Welfare analysis using life cycle paths of Long Term Care spending. Preliminary Draft.

Bram Wouterse[1,2,5], Arjen Hussem[1,3,5], and Albert Wong[4,5]

[1]CPB Netherlands Bureau for Economic Policy Analysis, The Hague, The Netherlands
[2]Erasmus School of Health Policy & Management, Erasmus University Rotterdam, Rotterdam, The Netherlands
[3]PGGM, Zeist, The Netherlands
[4]National Institute of Public Health and the Environment, Bilthoven, The Netherlands
[5]Network for Studies on Pensions, Ageing and Retirement, Tilburg, The Netherlands

October 4, 2017

Abstract

Long term care expenditures are a major financial risk for the elderly. The ageing of the population makes modeling the distribution of these expenditures over life and their effect on financial decision making of the elderly of great importance. Modeling long term care expenditures over the lifecycle is challenging because of their very uneven distribution, with a small proportion of elderly experiencing very high costs. We use a semi-parametric nearest-neighbor approach to estimate lifecycle paths of long term care spending. These paths are then used as inputs in a stochastic lifecycle decision model for retirees. We use an extensive administrative dataset to estimate the lifecycle paths. The data include information on long term care spending, household status, income, and wealth for the entire Dutch population. We apply the lifecycle model to the Dutch case of income and wealth dependent co-payments. We find that, first, even with income- or wealth-dependent co-payments, losses in both average and certainty equivalent consumption are highest for the groups with the least spending power. Replacing a co-payment that does not depend on income or wealth with income or income- and wealth dependent co-payments that raise the same revenue, does decrease welfare loss for groups with low spending power, while only leading to small increases in welfare loss for the highest spending power groups. Second, we find

that even in the case of a 30 percent co-payment rate, full annuitization of pension wealth is still optimal.

1 Introduction

The ageing of the population, and the resulting increase in the number of elderly with disabilities, has put the provision and financing of long term care (LTC) at the forefront of the policy debate. While some countries are only now starting to set up forms of social LTC insurance (Colombo and Mercier, 2012), other countries, such as the Netherlands, are trying to keep collective LTC systems sustainable by increasing the role of private financing through limiting overage and increasing co-payments.

Private financing of LTC costs affects the saving behavior of pensioners and changes the income distribution. First, elderly are confronted with the risk of potentially large costs during their remaining life for which they need liquid assets. This means that pensioners might increase their savings and choose to annuitize less of their pension wealth. Second, private financing affects the redistribution of income across pensioners with low and high means. Elderly with low incomes, on average, use more LTC than those with high incomes. Collective insurance thus redistributes income from high to low income groups, while increasing the role of private financing limits this redistribution. An assessment of the redistributive effects of private LTC financing should include the welfare losses due to increasing financial risks and should take the effects on behavior into account.

An important challenge in assessing the effects of private LTC financing on saving behavior is the modeling of the lifecycle distribution of LTC costs. The distribution of LTC costs is hard to model parametrically: a large part of the population does not have any LTC costs at all, while a small group of individuals experiences very high costs persisting over many years. Existing approaches use autoregressive models (De Nardi et al., 2010; French and Jones, 2004) or Markov models (Ameriks et al., 2011) to estimate time dynamics in LTC costs. However, these models require a variety of assumptions that most often cannot be justified on the basis of the data alone (Wong et al., 2016). As an alternative, we use the semi parametric nearest neighbor approach developed in (Wong et al., 2016; Hussem et al., 2016). The main advantage of this approach is its flexibility.

We use the nearest neighbor algorithm to estimate 20,000 synthetic lifecycle paths using Dutch data. These paths contain yearly LTC costs and income. We use the estimated paths as inputs in a stochastic lifecycle decision model for retirees. This model determines optimal consumption and saving behavior of elderly for different levels of initial wealth and pensions, taking into account their financial risk under different co-payment regimes for long term care. We use a simulation based algorithm developed by Koijen et al. (2010) to solve the model.

With the model we analyze the effects of different forms of income- and wealth-dependent co-payments that are implemented or considered by policy makers in the

Netherlands. We focus on the effects on welfare (certainty equivalent consumption) of the elderly across income groups. We also assess the effect of co-payments on the optimal annuitization rate of pension wealth.

2 The Dutch Long Term Care system

Unlike many other developed countries, the Netherlands has a long tradition of providing collective LTC through social insurance. It also has one of the most extensive collective LTC arrangements in the world (Colombo and Mercier, 2012). Almost all health spending is financed through collective social insurance or taxes. Since the LTC reform of 2015 (Maarse and Jeurissen, 2016), LTC is financed through two programs. Relatively light forms of care, such as personal assistance, are financed through the Social Support Act (WMO). The provision of this type of care is a responsibility of municipalities. They get a financial contribution out of the general means of the national government, depending on the composition of their population. Intensive forms of care for patients who are in permanent need of care (often in an institutional setting) are financed through a social insurance called the Long Term Care Act (WLZ).

