HEDG

HEALTH, ECONOMETRICS AND DATA GROUP

WP 18/24

The effect of co-payments in Long Term Care on the distribution of payments, consumption, and risk.

Bram Wouterse; Arjen Hussem and Albert Wong

August 2018

http://www.york.ac.uk/economics/postgrad/herc/hedg/wps/

The effect of co-payments in Long Term Care on the distribution of payments, consumption, and risk. 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 for Public Health and the Environment, Bilthoven, The Netherlands
[5]Network for Studies on Pensions, Ageing and Retirement, Tilburg, The Netherlands

July 13, 2018

Abstract

Population aging leads to concerns about the financial sustainability of collective long term care insurance systems. One way to keep public spending in check is by increasing the role of co-payments. An interesting feature of the co-payments that have been introduced in the Netherlands is that they are income- and wealth dependent. This dependency allows the fine-tuning of effects across income groups, but can also distort consumption decisions of the elderly.

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 flexible semi-parametric nearest-neighbor approach to estimate lifecycle paths of long term care spending. We apply this approach to an extensive administrative data set for the entire Dutch elderly population. The estimated paths are then used as inputs in a stochastic lifecycle decision model for singles at the retirement age.

We analyze the effects of different co-payment schemes on the distribution of LTC payments, consumption and risk across income groups. We find that, compared to a flat-rate co-payment, income- and especially wealth-dependent co-payments lead to much lower welfare costs for groups with low financial means. At the same time, the welfare costs of the groups with the highest means increase only slightly. Excluding a bequest motive leads to lower, and including health-state dependent utility to higher welfare losses due to co-payments compared to full insurance.

1 Introduction

The aging 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. Countries with extensive public LTC insurance, such as the Netherlands, are seeking to increase private financing to keep public spending in check (Colombo and Mercier, 2012). One way to do so, is by increasing the level of co-payments. An interesting feature of the Dutch system is that the co-payments are income- and wealth-dependent. LTC users pay full costs of LTC, up to a means-dependent maximum amount. This system differs from both flat-rate co-payments, that are independent of the financial means of the users, and means-testing, where LTC users can only apply for social LTC insurance when they have depleted their financial means. Even compared to means-based co-payment systems in other countries (see Colombo and Mercier (2012) for an overview), the Dutch system enables quite specific fine-tuning of the financial impact of co-payments across income- and wealth groups. At the same time, however, it can distort the saving decisions of the elderly.

Co-payments affect the distribution of spendable income across pensioners. Elderly with low incomes use more public LTC, on average, than those with high incomes (Bakx et al., 2016). Public insurance of LTC thus redistributes income from high to low income groups, while increasing the role of co-payments limits this redistribution. As in the case for income inequality (Aaberge and Mogstad, 2015), these redistributions are ideally measured over the whole lifecycle: the lifetime distribution of LTC costs is much more equal than that in one particular year. This means that a cross-sectional analysis might overestimate the distributional effects (Hussem et al., 2016). Further, an analysis of the effect of co-payments across income groups should not only include the average (ex-post) redistribution of payments and benefits, but also the effects on risk across income groups. As shown by McClellan and Skinner (2006), insurance against risk is an important part of the value of public health insurance, especially for low income groups. The insurance value is specifically important in LTC, where buying insurance on the private market is difficult or not possible at all (Brown and Finkelstein, 2007).

Co-payments also affect saving behavior. The introduction of high co-payments means that the elderly are confronted with the risk of potentially very substantial costs during their remaining life, for which they need to hold liquid wealth. Indeed, studies for the U.S. find that precautionary savings for LTC costs can, partially, explain why the elderly hold relatively large amounts of financial wealth, even at high ages, and why they choose not to fully annuitize this wealth (De Nardi et al., 2010; Ameriks et al., 2011; Peijnenburg et al., 2017; Kopecky and Koreshkova, 2014). Income- and wealth-dependent co-payments make the effect on saving behavior more complex: the elderly still want to hold liquid assets to pay for LTC costs and smooth consumption, while at the same time the co-payments are an implicit tax on wealth, reducing the desire to hold wealth at high ages.

An important empirical challenge in assessing the distributional effects of co-payments, is the modeling of the lifecycle distribution of LTC costs. This is hard to do 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; Jones et al., 2018) 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). A recent study by Hurd et al. (2017) for out-of-pocket LTC spending in the U.S. shows the relevance of a non-parametric approach. They compare their non-parametric method, based on matching, to a more standard parametric Markov model. They find that the risk (the chances of extreme use or costs) as estimated non parametrically is substantially greater than the risk as estimated by the parametric model.

In this paper, we use a semi-parametric nearest neighbor approach developed by Wong et al. (2016) and Hussem et al. (2016) to estimate the lifetime distribution of LTC costs. The main advantages of this approach are its flexibility and the ability to use it on short periods of panel data. It enables the modeling of the complex dynamics in LTC cost, and the joint dynamics in income, wealth, and other relevant variables. We apply the nearest neighbor algorithm to estimate 20,000 synthetic lifecycle paths using a rich administrative data that includes information on LTC spending, household status, income, and wealth for the entire Dutch elderly population.

The estimated paths serve as inputs in a stochastic lifecycle decision model for singles at the retirement age. 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 LTC. We use a simulation based algorithm developed by Koijen et al. (2010) to solve the model. The use of health dynamics in a lifecycle model adds additional challenges: the health model should be sophisticated enough to capture the complex dynamics in health, while at the same time be parsimonious enough so that its use in a structural life-cycle model is computationally manageable (De Nardi et al., 2017). The combination of the synthetic lifecycle paths with the approach of Koijen et al. (2010) fulfills both these requirements.

We analyze the distributional effects of different forms of income- and wealth-dependent co-payments that are implemented or considered by policy makers. We take it as a given that the government wants to finance a fixed percentage of aggregated

LTC costs out of co-payments [1]. The question we want to answer is how the design of the co-payment scheme, and especially the dependence on income and wealth, affects redistributions across income groups. This is relevant for the policy debate in the Netherlands, where wealth-dependent co-payments have only recently been introduced and are still a topic of debate. It is also relevant for other countries that are considering to introduce co-payments in LTC, but want to mitigate some of the distributional effects.

Our model allows us to analyze effects over the entire lifetime after the pension age and to incorporate the value of insurance. We concentrate on the effects on average LTC payments, average consumption, and certainty equivalent consumption across income- and wealth-groups. Our analysis is related to recent studies on the (value of) income transfers and insurance in Medicare (Khwaja, 2010) and Medicaid (De Nardi et al., 2016b), that both take a lifecycle perspective and include endogenous saving behavior. Although we model LTC costs as exogenous shocks, and thus cannot directly assess the effects of the co-payments on LTC use, we do investigate the effects on the price of LTC for the user at different margins.

This paper is organized as follows. In Section 2, we describe the Dutch LTC system and the role of co-payments in this system. In Section 3, we discuss the application of the nearest neighbor algorithm on Dutch LTC data, and we describe the estimated lifecycle paths. In Section 4, we introduce a lifecycle model for consumption and saving of retirees in case of LTC co-payments. We also explain the numerical approach that allows us to use the estimate lifecycle paths in this model. In Section 5 we show our results, and in Section 6 we discuss the implications of our findings and the main limitations of our approach.

2 The Dutch Long Term Care system

The Netherlands has one of the most extensive collective LTC arrangements in the world (Colombo and Mercier, 2012). In the period we investigate (before 2015), a social insurance, called the exceptional medical expenses act (AWBZ), covered a broad range of home care services (social support, personal care, nursing) and institutional care (nursing homes and residential care). The income-dependent premium for the AWBZ was collected through the income tax (including pension income) in the first and second income brackets.

Users of long term care pay a co-payment. This co-payment functions as a means-dependent deductible: users pay the full costs of LTC, up to a maximum amount. This maximum amount depends on the financial means of the individual, and differs according to the type of care (home car or institutional care) and living situation. The financial means are defined as net income and a fixed percentage of financial wealth. In 2013, this percentage was increased from 4 to 12 percent. Co-payments for home care

[1] This could be because of efficiency considerations (moral hazard, effects on labor supply), concerns about intergenerational redistribution (Wouterse and Smid, 2017), or some other reason

are lower than for institutional care: for home care, users pay a maximum amount that equals 15 percent of their means, while users of institutional care have to contribute up to 75 percent of their means. The details of the co-payment system are explained in Section 4.3.

