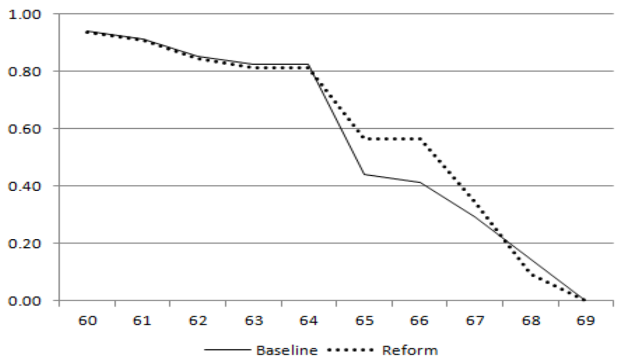

Figure 6: Cumulative probability of labor force participation

of work and earnings increase for single individuals when benefits are postponed until 70 years old. However, in their model, since they consider both consumption and savings decisions, consumption and net assets also fall after the age of 62, meaning that the adjustment of labor supply is not enough to compensate the fall in consumption and savings levels produced by postponing pension benefits. Rust and Phelan (1997) perform a counterfactual prediction of the impact of eliminating Social Security. They find that the optimal decision without benefits implies that individuals continue to working after the age of 65 and they interpret these results as evidence that Social Security rules have a strong incentive effect on individual behavior. Despite these similar results, there is also evidence of no or only a small impact on labor supply when the model accounts for borrowing constraints. In his analysis, French (2005) finds that shifting the early retirement age from 62 to 63 has almost no effect on labor force participation. The reason is because borrowing constraints do not bind for several individuals, so there is little disincentive effect on working produced by a retirement age of 62, so the change to 63 has a limited impact on retirement behavior.

In our case, we find that the reform has an impact on the effective retirement age but this impact is very small. We also find that healthy individuals are those who are most affected by the current reform and they decide to continue working for longer time. The small impact is partly explained by the design of the reform which basically aims to affect individuals gradually and more significantly to younger generations whereas leave (almost) unaffected the older ones (only individuals younger than 59 have to wait one year or more to receive benefits).

6 Conclusion

This paper uses a dynamic programming model to provide an empirical analysis of how pension system rules affect the labor force participation of older single individuals in the Netherlands. We model these labor supply responses in the presence of uncertainty on income, health status and lifespan. We use a panel based on administrative data which has the main advantage of having new and accurate information on pension entitlements. As far as we know, an analysis based on administrative micro data on pension entitlements has not been undertaken yet, because it has become only recently available in the last years.

Our model is able to capture the main patterns of the sample data. Overall we observe that as individuals become older they are more willing to sacrifice income for leisure which translates into more individuals deciding to leave the labor force. We also find that health status is a key determinant for the labor force participation (or retirement behavior). Not very surprisingly perhaps, individuals in good health tend to stay longer in the labor force than individuals in bad health.

The paper also discusses the simulation of the impact on retirement behavior of a soon to be implemented reform that gradually increases the legal retirement age. The idea of the reform is that the legal retirement age increases more for younger individuals than for older ones. Our simulation results show that there is an impact on the effective retirement age but this impact is very small. An individual's optimal behavior is to postpone retirement by around 3 months on average and there are differences by one's health status. Individuals in good health are those who are most affected by the reform and they decide to continue working for somewhat longer.

Although the model performs relatively well given the available information, several extensions of the model seem worthwhile to explore. Data limitations prevented us though to analyze these issues further. Firstly, one could consider more alternative labor supply choices. Some individuals might want to work only part-time to enjoy more leisure especially when they are close to retirement (see, e.g., Kantarci and van Soest, 2008, on the issue of partial retirement). This might bring better approximations to actual behavior of older individuals. Secondly, one could consider more health status categories. Our current definition of health status might underestimate the proportion of unhealthy individuals, so a more detailed health variable might be useful. Third, the savings decision as a second control variable could be introduced. To allow for savings decisions is an important aspect in life cycle models, since they allow individuals to smooth consumption and protect themselves against future shocks. The level of one's savings might also be a determinant for the early retirement decision in the sense that wealthier individuals might have enough savings to compensate reductions in their retirement income when leaving the labor force before the legal retirement age. Finally, another extension of our model could be to consider the intra-household decision making process in which we can distinguish between decisions made by one-person (single) and two-person (couple) households, where the latter were not included in the current analysis.

References

[1] Adda J. and Cooper R. (2003). *Dynamic Economics. Quantitative Methods and Applications.* The MIT Press.

[2] Aguirregabiria V. (2010). "Another look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior". *Journal of Business and Economics Studies*, 28, 201-217.

[3] Browning, M. and P.-A. Chiappori (1998), "Efficient intra-household allocations: a general characterization and empirical tests", *Econometrica*, 66, 1241-1278.