In this paper, we use data from before the reform. Before 2015, both home care and institutional care was financed through a collective insurance called the AWBZ. The premium for the AWBZ was collected through the income tax and is a percentage of taxable income (including pension income) in the first and second income brackets. A small part of the costs were covered by out-of-pocket payments.

In the current system, out-of-pocket payments play a larger role. They are income- and wealth dependent and also differ according to type of care and living situation. They can vary for home care from 247 euros per year plus 15 percent of spending power minus a 16.456 euros threshold to the maximum of the actual costs. For institutional care these payments can be 12.5 percent of spending power for people with a partner living at home to 75 percent of net income with a maximum of 26.983 euros per year (all amounts 2014). The spending power is defined as gross income plus 12 percent of financial wealth. For wealth a threshold is used of 21.139 euros which can go up to 49 euros.123 for people with low incomes.

3 Data

3.1 Source data

We use administrative data on LTC use from the Dutch Central Administrative Office (CAK). These data cover the period 2004-2006. The data include information on all formal LTC use in the Netherlands. The data contain information on the type of care (institutional care, nursing home care, personal home care, and support) and the amount of care used (in days for institutional care, and in hours for home care). We derive costs of LTC from use in hours/days in the CAK database and the tariffs provided by the Dutch Health Authority (NZA) for extramural care and derived from the CAK and CVZ annual reports for intramural care.

The LTC data is combined with other datasets through linkage the Dutch Municipal Register which contains basic information on everyone enlisted in a Dutch municipality. From this register we obtain date of death, age, sex and marital status. Additionally, we link the data to the Death Causes Registry for the entire Dutch population, and administrative data and survey data on all income sources from Statistics Netherlands (Regionaal InkomensonderzoekRIO), and marital status from Statistics Netherlands (Gehuwdheidsbestand, VRLHUWELIJKS-GESCHIEDENISBUS).

4 Long term care spending over the lifecycle

4.1 The nearest neighbor algorithm

As inputs for our model we use synthetic lifecycle paths estimated by Hussem et al. (2016). These lifecycle paths of LTC costs have been estimated using a nearest neighbor resampling method. Although there are many specific implementations, the idea behind nearest neighbor matching (NNM) is that we want to match an observation from one group (for instance a treatment group) to the most similar observation from another group (for instance the control group). NNM uses a distance metric to determine, based on the covariate values, which observation form the other group is the nearest. Some of the first implementations of NNM in a time series or panel context are by Farmer and Sidorowich (1987) and Hsieh (1991). The LTC paths we use have been estimated using the approach developed by Wong et al. (2016) who have implemented a nearest neighbor resampling method to estimate lifecycle paths of curative care costs.

The basic idea of the NNM algorithm is that we want to simulate N individual lifecycle realizations of LTC spending. Each simulated life cycle will consist of an age series $Z_i = \{Z^i_{a=0}, Z^i_{a=1}, ..., Z^i_{a=A_i}\}$. Z^i_a is a vector containing LTC spending and other variables of interest (for instance income) of individual i at age a. $a = 0$ denotes the starting age and A_i is the age of death. Our data is a relatively short panel containing observed values $Y^j_{a,t}$ for individuals $j = 1, .., J$ over time periods $t = 1, .., T$. The algorithm works as follows. Suppose we already have a simulated lifecycle path for an individual up to age A: $Z^i = \{Z^i_0, Z^i_1, .., Z^i_A\}$. To extend this lifecycle path to age $A+1$ we consider all individuals in our data who have age $A+1$ in period T. We pick the individual who's life history over the last p age years $Y^j = \{Y^j_{A-p+1,1}, ..., Y^j_{A,T-1}\}$ is most similar to $\{Z^i_{A-p+1}, ..., Z^i_A\}$. Note that, because we want to extend the lifecycle by one period, and the time length of the panel is T, we can use a maximum age lag p of $T-1$ years. When we have picked a individual j, we use $Y^J_{A+1,T}$ as our simulated realization of Z^i_{A+1}. Then, to obtain a realization for age $A+2$ we can repeat the procedure using all individuals in the data with age $A+2$ at time T, matching on the life history over ages $A-p+2$ to $A+p+1$. This procedure is repeated until i is matched to an individual who dies in period T.

To initialize the algorithm, we use all individuals with age $a=0$ at time T. For these individuals we have data on Y over $T-1$ ages before the starting age $a=0$. We include the information on the last $p-1$ ages in the simulated lifecycle path, so we start with $Z^i = \{Z^i_{-p+1}, , Z^i_0\}$.

To match a simulated lifecycle paths to an observation from the data we use k-nearest neighbor matching. We measure the distance between two p-long blocks $\mathbf{z}$ and $\mathbf{y}$ using a distance measure $d(\mathbf{z},\mathbf{y})$. Different distance measures, such as Euclidean distance, can potentially be used. We use the Mahalanobis measure that corrects for scale differences across the components of $\mathbf{y}$ and takes correlation between those components into account. The measure is defined as

$$d(\mathbf{z},\mathbf{y}) = \sqrt{(\mathbf{y}-\mathbf{z})^T\Sigma^{-1}(\mathbf{y}-\mathbf{z})}, \tag{1}$$

where Σ is an estimate of the covariance matrix of $\mathbf{y}$. Out of the k-nearest neighbors, one neighbor is randomly drawn. We use $k=2$.