In 2015, the long term care system has been reformed. The provision of home care is now mainly the responsibility of municipalities. They receive a financial contribution out of the general means of the national government Although there are now differences in the level of co-payments for home care across municipalities, the financing of long term care has remained largely the same: it is still largely based on income-dependent premiums or taxes, and users of care, in general, still pay a contribution with a maximum based on income and wealth.

3 Long term care spending over the lifecycle

3.1 Source data

We use administrative data on LTC use from the Dutch Central Administrative Office (CAK). These data cover the period 2008-2013. The data include information on all publicly financed 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 Dutch health care instutute (CVZ) annual reports for intramural care.

The LTC data is linked to other datasets using a unique personal identification number. The Dutch Municipal Register provides basic information on everyone enlisted in a Dutch municipality. From this register, we obtain date of death, age, sex and marital status. We use data from the tax services to obtain gross income, net financial wealth, and net housing wealth.

We select individuals who are alive at least up to January 1 2013, who are 67 or older in 2013, and who are single over the whole observation period. We purge the data from period effects (see Appendix A for the details).

3.2 The nearest neighbor algorithm

We estimate lifecycle paths of LTC use with 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 from 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). We use 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 lifecycle 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 (e.g. income, wealth) 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 of the variables of interest $Y^j_{a,t}$ 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 $Y^j = \{Y^j_{A-p+1,T-p}, ..., Y^j_{A,T-1}\}$ 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 an 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.

The time periods $t = 1, .., T$ and the number of lags p will generally depend on the data at hand. When the available panel data is long enough, the number of lags can be determined by comparing model performance across different choices of p. See Wong et al. (2016) for examples.

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 $a = 0$.[2] 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 path 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 $d(\mathbf{z}, \mathbf{y})$. We use the Mahalanobis measure, which corrects for the correlation between the components of y and differences in scale. This 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. Out of the k-nearest neighbors, one neighbor is randomly drawn.

In our main specification, we use 2 lags for the categorical variables, and 1 lag for the continuous variables. Thus, we only include the years 2010-2013. We set

[2]Obviously no information is available when the simulation starts at the age of 0.

$k = 2$, and start at age 70. The data is stratified by sex, age, and home-ownership, and matched on income, financial wealth, housing wealth, and LTC expenditures. We simulate 10,000 paths for women and 10,000 for men.

3.3 Estimation results

The lifetime distribution of LTC costs

In Appendix B, we assess how well the model fits the actual data. Based on this assessment, we conclude that the algorithm provides a sufficiently credible reproduction of real life cycles of health care costs, income and wealth to study the effects of co-payments for LTC.

Table 1 shows statistics of the estimated lifetime LTC costs. On average, a 70-year old single uses almost 32,000 euros of home care and 45,000 euros of nursing home care over the rest of his or her life. The costs are distributed very unevenly: 19 percent of the elderly does not use any home care, while 5 percent of the elderly uses more than 138,000 euros of home care. Almost half of the elderly (49 %) do not use any nursing home care, while the top 5 percent use 253,000 euros of nursing home care or more. 14 percent of the elderly uses neither home care nor nursing home care, while the 5 percent of the elderly that use the most LTC overall, have total LTC costs of 326,000 euros or higher.

Figure 1 shows the age pattern of LTC use. The top part shows the average spending by age. Until the age of 80, this amount is limited to 2,500 euros annually for home care and the same amount for nursing home care. For home care, the average costs rise gradually to 5,000 euros for the age of 95. The increase for nursing home care is much steeper, and average cost go up to about 17,000 euros at the highest ages. The bottom figure shows the composition of the population by age in five groups: no costs, low costs ($< 5,317$), medium costs ($5137 - 22,500$), high costs ($> 22,500$), and deceased. The rising age pattern is explained by both an increase in the percentage of people using LTC (among the survivors) and an increase of the average costs per user (the relative size of the high cost group increases with age).

Distribution of LTC costs across total wealth groups

Our estimates contain individual income and wealth for each lifecycle path. To simplify both the analysis and the interpretation, we group all individuals in wealth deciles and income quintiles. We assign each individual with a fixed income stream (y), equal to average income at 70 within his income group, and initial financial wealth at 70, equal to average financial wealth at 70 within his wealth group.

In our presentation of the results, we focus on te distribution of LTC costs across total pension wealth groups. We define total pension wealth as the sum of the expected [3] present value of the income stream over the rest of life and initial financial wealth

[3] The expectations are equal to the average per income and gender group.

Table 1: Descriptive statistics for lifetime LTC costs at 70, for the whole population and by total pension wealth group 1 to 5. All amounts are discounted using a discount factor of 1.5 % (see Table 2)

		mean	std	% no costs	p25	p50	p75	p95
all	Home care	31,679	57,523	19	622	8,509	37,317	138,148
	Nursing home	45,130	98,193	49	0	262	34,717	252,921
	Total	76,810	118,051	14	2,620	24,979	102,960	326,098
1	Home care	39,907	71,088	23	272	10,027	46,384	187,657
	Nursing home	52,257	110,506	49	0	466	43,461	287,309
	Total	92,164	138,297	16	2,500	30,399	128,076	391,408
2	Home care	37,920	65,423	19	750	11,202	48,461	162,019
	Nursing home	51,667	105,605	47	0	1,087	47,527	271,743
	Total	89,587	125,908	13	3,790	36,401	124,978	354,229
3	Home care	29,927	56,238	20	510	7,924	35,188	128,830
	Nursing home	45,918	99,352	48	0	520	39,699	252,151
	Total	75,845	117,593	14	2,394	25,464	101,456	328,760
4	Home care	27,472	46,758	18	690	7,930	34,151	121,281
	Nursing home	43,003	91,604	48	0	564	33,911	245,344
	Total	70,476	106,098	13	2,574	22,285	97,075	295,155
5	Home care	22,667	39,558	17	890	7,122	27,119	97,673
	Nursing home	32,256	78,598	53	0	0	18,583	196,611
	Total	54,924	91,945	13	2,202	15,996	64,430	246,663

at 70. We group individuals in five total pension wealth quintiles and show average results for each group. In Appendix D, we also show results for specific combinations of pension income and initial wealth.

The top part of Figure 2 shows lifetime income and initial wealth across total pension wealth quintiles. The figure shows that higher total pension wealth quintiles have, naturally, both more remaining lifetime income and higher initial financial wealth at the age of 70. Financial wealth for this cohort of 70-year old singles seems to be remarkably low, also compared to all 70-year old households (see e.g. Hussem et al. (2017)), but this seems to be in line with other sources[4].

The bottom part of Figure 2 shows the life expectancy, and expected number of years with use of home care and nursing home care, for each total pension wealth group. Despite va lower life expectancy, the elderly with the least financial means

[4]statline.cbs.nl

spend more lifeyears, on average, in need of home care and nursing home care. This also results in the highest expected LTC costs for these groups. The statistics of the estimated lifetime LTC costs across total wealth groups can also be found in Table 1. The total LTC costs for the quintile with the lowest total wealth are 92,000 euros on average (Table 1). For the highest quintile, this is 55,000 euros. Although groups with low means also have the highest probability of using any LTC, the difference in costs is mainly driven by intensity of use [5]: within the lowest wealth quintile, the 5 percent users with the highest cost spend 390,000 euros of LTC or more. For the highest wealth quintile, this is 247,000 euros or more.

[5]Differences in average discounted costs across income groups are also partly explained by differences in timing. High total wealth groups live longer, and thus, on average, use LTC at higher ages than low groups. Differences in timing explain about 10 % of the total difference in discounted costs: with discounting the lowest wealth group has average costs that are 40 % higher than for the highest wealth group, without discounting this is 37 %.