[4] Cherchye, L., B. De Rock and F. Vermeulen (2009), "Opening the black box of intra-household decision-making: Theory and non-parametric empirical tests of general collective models", *Journal of Political Economy* 117, 1074-1104.

[5] Cherchye, L. and F. Vermeulen (2008), "Nonparametric analysis of household labor supply: goodness of fit and power of the unitary and the collective model", *Review of Economics and Statistics*, 90, 267-274.

[6] Chiappori, P.-A. (1988), "Rational household labor supply", *Econometrica*, 56, 63-89.

[7] Eckstein Z. and Wolpin K. (1989). "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey". *The Journal of Human Resources*, 24, 562-598.

[8] Eckstein Z and Lifshitz O. (2011), "Dynamic Female Labor Supply". *Econometrica*, 79, 1675-1726.

[9] European Commission (2012). *White Paper. An Agenda for Adequate, Safe and Sustainable Pensions.* Brussels.

[10] Ewisk C. (2005). "Reform of Occupational Pensions in the Netherlands". *De Economist*, 153, 331-347.

[11] French E. (2005). "The Effects of Health, Wealth and Wages on Labor Supply and Retirement Behavior". *Review of Economic Studies*, 72, 395-427.

[12] French E. and Jones J. (2011). "The Effects of Health Insurance and Self-Insurance on Retirement Behavior". *Econometrica*, 79, 693-732.

[13] Gregory, A. and M. Veall (1985), "Formulating Wald Tests of Nonlinear Restrictions", *Econometrica*, 53, 1465-68.

[14] Gustman A. and Steinmeier T. (1986). "A Structural Retirement Model". *Econometrica*, 54, 555-584.

[15] Gustman A. and Steinmeier T. (2000). "Retirement in a Family Context: A Structural Model for Husbands and Wives". *Journal of Labor Economics*, 18, 503-545.

[16] Gustman A. and Steinmeier T. (2004). "Social Security, Pensions and Retirement Behavior within the Family". *Journal of Applied Econometrics*, 19, 723-738.

[17] Gustman A. and Steinmeier T. (2006). "Social Security and Retirement Dynamics". Michigan Retirement Research Center. University of Michigan. Working paper 2006-121.

[18] Gruber J. and Wise D. (2005). "Social Security Programs and Retirement around the World: Fiscal implications". National Bureau of Economic Research, Working Paper 11290.

[19] Heyma A. (2004). "A Structural Dynamic Analysis of Retirement Behavior in the Netherlands". *Journal of Applied Econometrics*, 19, 739-759.

[20] Kantarci, T. and A. Soest (2008), "Gradual Retirement: Preferences and Limitations", *De Economist*, 156, 113-144.

[21] Karlstrom A., Palme M. and Svensson I. (2004). "A Dynamic Programming Approach to model the Retirement Behavior of Blue-collar workers in Sweden". *Journal of Applied Econometrics*, 19, 795-807.

[22] Keane M. and Wolpin K. (2009). "Empirical Applications of Discrete Choice Dynamic Programming Models". *Review of Economics Dynamics*, 12, 1-22.

[23] Michaud P. and Vermeulen F. (2011). "A Collective Labor Supply Model with Complementarities in Leisure: Identification and Estimation by means of Panel Data". *Labour Economics*, 18, 159-167.

[24] OECD (2009). *Ageing and Employment Policies - Statistics on average effective age of retirement.* http://www.oecd.org/document/47/0,3746,en_2649_33927_39371887_1_1_1_1,00.html. Downloaded in February, 2012.

[25] Rust J. (1989). "A Dynamic Programming Model of Retirement Behavior" in the Economics of Aging, ed. by D. A. Wise. Chicago: University of Chicago Press, 359-398.

[26] Rust J. (1990). "Behavior of Male Workers at the end of the Life-Cycle: An Empirical Analysis of States and Controls" in *Issues in the Economics of Aging*", ed. by D. A. Wise. Chicago: University of Chicago Press, 317-375.

[27] Rust J. (1994). "Structural Estimation of Markov Decision Processes", in *Handbook of Econometrics.* Vol. 4, ed. by R. Engle and D. McFadden, 3082-3139.

[28] Rust J. (2006). "Dynamic Programming". University of Maryland.

[29] Rust J. and Phelan C. (1997). "How Social Security and Medicare affect Retirement Behavior in a World of Incomplete Markets". *Econometrica*, 65, 781-831.

[30] van der Klaauw W. and Wolpin K. (2008). "Social security and the retirement and savings behavior of low-income households". *Journal of Econometrics*, 145, 21-42.

[31] Wooldridge J. (2002). "Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity". *Journal of Applied Econometrics*, 20, 39-54.