Since Hussem et al. (2016) have data for three years (2004-2006), they use two lags ($p=2$). They stratify the data by sex, age, and household status (e.g. men are only matched to other men) and match on income and LTC expenditures. They simulated 10,000 paths for women and 10,000 for men. These paths start at birth, but we start the simulation exercise at the pension age (65) and thus select those individuals who are alive at 65. A more detailed description of the matching procedure can be found in (Hussem et al., 2016).

4.2 Including information on wealth

The lifecycle paths that we currently use does not include information on wealth. Therefore, we include wealth ex-post (similarly to Hussem et al. (2017)). We use the relationship between income and financial wealth for all Dutch singles 70 year olds in 2012 [1]. We group these individuals in income and wealth quintiles. This gives us 25 combinations of income and wealth groups. For each of these combinations we estimate the average financial wealth.

We also group each individual in the lifecycle paths to an income quintile, based on his or her income at age 70. Then we can assign initial wealth at 70 using the relationship between income and wealth we have estimated for the whole Dutch population of 70 year olds. We use each individual from the lifecycle paths five times, using the averages of the five wealth quintiles as starting values. We weight the outcomes based on the relative sizes of each combination of income and wealth in the Dutch population.

[1]Data obtained from CBS

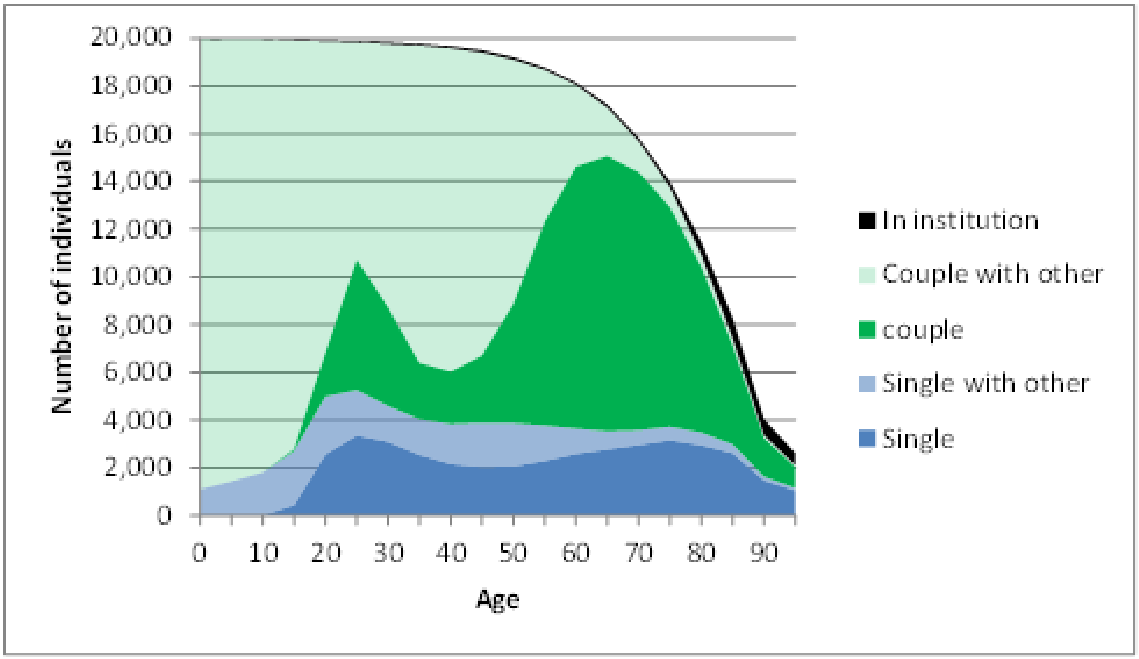

Figure 1: Household type composition of the simulated population by age.

4.3 Estimation results

As described in (Hussem et al., 2016) the resulting lifecycles can be interpreted as a cohort of newborns in 2006 and how they would accumulate income and LTC costs during an entire lifetime, under assumption that the levels of as well as transition probabilities of income and costs remain constant at those observed in the calendar year 2006, as well as the state of the long-term care system in 2006. The method replicates the 'true' transitional probabilities (e.g. death rate or probability of becoming a LTC-user), and the distribution of income and LTC from the source data conditional on the variables considered. As the population composition (e.g. by age and household type) in the synthetic life-cycle paths differs from the source data the total annualized amounts in the life-cycle paths are not equal to those amounts in the source data. The life-cycle paths are representative for a stationary population based on the behavior and institutions of the 2004-2006 period. Future changes in mortality and LTC use are therefore not taken into account explicitly .

Figure 1 shows the resulting 'stationary' population, consisting of people living as a single, couple or in an institution. (In this version of the paper we include individuals from age 65 onwards, and we ignore household status (treat everyone as single).)

Table 1 shows the average annual and lifetime LTC costs in the lifecycle paths.