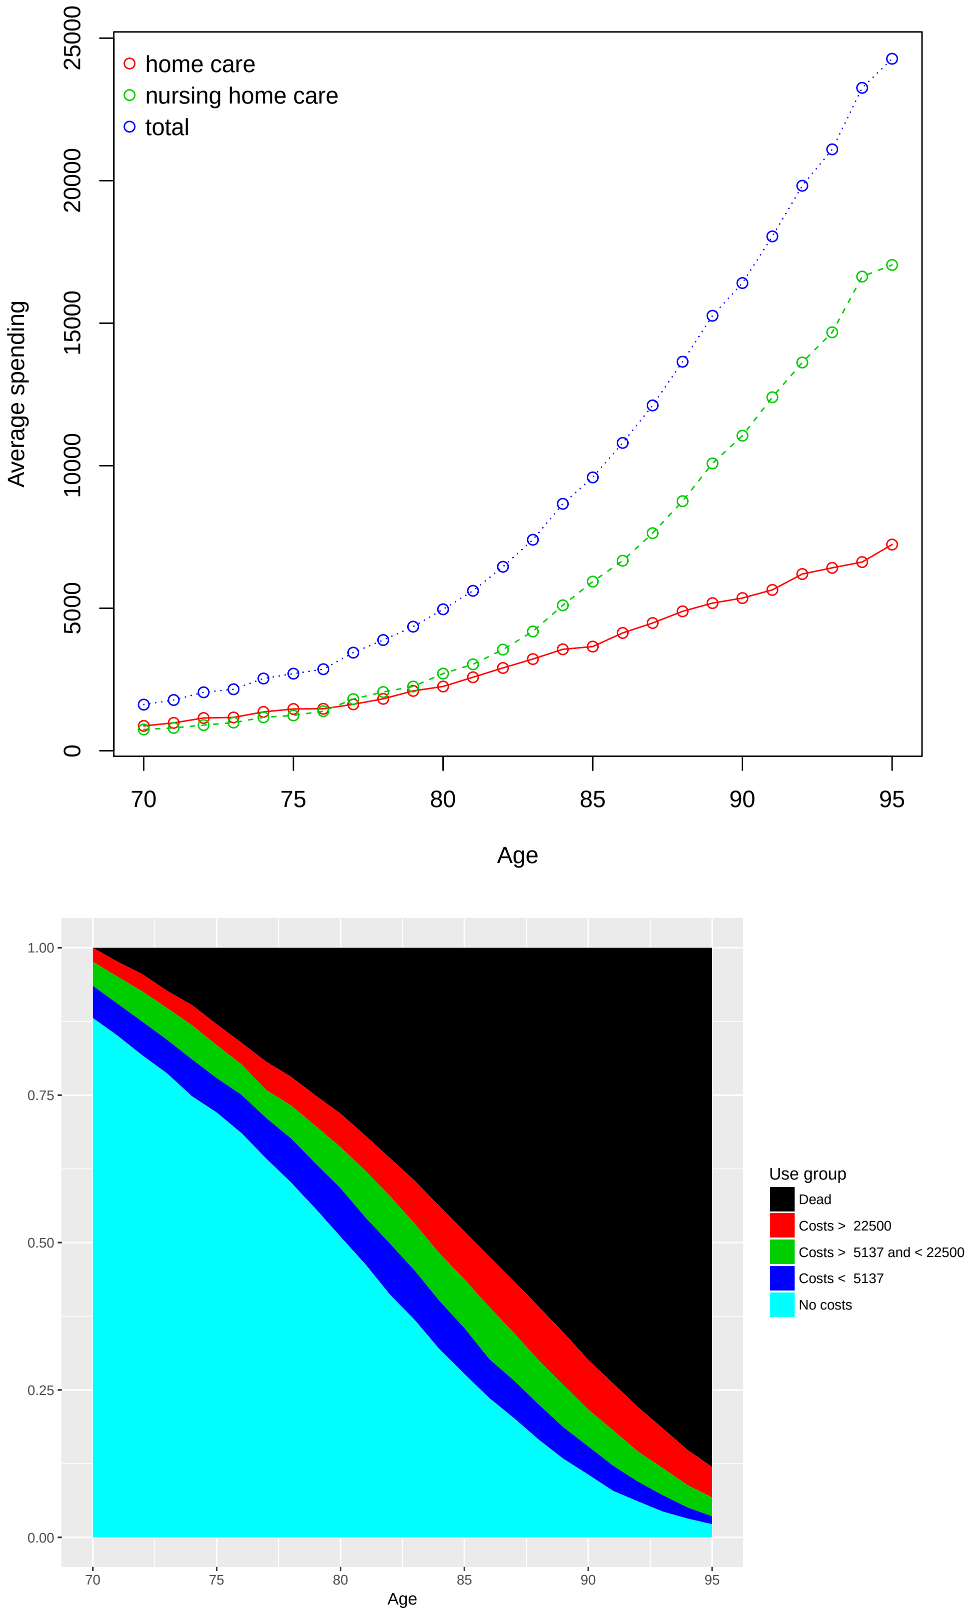

Figure 1: LTC costs by age. Average (top) and categories (bottom)

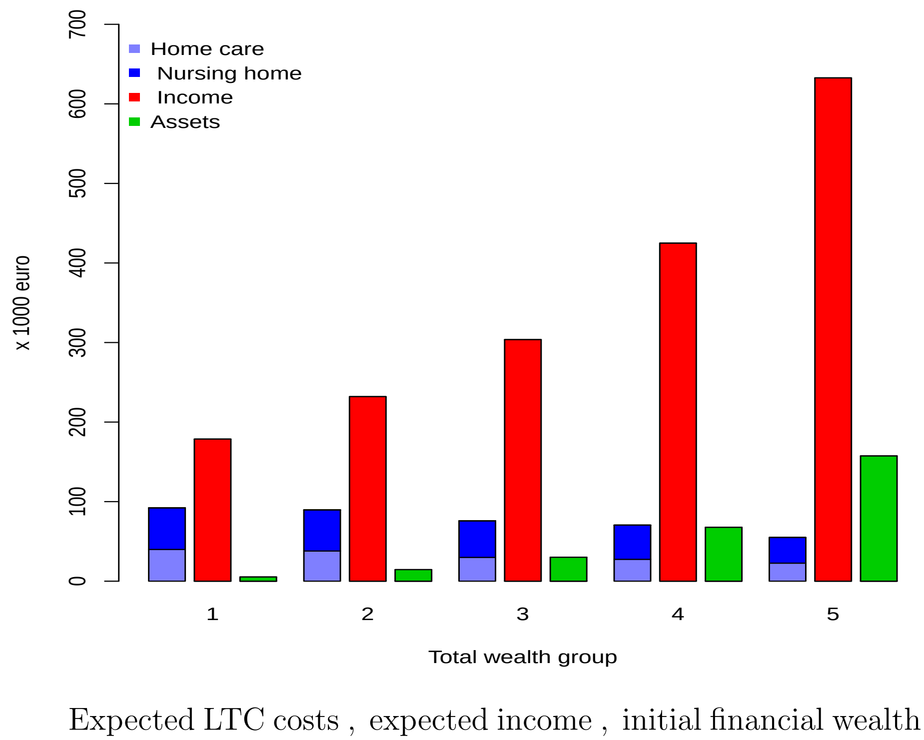

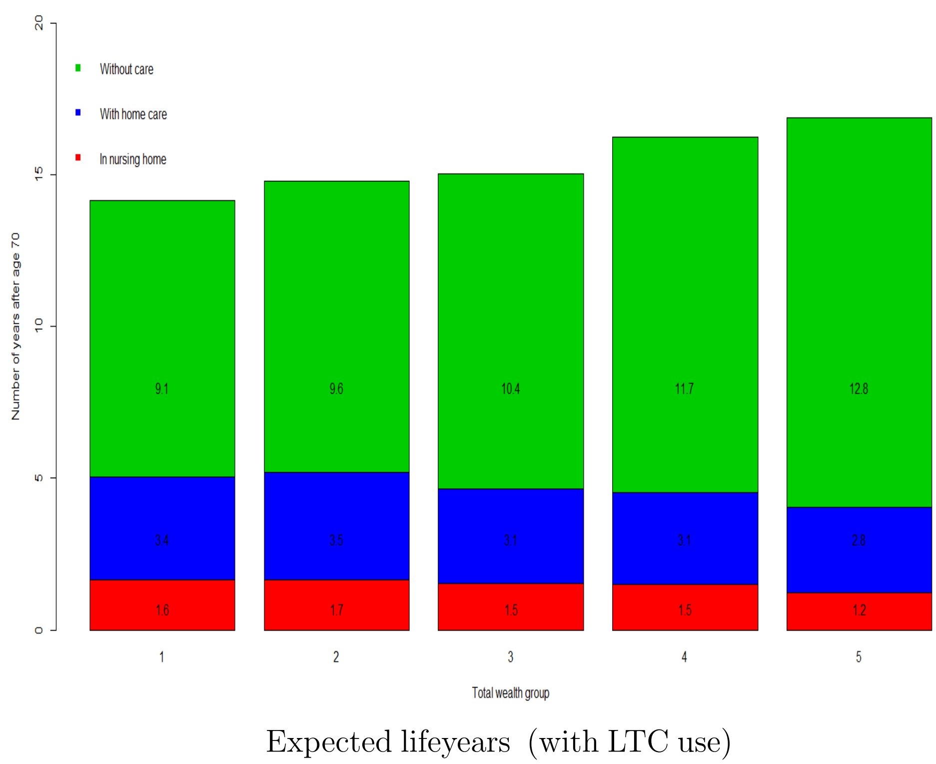

Figure 2: Total pension wealth groups. Expected (discounted) LTC costs, expected income, and initial financial wealth, at 70 (top). Expected lifeyears (with and without) care at 70. By quintile of total pension wealth.