Figure 2 shows lifetime LTC costs, lifetime income and initial wealth across spending power quintiles. Spending power is defined as the certainty equivalent consump-

Table 1: LTC costs in the lifecycle paths, from age 65 onwards

	average	std/average
per year	1,320	5.8
lifetime	34,323	2.5

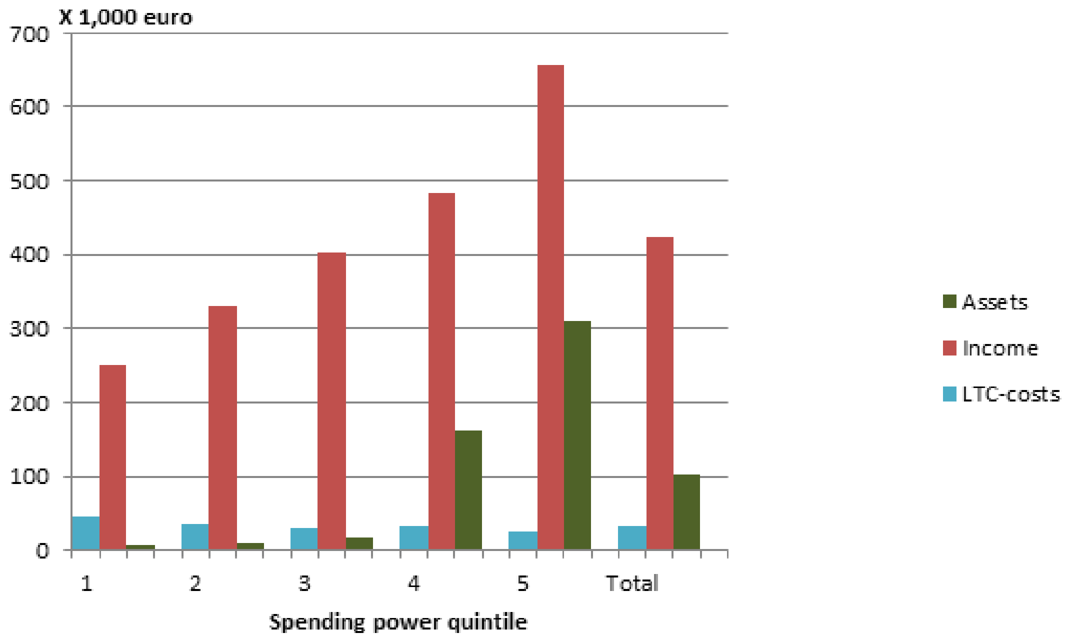

Figure 2: The relationship between LTC and spending power on age 65, per quintile of spending power.

tion in the baseline (without any co-payment). [2]

5 A model of lifecycle consumption after retirement

5.1 The model

The estimated lifecycle paths provide a semi-parametric distribution function of LTC expenses and mortality. We implement a standard life cycle model with rational and forward looking individuals to model consumption and saving behavior given this distribution. Mortality risk and the development of LTC spending over life are based on the lifecycle paths. Consumption and saving behavior, conditional on initial wealth, are determined by the lifecycle model. To be able to use the semi-parametric lifecycle paths in the optimization problem, we implement a simulation-based maximization

[2] See Section 5.1 for the definition of certainty equivalent consumption.

algorithm developed by Koijen et al. (2010).

The basic model

We model the consumption and savings decisions of individuals after retirement. An individual starts at the pension age, $t = 0$, with initial wealth W_0. He uses this wealth to finance consumption over the remaining time periods $t \in 1, ..., T$. The individual faces uncertainty about the duration of remaining life and the amount of LTC co-payments. We assume that the individual only derives utility out of consumption (later we will introduce a bequest motive). The individual wants to maximize his expected utility over remaining lifetime. With a time-separable utility function the individual's maximization problem then is:

$$\mathrm{E}(V_0) = \mathrm{E}\left[\sum_{t=0}^{T}\left(\beta^t u(c_t)\prod_{s=0}^{t} p_s\right)\right], \tag{2}$$

withp_s the probability of surviving period s, and β the discount factor.

Each period, the individual has to choose the amount of his wealth W_t he wants to consume now (c_t), and the amount he want to save for later (m_t). He faces the following annual budget constraint:

$$c_t + m_t + h_t = W_t. \tag{3}$$

We impose the borrowing constraint $W_t \leq 0$.The timing is such that first h_t has to be paid, and then the individual decides how to divide his remaining wealth between c_t and m_t. We treat the level of private LTC spending, h_t, as given: the individual does not weight utility gained from h_t against utility from c_t, but instead h_t is an exogenous shock in W_t.

The utility function is defined as a standard CRRA function:

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma}. \tag{4}$$

This implies that individuals want to smooth consumption evenly over the lifecycle. Wealth growths with the risk free interest rate r, so that

$$W_{t+1} = m_t r. \tag{5}$$

Extensions

We extend the model in three ways. First, we include a bequest motive. We assume that the individual derives utility from the level of wealth W_{death} he leaves at time of death. We use the same bequest function as De Nardi et al. (2010):

$$g(W_{death}) = \theta\frac{(W_{death} + \xi)^{1-\gamma}}{1-\gamma}, \tag{6}$$

where θ determines the strength of the bequest motive and ξ the curvature of the bequest function.

Second, we allow for (partial) annuitization. Individuals can choose to annuitize a share α of their initial wealth at the age of retirement. In that case, they will receive an actuarially fair constant annual income y_α:

$$y_\alpha = \frac{\alpha W_0}{\sum_{t=0}^{T} \left[(1+r)^t \prod_{s=1}^{s=t} p_s \right]}. \tag{7}$$

The budget constraint then becomes

$$c_t + m_t + h_t = W_t + y_\alpha. \tag{8}$$

The individual's maximization problem now not only involves choosing consumption in each period, but also deciding on the share of initial wealth he wants to annuitize.

Third, we allow the level of co-payments to depend on wealth and pension income. Let hc_t be the total LTC spending an individual needs in period t. This spending is exogenous. Private LTC spending, h_t, is not necessarily equal to hc_t, but depends on the co-payment rules set by the government. We use the following general co-payment rule:

$$h_t = \min(\tau hc_t, \nu_y y_\alpha + \nu_w W_t). \tag{9}$$

The government sets the parameters τ, ν_y, and ν_w. The parameter τ determines what share of total health care spending has to be paid by the individual himself. The parameters ν_y and ν_w determine the maximum share of income and wealth that have to be spend on co-payments. The way the government sets the co-payment rules affects the optimization problem of the individuals. Especially, when $\nu_y > 0$ or $\nu_w > 0$, co-payments are no longer fully exogenous since they depend on the annual savings and the annuitization share of initial wealth chosen by the individual.

Outcome

The main outcome measure we will use to present welfare effects of different financing schemes across groups is certainty equivalent consumption:

$$CEC = u^{-1} \left(\frac{\mathrm{E}(V_0)}{\sum_{t=0}^{T} \beta^t \left(\prod_{s=0}^{t} p_s \right)} \right) \tag{10}$$

More specifically, we will show changes in CEC compared to a baseline financing scheme without any co-payments. For group g and alternative a we will show

$$\frac{CEC_{g,a} - CEC_{g,0}}{CEC_{g,0}}, \tag{11}$$

with $CEC_{g,0}$ is the CEC for group g under the baseline scheme.

5.2 Numerical approach

The basic model

The individual's maximization problem can be solved using a dynamic programming approach. In this approach, the lifecycle optimization problem is divided into smaller yearly optimization problem. The algorithm starts at the last time period T, and is then solved backwards recursively using Bellman equations. In each period the optimization problem can be written as

$$\max\left[\mathrm{E}(U_t) = u(c_t) + \mathrm{E}_t(V_{t+1}(m_t)\right]. \tag{12}$$

We solve this problem using the approach developed by Koijen et al. (2010). This approach has been applied to LTC financing for the U.S. by Peijnenburg et al. (2015). The approach combines the method of endogenous gridpoints (Carroll, 2006) with a simulation based approximation of the expected values (Brandt et al., 2005). A simulation based approach is well suited to use in combination with the lifecycle paths. Other approaches generally approximate the stochastic processes by a limited number of discrete states. Instead, the method of Koijen et al. (2010) allows us to directly use the lifecycle paths as inputs.

Often, the maximization problem in Equation (12) is solved for a finite number of possible values of wealth (on a grid) *at the beginning* of a period W_t. The solution for other values of W_t is then obtained by (linear) interpolation between the gridpoints. Instead, in the endogenous gridpoints method, a grid is used for m_t the amount of wealth *at the end* of period t *after* consumption and health spending. Optimal consumption c_t^* in period t is then determined given the amount of wealth that is left at the end of t. This method avoids the need for numerical optimization to determine c_t^*. An "endogenous" gridpoint for W_t is determined afterwards by summing up m_t, optimal consumption c_t^*, and LTC spending h_t.

To see how the algorithm works, let's start in the final period T. If an individual is still alive at period T, he consumes all his remaining wealth. So optimal consumption is given by:

$$c_T^* = W_T - h_T, \tag{13}$$

and $u_T^* = u(c_T^*)$.

For period $T-1$, we define a fixed grid with $j = 1,..,J$ gridpoints $m_{j,T-1}$ for wealth after consumption and LTC spending. Because the wealth level after consumption in $T-1$ is already known, the corresponding level of consumption c_{T-1}^* is given by the first order condition:

$$c_{j,T-1}^* = \left(\mathrm{E}(\beta c_T^{*-\gamma} r | m_{j,t-1})\right)^{-\frac{1}{\gamma}}. \tag{14}$$

This is the standard Euler condition implying that individuals want to smooth consumption evenly over remaining lifetime.