4 A model of lifecycle consumption after retirement

4.1 The model

The estimated lifecycle paths provide a semi-parametric distribution function of LTC costs and mortality. We implement a standard lifecycle 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 algorithm developed by Koijen et al. (2010).

The basic model

We model the consumption and savings decisions of individuals after retirement. We start with a relatively simple model, and add several extensions (bequests, health-state depent utility of consumption, annuities) in the next section. 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 from consumption. The individual wants to maximize his expected utility over remaining lifetime. With a time-separable utility function the individual's maximization problem then is:

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

with p_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 wants to save for later (m_t). The individual is also faced with co-payments for LTC costs h_t. He faces the following annual budget constraint:

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

We impose the borrowing constraint $W_t \geq 0$.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 $r - 1$, so that

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

Extensions

We extend the model in four ways. First, we allow for a fixed pension income stream (state pension and/or annuity). The budget constraint then becomes

$$c_t + m_t + h_t = W_t + y. \tag{6}$$

Second, we allow the level of co-payments to depend on wealth and pension income. Let H_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 H_t, but depends on the co-payment rules set by the government. We use the following general co-payment rule:

$$h_t = \min[\tau H_t, \nu_y y + \nu_w W_t - \delta, \mu]. \tag{7}$$

This general rule allows us to emulate the Dutch co-payment system, but also to include other variants, such as a nominal co-payment independent of spending power. The government sets the parameters τ, ν_y, ν_w, δ, and μ. 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 are the maximum shares of income and wealth that have to be spent on co-payments. The parameter δ is a sort of deductible: a fixed amount of income that is exempted from the co-payments. There government can also set a maximum μ on annual co-payments.

The way the government sets the co-payment rules affects the optimization problem of the individuals. When $\nu_w > 0$, co-payments are no longer fully exogenous since they depend on the annual savings chosen by the individual.

As a comparison to the variants with co-payments, we also include a variant where health care spending is financed out of a, income-dependent, premium π. In that case individuals will pay a contribution to the health care system, regardless of their own use, but depending on their pension income: πy.

Third, 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):

$$g(W_{death}) = \theta \frac{(W_{death} + \xi)^{1-\gamma}}{1 - \gamma}, \tag{8}$$

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

Fourth, we allow for health state-dependent utility. The utility an individual derives from non-health care consumption could depend on his health status (disability).

Finkelstein et al. (2013) find that an increase in the number of chronic diseases has a significant negative impact in the marginal utility of consumption. A priori, however, the effect of poor health could go both ways: individuals might derive less utility from things like eating out or recreation, but at the same demand for things like domestic help, wheelchairs, and stairlifts might increase (Meyer and Mok, 2009). Indeed, as pointed out by Peijnenburg et al. (2017), there is no consensus in the empirical literature on the size and even the sign of the effect.

We use a health dependent utility function only for individuals who use nursing home care, but not for individuals who use home care. A negative effect on the marginal utility of consumption is more likely for nursing home care users, as this type of care is relatively comprehensive and encompasses most additional consumption needs (housing, cleaning) related to disability.

To include state-dependent utility, we use the following commonly used adaptation of the utility function in Equation (4) (Palumbo, 1999; De Nardi et al., 2010; Peijnenburg et al., 2017):

$$u(c_t) = (1 - \kappa he_t)\frac{c_t^{1-\gamma}}{1-\gamma}. \tag{9}$$

The variable he is a dummy indicator for poor health, which we define as an individual having any nursing home care in period t. The parameter κ determines the relative change in the marginal utility of consumption in poor health ($he_t = 1$) compared to good health ($he_t = 0$). When $\kappa < 0$, marginal utility is lower in poor health. When $\kappa = 0$, marginal utility is equal in both health states.

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 the averages of this measure $CEC_{g,v}$ for each total pension wealth quintile $g = 1, .., 5$ across policy variants v.

4.2 Numerical approach

The individual's maximization problem can be solved using dynamic programming. The lifecycle optimization problem is divided into smaller yearly optimization problems. The algorithm starts at the last time period T, and is then solved backwards recursively. We solve this problem using the approach developed by Koijen et al. (2010), that has been applied to LTC financing in the U.S. by Peijnenburg et al. (2017). The approach combines the method of endogenous gridpoints (Carroll, 2006) with a simulation based approximation of the expected values (Brandt et al., 2005). The approach is well suited to use in combination with the semi-parametric estimation of the lifecycle paths. Most approaches approximate the stochastic processes (mortality, LTC costs) 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.

Specifically, solving the maximization problem involves the estimation of decision rules (the optimal amount of consumption in period t given initial wealth W_t *at the beginning* of t) over a grid of values for W_t. In the endogenous gridpoints method, these decision rules are determined by finding the optimal consumption c_t^* for a grid of values for wealth m_t *at the end* of period t (after consumption and health care costs). Given that we already have the optimal consumption rules for period $t+1$, optimal consumption in t given m_t is determined by the Euler condition:

$$c_t^*|m_t = (\mathrm{E}(\beta c_{t+1}^{*-\gamma} r|m_t))^{-\frac{1}{\gamma}}. \tag{11}$$

The most relevant part of the method of Koijen et al. (2010), in our context, is that $\mathrm{E}(\beta c_T^{*-\gamma} r|m_t)$ is estimated using a simulation approach. The lifecycle-paths provide a large number of random draws from the stochastic process determining mortality and LTC spending. The expected values in Equation (11) can thus be estimated with regression analysis using the realizations of $c_{t+1}^*|m_t$ in each path. Appendix C provides a detailed overview of the numerical procedure.

4.3 Implementation

We use the lifecycle paths and the lifecycle model to evaluate different co-payment schemes for the Dutch collective long term care insurance.

Policy variants

In the Dutch system, co-payments depend on both income and wealth. We emulate the Dutch co-payment scheme in 2015 using the formula in Equation (7). In this scheme, 75 % of income and 9 % of financial wealth[6] is included in the co-payment for nursing home care and 15 % of income and 2 % of financial wealth for home care. There is a deductible of 4,500 euros for nursing home care and 16,600 euros for home care. There also is a maximum co-payment of 27,000 euros.

We introduce two alternative co-payment schemes. First, a flat-rate co-payment: co-payments are a fixed percentage τ of an individual's annual LTC costs, independent of his income and wealth. Second, an income-dependent co-payment: co-payments are a share ν_y of income, but do not depend on an individual's wealth. This variant resembles the co-payment scheme in place before 2013. During the 2017 Dutch election campaign, some political parties proposed to return to a co-payment system only depending on income (CPB, 2017).

[6] 12 % of financial wealth is added to the income definition used to calculate the co-payment. As 75 % of this income definition is included, this means that 0.12*0.75 = 9 % of financial wealth is included.

Additional to these co-payment schemes, we introduce two full insurance variants. These insurance variants are not meant to compare the welfare effects of full insurance to co-payments (as this would require a welfare analysis including moral hazard effects), but to use as baselines to compare the distributional effects of the co-payment variants to.

The parameters of the other variants are set in such a way that they finance an equal amount of aggregated nursing home and home care costs as the variant based on the current system. We focus only on the financing of the share of health care costs financed out of co-payments in the current system (which in our model is equal to 27 percent of total LTC costs). This means that in the variants with premiums instead of co-payments, these only have to raise the amount of LTC costs currently financed out of co-payments. We do this to make a fair comparison. We only model the retirement phase of life. Variants that would, for instance, raise less revenues than the current policy would imply the redistribution of costs from the older generations, who are in the model, to the younger (pre-retirement) generations, who are not in the model.

In the first insurance variant, individuals pay a fixed percentage π of their pension income as an annual premium. The premium is independent of an individual's (expected) LTC costs. This variant reflects a base case, where the amount of LTC costs currently financed out-of-pocket are financed out of the social LTC insurance premium. This premium is income dependent in the Netherlands[7]. In the second insurance variant, we include an actuarially fair premium π_g per total pension wealth group g. This means that the premiums paid by group g exactly finance the costs of group g.