To determine $\mathrm{E}(\beta c_T^{*-\gamma} r | m_{j,t-1})$ we use a simulation approach. The idea is similar to a Monte-Carlo approach: the lifecycle-paths give us a large number of random draws from the stochastic process determining mortality and LTC spending. The expected value can then be estimated by averaging over these draws. Note that for each individual (path) the realized consumption in T conditional on m_{T-1} is given by the fact that (if still alive) the individual will consume all the wealth he has left: $(W_T | m_{j,t-1}) = m_{j,t-1}$ and $(c_T^* | m_{j,t-1}) = m_{j,t-1} - h_T$. To determine the expected value we regress these realizations of consumption at T on (a polynomial expansion) of the state variables (background characteristics and LTC spending) at time $T-1$. This gives

$$\mathrm{E}(\beta c_T^{*-\gamma} r | m_{j,T-1}) \simeq \theta f(x_{T-1}), \tag{15}$$

with x_{t-1} a vector with the state variables in period $t-1$ and $f()$ a polynomial expansion of some order.

The expected values are then obtained by using the predictions from the regression model (conditional on the state variables), and this also provides the optimal level of consumption in period $T-1$ given m_{T-1}. We have to perform this procedure for each gridpoint, and thus have to run a regression for each gridpoint.[3]

Now that we have the optimal consumption levels $c_{j,T-1}^*$ for each fixed gridpoint for wealth $m_{j,T-1}$ *at the end* of period $T-1$, we can create a grid with endogenous gridpoints for wealth $W_{j,T-1}$ *at the beginning* of period $T-1$. These are given by

$$W_{j,T-1} = c_{j,T-1}^* + h_{T-1} + m_{j,T-1}. \tag{16}$$

The level of initial wealth at the beginning of $T-1$ is determined by the level of wealth that is saved at $T-2$. So we now have the set-up for the iterative algorithm. Because the endogenous gridpoints $W_{j,T-1}$ are not necessarily the same as the gridpoint we use for m_{T-2}, we use linear interpolation to obtain the levels of optimal consumption in $T-1$ belonging to the gridpoints $m_{j,T-2}$ for wealth saved at the end of period $T-2$ for each individual (path). This then, allows us to estimate expected optimal consumption at $T-1$ using the same regression as in Equation (15). This gives optimal consumption in $T-2$. And this in turn determines the endogenous gridpoints for W_{T-2}. We can iteratively perform this algorithm for periods down to $t = 1$. In the end, we have the optimal consumption at each period for the endogenous gridpoints $W_{1,t}, ..W_{J,T}$. We have a series of (different) endogenous gridpoints and optimal consumption for each individual (path) i.

Now that we have the consumption rules, we can use these to simulate consumption and saving behavior of the individuals in the lifecycle sample. We do this by assigning an amount of initial wealth at the start of the first period to each individual. We can then simulate forward.

[3] This is also one of the reasons we use the endogenous gridpoints method. Else, we would have to run a regression at each iteration step of the numerical maximization of Equation (12) for each gridpoint in each period $m_{j,t}$.

Table 2: Values of parameters

r	1.015
β	0.985

Extensions

Including a fixed pension income in the optimization procedure is straightforward. The optimization problem does not change, only the budget restriction has to be adjusted to include the income stream. To estimate optimal annuitization shares we run the algorithm over a grid of values for α between 0 and 1. The value of α determines which share of each individual's actual initial wealth (taken from the lifecycle data) is annuitized. Comparison of certainty equivalent consumption over the values of α then gives the optimal annuitization share.

5.3 Policy variants

The effects of (income- and wealth-dependent) co-payments on welfare

In the Dutch system co-payments depend on both income and wealth. To study the effect of introducing or abolishing the income and/or wealth dependency of the co-payments we define the following three policy variants. These variants are constructed in such a way that they all finance 13 percent of total LTC costs for the whole population.

- Variant 1. A co-payment with a maximum of 13 % of income plus 12 % of wealth.
  $\tau = 1$, $\nu_y = 0.12$, $\nu_w = 0.12$
- Variant 2: A co-payment with a maximum of 20 % of income.
  $\tau = 1$, $\nu_y = 0.2$, $\nu_w = 0$
- Variant 3: A co-payment of 13 % of LTC costs.
  $\tau = 0.2$, $\nu_y = 1$, $\nu_w = 1$

In all cases, an individual does not pay more than the actual LTC costs. The safety net is set at 7,000 euros.

We, for now, do not include bequest motives. We assign a fixed income stream (y_a) to each individual path based on income at the starting age (70). To assign (not annuitized) initial wealth to individuals we use the approach described in Section 4.2.

The other parameters are set as described in Table 2.

Table 3: Average loss of consumption due to co-payments.

	Co-payment	C	CEC
Variant 1	13 % income + 12 %wealth	0.7	1.0
Variant 2	20 % income	1.0	1.2
Variant 3	13 % LTC costs	1.4	2.1

The effect of co-payments on annuitization

There has been some discussion on the role of pensions in case of an increase in private LTC spending. The reasoning is that when private LTC costs increase individuals want to hold more of their pension wealth in the form of liquid assets instead of in annuities in order to pay for unexpected costs. As discussed by Peijnenburg et al. (2015), the optimal annuitization share will depend on the timing of LTC costs: when LTC costs occur relatively early in retired life, full annuitization is not optimal, while when LTC costs occur at the end of life (near) full annuitization might still be optimal.

We analyze the effect of co-payments on annuitization by comparing different levels of annuitization shares for two co-payment variants: no co-payments and a co-payment of 30 % of LTC costs. We first determine initial pension wealth for each individual based on their pension income in the lifecycle paths. Then, we run the model for different annuitization shares of this initial pension wealth (other wealth is never annuitized) for both variants.

6 Results

6.1 Income and wealth dependent co-payments

As the co-payments are income and wealth dependent, we analyze the differences across five spending power quintiles.

The variants are compared to a base case where co-payments are zero with no extra premiums or taxes. Table 3 shows the average loss of consumption over all individuals compared to the base case, both in euros (C) and in certainty equivalent consumption (CEC).

The effects on average lifetime consumption per spending power quintile are shown in Figure 3. Although the same amount of revenues is raised in each variant, average loss in lifetime consumption is largest in variant 3 for all spending power groups. The reason for this, is that the variant also affect saving behavior: variant 3 increases precautionary savings, so that people, on average, die with higher amounts of remaining wealth, thus decreasing average lifetime consumption. A similar mechanism also explains why average consumption loss is larger for almost all groups in variant 2 compared to 1.

The loss in average lifetime consumption is largest for the lowest spending power quintiles for each variant, because lifetime LTC is highest for this group. Consumption

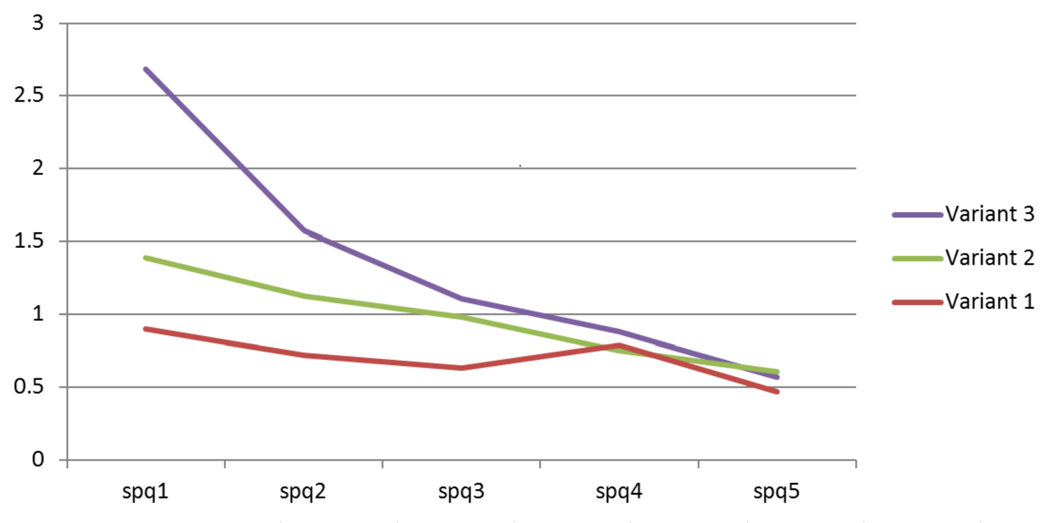

Figure 3: Loss of average consumption by introducing co-payments for LTC, per quintile of spending power.

loss does become much smaller for this group when going from a generic co-payment to an income- or income- and wealth-dependent co-payment.

Figure 4 shows the loss in certainty equivalent consumption. the pattern is generally similar to the loss in average consumption. However, losses are higher since utility losses due to uncertainty are now also included. These losses are highest for the lowest spending power groups. A difference with Figure 3 is that in terms of CEC the highest spending power groups are worst off in variant 1, instead of 3.

6.2 The effect of co-payments on annuitization

Figure 5 shows certainty equivalent consumption for different spending power groups at different levels of annuitization of pension wealth in the case of no co-payments and in case of a 30 percent co-payment. Even with co-payments full annuitization of pension wealth is optimal. This is due to the fact that institutional LTC costs occur mostly at the end of life. However, the welfare gain of annuitization is smaller in case of co-payments compared to no co-payments. This is especially the case for individuals in low spending power groups. For these groups, there is a tradeoff: less annuitization means more liquid assets that can be used to pay for LTC costs, but less insurance against mortality risk. Higher spending power groups have a higher amount of initial financial wealth, apart from pension wealth, that they can use to pay for LTC costs. For these groups the value of additional liquid assets to pay for LTC costs is thus limited.