The policy parameters for the, in total, five variants are shown in Table 2. In all cases, co-payments do no exceed the actual LTC costs. The uniform income dependent premium π is 6.22 %. The actuarially fair premiums range from 2,020 euro for the lowest total pension wealth group to 1,086 for the highest group. The consumption floor is set at 7,000 euros.

Other parameters

The other parameters are set as described in Table 3. In the main specification we include a bequest motive by setting $\theta = 2.3$. We set the risk aversion parameter γ to 3. We perform two sensitivity analyses. In the first one, we do not include a bequest motive. In the second one, we introduce state-dependent utility of consumption. We set $\kappa = 0.2$ which means that marginal utility of consumption is 20 percent lower for individuals living in a nursing home than for others. De Nardi et al. (2010) choose a similar value for κ and it seems to be at the more extreme side of the range of values found by Finkelstein et al. (2013).

[7]This premium is actually only levied over the first two income brackets, so the maximum premium payment is capped at some income level. However, as the social insurance premium is in practice treated as an integral part of the tax system, it seems to make more sense to model the financing of additional revenues through an uncapped income dependent premium.

Table 2: Policy parameters in each variant, for the main specification. The co-payment rule (Equation 7) is : $h_t = \min[\tau H_t, \nu_y y + \nu_w W_t - \delta, \mu]$

Variant		τ	ν_y	ν_w	δ	μ
1 Uniform premium	nursing home	0	0	0	0	0
1 Uniform premium	home care	0	0	0	0	0
2 Act. fair premium	nursing home	0	0	0	0	0
2 Act. fair premium	home care	0	0	0	0	0
3 Inc dep. co-pay	nursing home	1	0.84	0	4500	27000
3 Inc dep. co-pay	home care	1	0.25	0	16500	27000
4 Inc and wealth dep co-pay.	nursing home	1	0.75	0.09	4500	27000
4 Inc and wealth dep co-pay.	home care	1	0.15	0.02	16500	27000
5 Flat-rate co-pay	nursing home	0.38	0	0	4500	27000
5 Flat-rate co-pay	home care	0.16	0	0	16500	27000

Table 3: Values of parameters in different specifications

	main	no beq.	state dep. util.
r	1.015	1.015	1.015
β	0.985	0.985	0.985
γ	3	3	3
θ	2.3	0	2.3
ξ	0	0	0
κ	0	0	0.2

5 Results

5.1 Results for the main specification

Table 4 shows the certainty equivalent consumption (CEC) in each payment variant (v) fore each total pension wealth groups (in this section sometimes abbreviated as wealth groups). The table also shows average annual consumption, LTC payments (either premiums or co-payments) and bequests (the amount of wealth they leave after death). There are considerable differences in CEC across total pension wealth groups. In the baseline variant (a uniform income dependent premium), the group with the lowest financial means has a CEC of 12,408 euros, while the highest group has 44,393 euros. LTC payments also vary considerably, across wealth groups and across policy variants. An actuarially fair premium for the lowest wealth group would be 2,013 euros. In most variants, with the exception of the flat-rate co-payment, they pay far less (between 842 and 967 euros). For the highest wealth group, things are reversed. Their actuarially fair premium is 1,152 euros, and they pay considerably more in all variants, except the flat-rate co-payment.

The design of the co-payment matters: all variants raise the same revenue, but the distribution of costs, in terms of effects on CEC, across groups differs greatly.

Table 4: Certainty equivalent consumption and average consumption, LTC payments and bequests per policy variant and total pension wealth group (1-5). Main specification

		1	2	3	4	5
1 Uniform premium	CEC	12,408	15,558	20,493	28,296	44,393
	$Consumption$	13,000	16,517	21,790	30,553	47,090
	LTC	874	1,097	1,402	1,871	2,690
	$Bequest$	732	1,403	1,982	2,970	5,064
2 Act. fair premium	CEC	11,355	14,768	20,212	28,502	45,697
	$Consumption$	11,910	15,764	21,564	30,936	48,559
	LTC	2,013	1,897	1,640	1,463	1,152
	$Bequest$	672	1,345	1,967	3,001	5,153
3 Inc dep. co-pay	CEC	12,167	14,766	19,187	26,586	44,023
	$Consumption$	13,144	16,271	21,336	30,137	47,365
	LTC	967	1,193	1,559	1,832	1,878
	$Bequest$	640	1,625	2,358	3,538	5,736
4 Inc and wealth dep co-pay.	CEC	12,255	14,887	19,499	27,134	44,011
	$Consumption$	13,248	16,357	21,535	30,342	47,534
	LTC	842	1,223	1,553	1,880	1,868
	$Bequest$	509	1,435	2,104	3,219	5,534
5 Flat-rate co-pay	CEC	10,744	13,259	18,037	26,431	44,495
	$Consumption$	11,503	14,640	20,210	29,668	47,581
	LTC	2,669	2,176	1,723	1,454	1,076
	$Bequest$	1,643	2,782	3,503	4,410	6,264

Figure 3 zooms in on these effects. The other policy variants are compared to the base case (variant 1) where co-payments are zero and LTC costs are paid out of a uniform income dependent premium. The figure shows the loss of CEC for group g in variant v compared to variant 1:

$$\frac{CEC_{g,v} - CEC_{g,1}}{CEC_{g,1}}.$$

Flat-rate co-payments, independent of income and wealth, lead to the highest loss in CEC for the lower wealth groups. Income, and especially income- and wealth dependent co-payments lead to a much lower loss in CEC for these groups, while the loss for the highest group increases only slightly. Although we are mainly interested in the effect of different forms of co-payments, the actuarially fair premium offers a useful comparison. In contrast to the co-payments, the difference between the baseline (uniform premium) and actuarially fair premium is only in the redistribution of average costs across groups; no additional risk is introduced. Lower total pension wealth groups, especially the first quintile, are worse off in this variant compared to the base case with a uniform premium: LTC use is higher for low income groups, and whereas a uniform premium redistributes costs from low to high groups, an actuarially fair premium does not. Strikingly, the actuarially fair premium leads to the lowest welfare for the lowest wealth group of all variants, except the flat-rate co-payment.

An interesting finding is that, in case of flat-rate co-payments, wealth quintiles 2 and 3 have a higher welfare loss than quintile 1. The income of members of quintile 1 is relatively close to the consumption floor, so that the net payments (co-payments minus the income transfer due to an income drop below the consumption floor) are relatively limited. Groups 2 and 3 hardly benefit from the consumption floor and do have to pay the high flat-rate co-payments.

Relevant to note is that the results by total pension wealth groups mask the effect of the income and wealth dependent co-payment variant. Almost all groups seem to benefit going from an only income based to an income and wealth based co-payment: only for the highest group this leads to a small loss in CEC. There is, however, heterogeneity in the amount of pension income and financial wealth *within* the groups. As the results for specific combination of income and wealth (Appendix D) show: the elderly that are most negatively affected by an income- and wealth-dependent co-payment are those with relatively high financial wealth and low pension income.

Including risk in the assessment of different variants, instead of only looking at effects on average consumption, is relevant. The effect of a co-payment variant on CEC can differ considerably from the effect on *average* consumption (see Table 4). Co-payments introduce uncertainty which leads to welfare losses. In some cases, the losses due to uncertainty can be very substantial. Compare, for instance, the difference between the flat-rate co-payment and the uniform premium for wealth group 3: average consumption is 1,580 euros lower in the second case, while the loss in CEC is 1.5 times larger (2,456 euros).

Given that individuals are risk averse in the model, co-payments always lead to a welfare loss compared to (an actuarially fair) insurance. The fact that we also find this, is thus not so interesting in itself. However, the way this uncertainty affects the distribution across groups is: in co-payment variants, like the flat-rate, that put a relatively large burden on groups with low means, these groups are doubly affected: they face increased average costs, but also increased risk which is also relatively costly for them.

Including the effects on savings is relevant as well. This can be seen by comparing the bequests across policy variants (Table 4). The bequest is the amount of wealth left at death. The amounts shown here are the (discounted) annualized averages. Co-payments lead to precautionary savings: all co-payment variants have higher bequests than the insurance variants. The flat-rate co-payment leads to the highest additional savings: bequests for the lowest wealth group, for instance, almost double from 874 euros in the baseline variant to 1,643 euros in the flat-rate co-payment variant. Also relevant is that the income and wealth dependent co-payment leads to a reduction in savings compared to the income dependent co-payment. The fact that co-payments depend on wealth put an implicit tax on wealth, making savings less attractive.