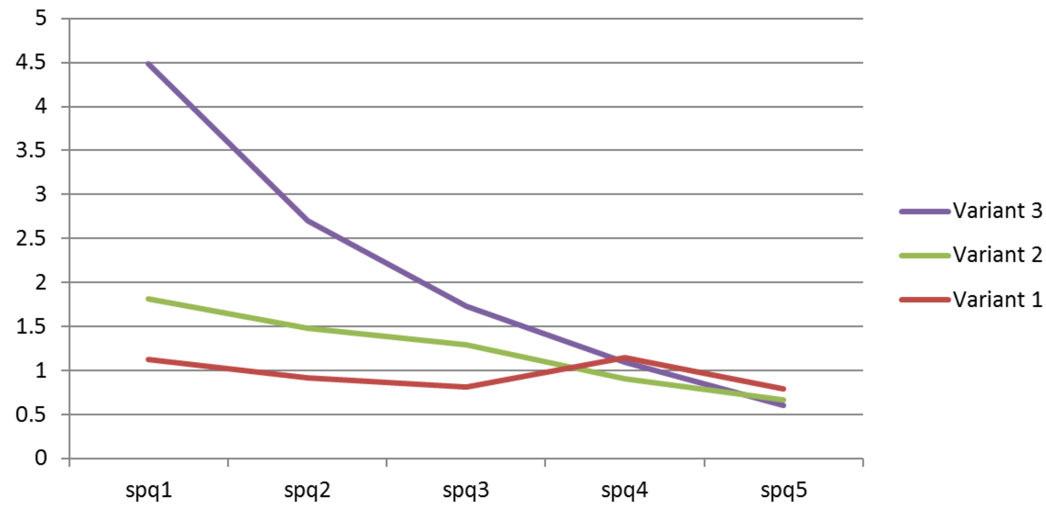

Figure 4: Loss of certainty equivalent consumption by introducing co-payments for LTC, per quintile of spending power.

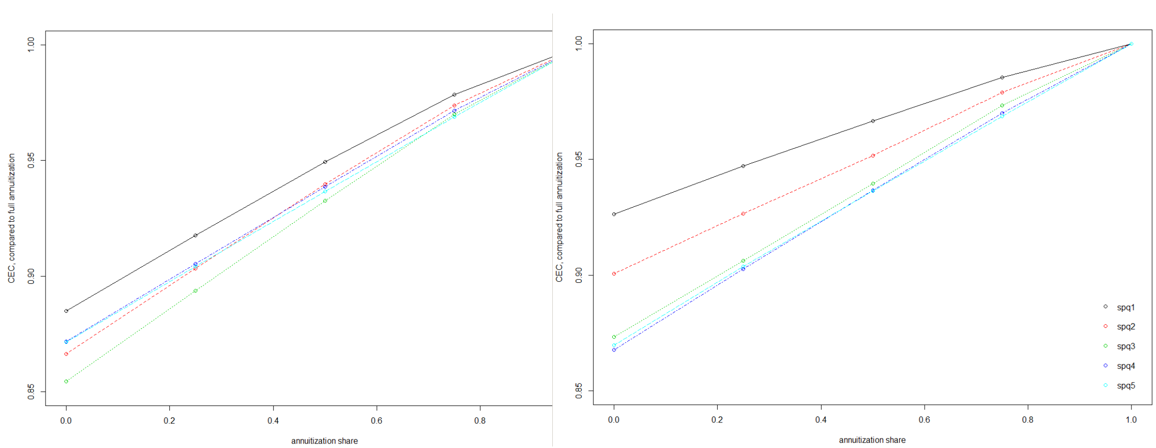

(a) No co-payments

(b) A 30 % co-payment

Figure 5: CEC for different levels of annuitization of pension wealth, by spending power quintile

References

Ameriks, J., Caplin, A., Laufer, S., and Van Nieuwerburgh, S. (2011). The joy of giving or assisted living? using strategic surveys to separate public care aversion from bequest motives. *The journal of finance*, 66(2):519–561.

Brandt, M. W., Goyal, A., Santa-Clara, P., and Stroud, J. R. (2005). A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies*, 18(3):831–873.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, 91(3):312–320.

Colombo, F. and Mercier, J. (2012). Help wanted? fair and sustainable financing of long-term care services. *Applied Economic Perspectives and Policy*, 34(2):316–332.

De Nardi, M., French, E., and Jones, J. B. (2010). Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, 118(1):39–75.

Farmer, J. D. and Sidorowich, J. J. (1987). Predicting chaotic time series. *Physical review letters*, 59(8):845.

French, E. and Jones, J. B. (2004). On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19(6):705–721.

Hsieh, D. A. (1991). Chaos and nonlinear dynamics: application to financial markets. *The journal of finance*, 46(5):1839–1877.

Hussem, A., Ter Rele, H., and Wouterse, B. (2017). Inkomens en vermogensafhankelijke eigen bijdragen in de langdurige ouderenzorg een levensloopperspectief. Design Paper in press, Netspar.

Hussem, A., Van Ewijk, C., Ter Rele, H., and Wong, A. (2016). The ability to pay for long-term care in the Netherlands: A life-cycle perspective. *De Economist*, pages 1–26.

Koijen, R. S., Nijman, T. E., and Werker, B. J. (2010). When can life cycle investors benefit from time-varying bond risk premia? *Review of Financial Studies*, 23(2):741–780.

Maarse, J. H. and Jeurissen, P. P. (2016). The policy and politics of the 2015 long-term care reform in the netherlands. *Health Policy*, 120(3):241–245.

Peijnenburg, K., Nijman, T., and Werker, B. J. (2015). Health cost risk: A potential solution to the annuity puzzle. *The Economic Journal*.

Wong, A., Boshuizen, H., Polder, J., and Ferreira, J. A. (2016). Assessing the inequality of lifetime healthcare expenditures: a nearest neighbour resampling approach. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*.