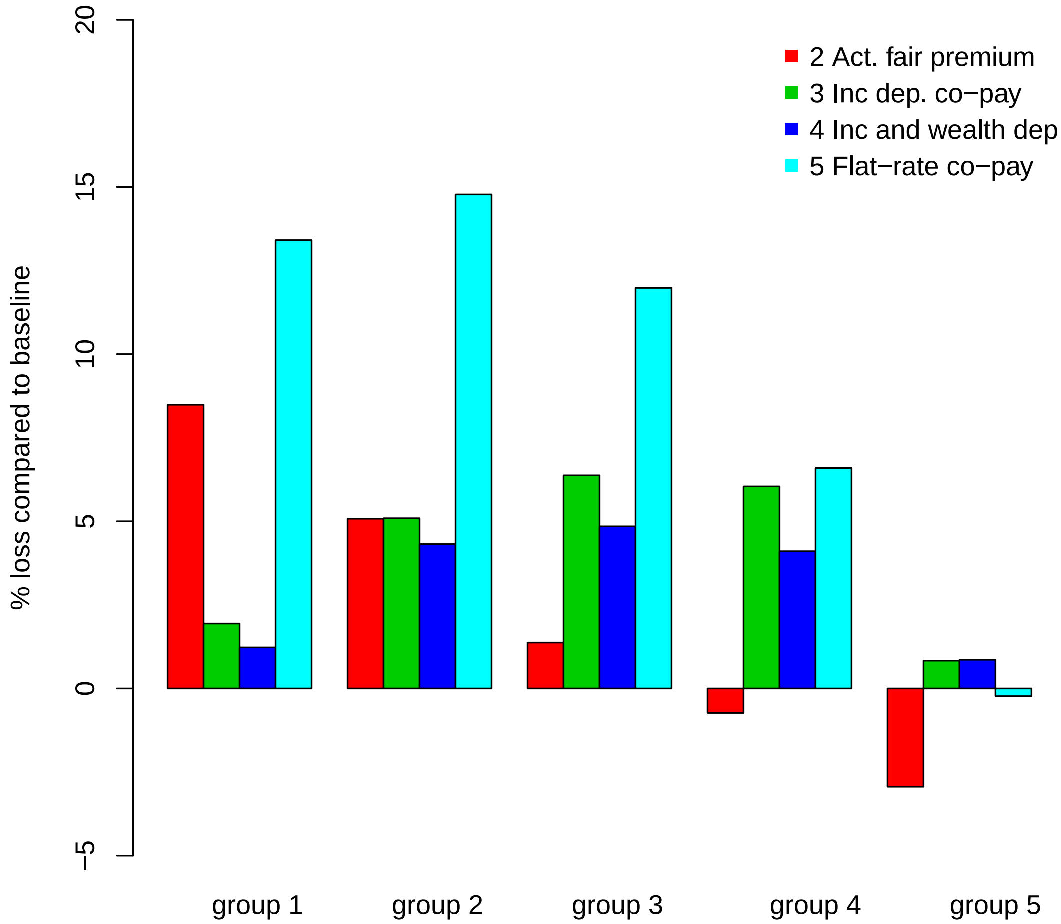

Figure 3: How large is the loss in CEC compared to the baseline policy with a uniform premium and no co-payments ($(CEC_v - CEC_1)/CEC_1$). Per policy variant (v) and total pension wealth group (1-5). Main specification

5.2 Sensitivity analysis

Table 5 shows $\frac{CEC_{g,v}-CEC_{g,1}}{CEC_{g,1}}$ by wealth group for the main specification and for two sensitivity analyses. In the first sensitivity analysis, we exclude the bequest motive. We do this by setting $\theta = 0$. We recalibrate the co-payment parameters so that all variants again raise the same revenue as the variant based on the current system ($v = 4$). The absence of a bequest motive generally increases the utility loss from introducing co-payments (compared to full insurance): as financial wealth left at death now has no value, the precautionary savings induced by co-payments are more costly in terms of utility than in the case with a bequest motive.

Table 5: The loss in CEC compared to the baseline policy (v_1) with a uniform premium and no co-payments. By total pension wealth group (1-5) and policy variant. For different parameter specifications.

	1	2	3	4	5
Main specification					
2 Act fair premium	8.49	5.08	1.37	-0.73	-2.94
3 Inc dep co-pay	1.94	5.09	6.37	6.04	0.83
4 Inc and wealth dep co-pay	1.23	4.32	4.85	4.1	0.86
5 Flat-rate co-pay	13.41	14.78	11.98	6.59	-0.23
No bequests					
2 Act fair premium	8.23	4.78	1.36	-0.65	-2.8
3 Inc dep co-pay	6.78	10.86	13.97	15.39	4.21
4 Inc and wealth dep co-pay	6.67	8.89	8.89	8.17	3.84
5 Flat-rate co-pay	15.75	19.65	16.22	10.63	2.17
State dependent utility					
2 Act fair premium	8.38	5	1.34	-0.74	-2.94
3 Inc dep co-pay	-3.89	3.08	3.89	3.67	0.17
4 Inc and wealth dep co-pay	-3.9	3.13	3.3	3.06	0.29
5 Flat-rate co-pay	7.66	12.34	9.86	5.32	-0.8

As a second sensitivity analysis, we assess the influence of health state-dependent utility by setting κ in Equation (9) to 0.2. This means that the marginal utility of consumption is 20 percent lower in poor health compared to that in good health. The loss in CEC in the variants with co-payments ($v = 3, 4, 5$) compared to the variant without co-payments ($v = 1$) is smaller than in the baseline specification: the marginal utility is now lower in poor health, which means that it is no longer optimal to fully smooth consumption between health states. Replacing a full insurance ($v = 1, 2$) with a co-payment system allows individuals to shift consumption towards years in good health. However, even with state-dependent utility, co-payments lead to a loss in CEC compared to full insurance, for all wealth groups: the utility loss due to risk is apparently larger than the gain from shifting consumption to good health years.

5.3 The marginal price of LTC

In the model, we treat LTC costs as exogenous shocks and thus changes in the co-payment do not affect the use of LTC. To shed some light on how different co-payment variants might affect the use of LTC, Figure 4[8] shows the marginal co-payment for home care and nursing home care across levels of LTC use, for each co-payment variant and across total pension wealth groups. The most relevant difference is between the income and wealth-dependent variants on the one hand, and the fixed co-payment variant on the other. The income- and wealth-dependent variants are basically systems with a deductible: LTC users completely pay LTC costs out of pocket until they reach an (income- and wealth-dependent) ceiling. Marginal costs of LTC use above the ceiling are zero. This is reflected in the marginal costs in Figure 4: at low levels of LTC use, average marginal co-payments are 1. For higher levels of LTC use, average marginal costs drop quickly towards zero, especially for low wealth groups, as more and more users reach their individual (income- and wealth-dependent) co-payment ceiling.

In contrast, the fixed co-payment variant is really a co-payment system: users never pay full costs, but only a fixed share (35 percent in our case). Although there is a maximum co-payment in this variant as well, this is reached at a much higher level of LTC costs. This can also be seen in the figure: compared to the income- and wealth dependent variants, average marginal costs at low levels of LTC use are lower, but they drop less quickly and are greater than zero over the whole range of LTC use.

The variants also differ in the marginal price that is paid across total pension wealth groups. The fixed co-payment variant leads to lower marginal costs for individuals with low financial means and low LTC use, but extends the range of use over which these co-payments are positive: leading to higher average costs for groups with low means as a whole. An income- and wealth-dependent co-payment shifts marginal costs, though in a limited way, from individuals whose total means consist mostly of income (generally low wealth groups) to individuals whose total means consist for a significant part of financial wealth.

[8]This figure is made using the simulation results from the main specification. We have pooled results over all lifeyears. As there are differences, in income and wealth and thus also in required co-payments, within observations belonging to the same total pension wealth group, we estimated the average marginal costs per 500 euro bin of LTC use.

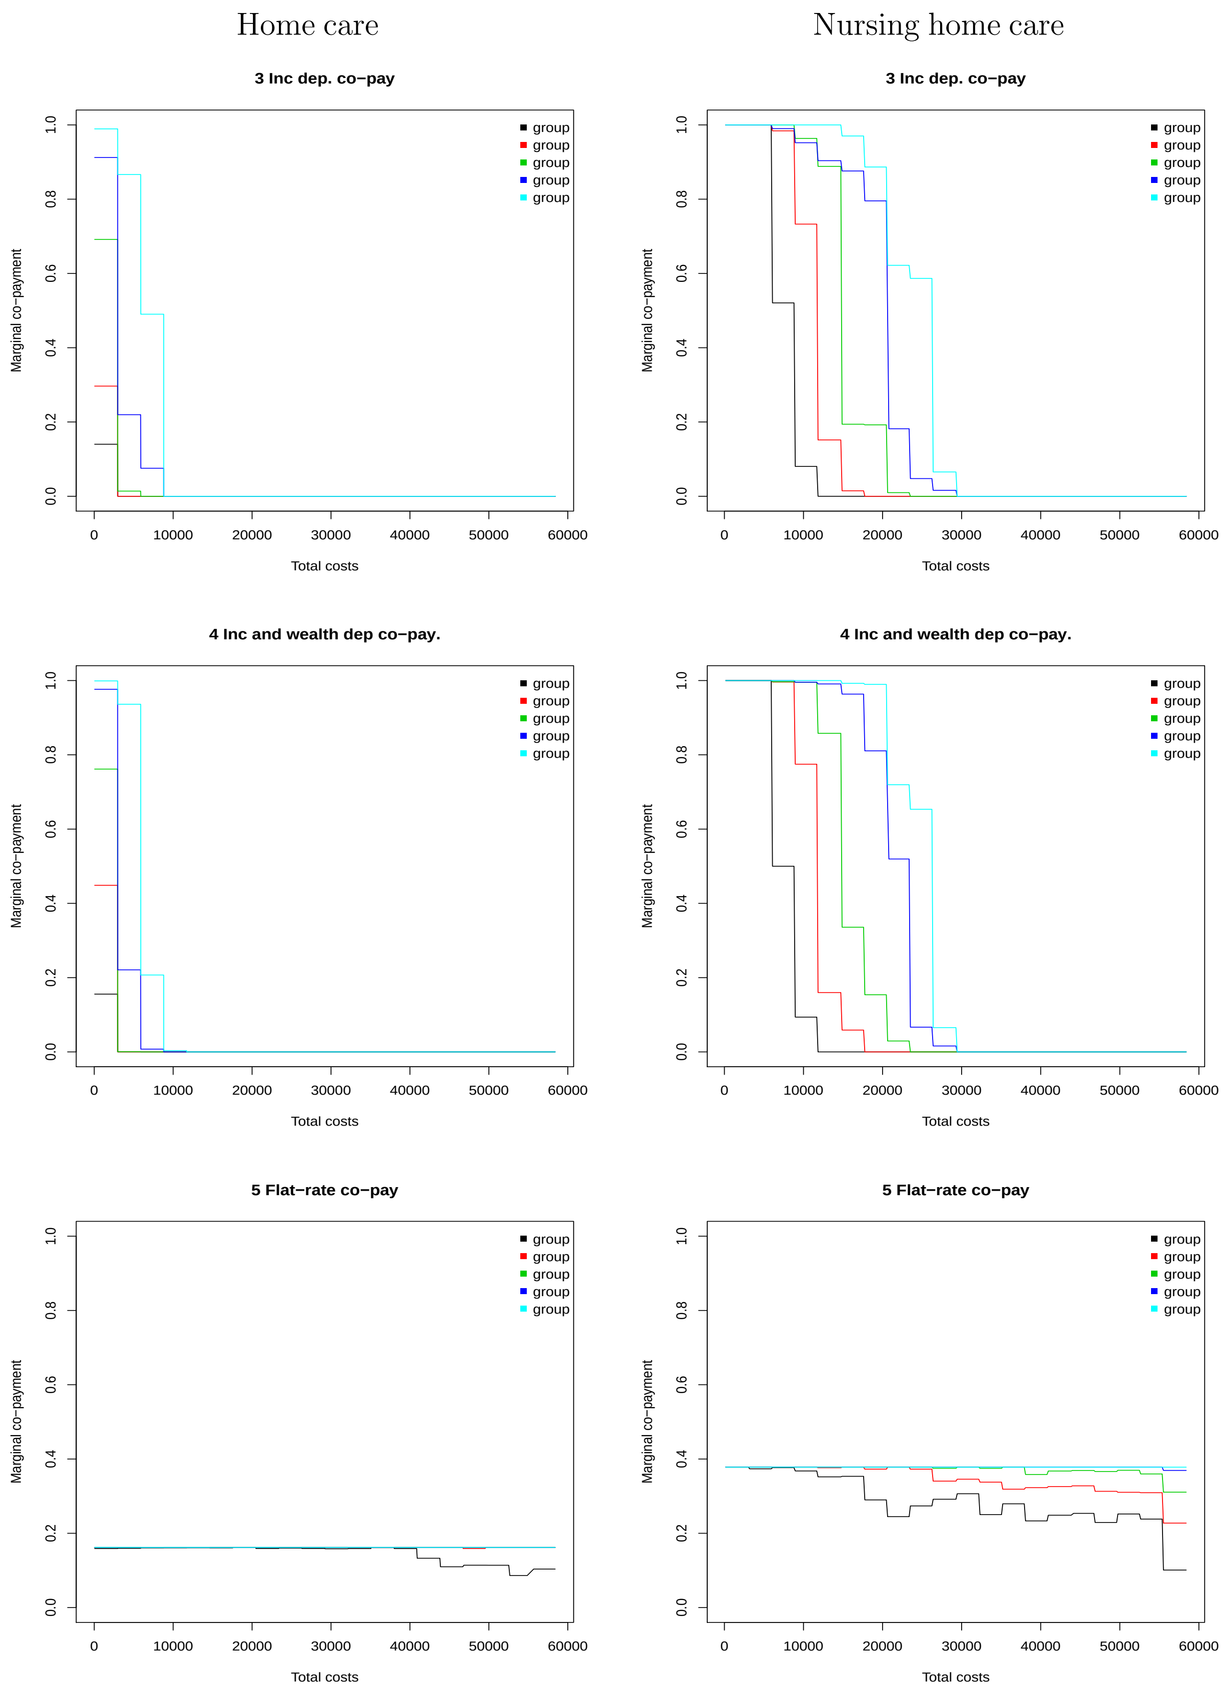

Figure 4: The marginal annual co-payment for different levels of LTC use in each policy variant, by total pension wealth group.

6 Discussion

Co-payments in social LTC insurance are a way to increase private financing and keep public spending in check. We have evaluated the distributional effects of different co-payment system. We have focused on the Dutch system, where LTC co-payments are based on a fixed share of income and wealth,a nd compared this to other variants, raising the same revenue, based only on income or completely independent of the financial means of the users. We have used a lifecycle model for singles at the retirement age. The model has allowed us to take the effects on risk and saving behavior into account.

We have estimated the longitudinal dynamics of LTC costs in the Netherlands, using the semi-parametric nearest-neighbor approach. We have shown that this approach has advantages compared to parametric approaches, especially because of its flexibility and the possibility to easily include the relationship between dynamics in LTC costs and other variables of interest, such as income. We have also shown that, despite its flexibility, it can still be applied in the context of a structural lifecycle model using the simulation based numerical optimization procedure developed by Koijen et al. (2010).

We have found that the Dutch system of income- and wealth dependent co-payment drastically redistributes the costs of co-payments from the elderly with the lowest financial means to those with the highest means. A co-payment based on a fixed share of LTC costs, independent of income and wealth and raising the same revenues as the current system, would lead to substantial losses in certainty equivalent consumption for the lowest income groups: compared to full insurance, elderly in the second total pension wealth quintile would loose 15 percent of their certainty equivalent consumption, while under the current co-payment system this is only 5 percent. At the same time, for elderly in the highest wealth quintile the loss of certainty equivalent consumption, compared to full insurance, in the current system is 0.83 percent, and this would only slightly decrease in the case of fixed co-payments.

Our analysis underlines the point of McClellan and Skinner (2006) that including risk in the analysis of redistributional effect of care systems is important. Especially for low income groups, the welfare losses induced by co-payments, as measured by the drop in certainty equivalent consumption compared to full insurance, are much larger than the drop in average consumption. Including the effects on saving behavior is also relevant: as a reaction to co-payments (especially in case of a flat-rate) individuals increase precautionary savings. Moreover, when co-payments depend on the level of private wealth, co-payments function as an implicit tax on wealth, inducing individuals to reduce their savings.

We have focused on the distributional effects of co-payments during the retirement phase. Not modeling the working phase of life helps to keep the model tractable and computationally manageable. This has enabled us to include a relatively large amount of detail, both in the dynamics in LTC costs and the policy variants. The costs of not modeling the working life phase is that we might overestimate the welfare losses due to co-payments, as individual might increase saving before retirement as a precaution. At the same time, the current generation of elderly has not been able to anticipate the introduction of the co-payments. In that sense, a lifecycle model starting at the retirement age better reflects their situation.

We have treated LTC costs as exogenous shocks. Although this is not an uncommon approach (e.g. Peijnenburg et al. (2017)), there are studies that do include endogenous use of health care (e.g. De Nardi et al. (2016b)). In theory, co-payments reduce inefficient use of care (moral hazard) as they increase the marginal costs of care use for the patient. An issue in modeling LTC use as an endogenous decision is that the empirical literature on the effects of co-payments (or prices) on LTC use is very limited (Konetzka et al., 2014). Grabowski and Gruber (2007) find no evidence of a moral hazard effect within the context of Medicaid. Studies looking at private LTC insurance do find that having LTC insurance raises use of LTC (Li and Jensen, 2011; Konetzka et al., 2014). These studies provide insights in the effect of (full) insurance versus no insurance, but tell little about two things that are crucial in comparing the three co-payment schemes that we analyze: the income elasticity of LTC demand and the effect of marginal price at different levels of LTC use.

Theoretical findings, on the hand, suggest that co-insurance rates should be high in good health states, in which elasticity of demand is high, and low in poor health states, in which elasticity of demand is low (Drèze and Schokkaert, 2013; Blomqvist, 1997). This suggests that the income- and wealth-dependent co-payment schemes, with marginal prices of 1 at low levels of use and 0 at high levels, might be more effective in reducing moral hazard than a fixed-rate co-payment. On the other hand, the flat-rate might be more optimal as it puts relatively high co-payments on groups with low financial means. In standard models (e.g. De Nardi et al. (2010, 2016b)), these groups are more price sensitive (as their opportunity costs are higher). Again, this insight seems of limited value to assess effects on use in the Dutch context. High income groups might for instance be more price sensitive for public LTC, as they have more possibilities to substitute with private care.

7 Conclusion

Income and wealth dependent co-payments provide substantially more value of insurance than flat-rate co-payments, that do not depend on the financial means of the LTC user. Not only for the elderly with low financial means, but also for elderly in the middle groups. Elderly with little financial means benefit from an income- and wealth-dependent co-payment, compared to a flat-rate co-payment, both because their LTC payments are lower on average and because they are exposed to less financial risk. Elderly with more financial means have to pay more on average, but only for the twenty percent of the elderly with the highest means does this outweigh the costs of the additional risk that comes with the flat-rate co-payment. Unless one expects that a flat-rate co-payment leads to substantially less distortions during working life, or substantially decreases moral hazard, the welfare case for an income- and wealth-dependent co-payment seems strong.

References

Aaberge, R. and Mogstad, M. (2015). Inequality in current and lifetime income. *Social Choice and Welfare*, 44(2):217–230.

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.

Bakx, P., O'Donnell, O., and Doorslaer, E. (2016). Spending on health care in the netherlands: not going so dutch. *Fiscal Studies*, 37(3-4):593–625.

Blomqvist, Å. (1997). Optimal non-linear health insurance. *Journal of Health Economics*, 16(3):303–321.

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.

Brown, J. R. and Finkelstein, A. (2007). Why is the market for long-term care insurance so small? *Journal of Public Economics*, 91(10):1967–1991.

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.

CPB (2017). *Keuzes in kaart 2018-2021*. CPB, The Hague.

De Nardi, M., Fella, G., and Pardo, G. P. (2016a). The implications of richer earnings dynamics for consumption and wealth. Technical report, National Bureau of Economic Research.

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.

De Nardi, M., French, E., and Jones, J. B. (2016b). Medicaid insurance in old age. *The American Economic Review*, 106(11):3480–3520.

De Nardi, M., Pashchenko, S., and Porapakkarm, P. (2017). The lifetime costs of bad health. Technical report, National Bureau of Economic Research.

Drèze, J. H. and Schokkaert, E. (2013). Arrows theorem of the deductible: moral hazard and stop-loss in health insurance. *Journal of Risk and Uncertainty*, 47(2):147–163.

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

Finkelstein, A., Luttmer, E. F., and Notowidigdo, M. J. (2013). What good is wealth without health? the effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, 11(suppl_1):221–258.

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

Grabowski, D. C. and Gruber, J. (2007). Moral hazard in nursing home use. *Journal of Health Economics*, 26(3):560–577.

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

Hurd, M. D., Michaud, P. C., and Rohwedder, S. (2017). Distribution of lifetime nursing home use and of outofpocket spending. *PNAS*, 114(37):9838–9842.

Hussem, A., Ter Rele, H., and Wouterse, B. (2017). Inkomens- en vermogensafhankelijke eigen bijdragen in de langdurige ouderenzorg een levensloopperspectief. Design Paper 87, 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.

Jones, J. B., De Nardi, M., French, E., McGee, R., and Kirschner, J. (2018). The lifetime medical spending of retirees. Technical report, National Bureau of Economic Research.

Khwaja, A. (2010). Estimating willingness to pay for medicare using a dynamic lifecycle model of demand for health insurance. *Journal of Econometrics*, 156(1):130–147.

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.

Konetzka, R. T., He, D., Guo, J., and Nyman, J. A. (2014). Moral hazard and long-term care insurance. *Health & Healthcare in America: From Economics to Policy. Ashecon*.

Kopecky, K. A. and Koreshkova, T. (2014). The impact of medical and nursing home expenses on savings. *American Economic Journal: Macroeconomics*, 6(3):29–72.

Li, Y. and Jensen, G. A. (2011). The impact of private long-term care insurance on the use of long-term care. *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, 48(1):34–50.

McClellan, M. and Skinner, J. (2006). The incidence of medicare. *Journal of Public Economics*, 90(1):257–276.

Meyer, B. and Mok, W. K. (2009). Disability, earnings, income and consumption. Technical report, National Bureau of Economic Research.

Non, M. (2017). Co-payments in long-term home care: do they affect the use of care? CPB Discussion Paper, CPB Netherlands Bureau for Economic Policy Analysis.

Palumbo, M. G. (1999). Uncertain medical expenses and precautionary saving near the end of the life cycle. *The Review of Economic Studies*, 66(2):395–421.

Peijnenburg, K., Nijman, T., and Werker, B. J. (2017). Health cost risk: A potential solution to the annuity puzzle. *The Economic Journal*, 127(603):1598–1625.

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)*.

Wouterse, B. and Smid, B. (2017). How to finance the rising costs of long-term care: Four alternatives for the Netherlands. *Fiscal Studies*, 38(3):369–391.

A Purging period effects from the source data

To purge the data from period effects, we rescale the values of the LTC, income, and wealth variables in earlier years to 2013 levels. As period effects may not only affect the mean, but also the shape of the distribution, we perform the following procedure for each variable. First, we divide the variable in 200 quantiles for each year and spline these quantiles (cubic splines with 10 knots). Then, we regress the variable on these splined quantiles. This gives a smooth estimate of the value of the variable over its entire distribution for each year. Finally, we use differences between the estimated value in the year of the observation and in 2013 to determine a scale factor for each quantile. We scale the original values of the variable using these scale factors. For financial wealth and income we directly use this procedure. For the variables with a lot of zeros (housing wealth, LTC use), we use a two-part procedure. We first determine sampling weights, using the difference between the share of zeros in the year of observation and 2013, and then use the regression on quantiles for the observations with values greater than zero.

B Assessment of the model fit

In this section we describe an assessment study to assess the performance of our proposed algorithm. The assessment follows the one described by Wong et al. (2016), in which the simulated life cycles generated for the assessment fall within the period covered by the data, such that removal of period effects is not necessary.

B.1 The assessment study setup

For a group of individuals with gender g and age a in year 1 (2009), we simulate the next four years of the life cycles (2010,..., 2013). The first observations of the lifecycles are obtained by drawing randomly from all the individuals with gender g and age a in year 1. The lifecycles are then extended to year $t+1$ by making use of the individuals history up to year t. We make use of the settings as described in the methods section (number of nearest neighbors $k = 2$, lag history for categorical variables 2 and lag history for continuous variables is 1). The simulated life cycles are then compared with the actual cycles of the individuals with age a in 2010 in several ways. For the continuous variables, the marginal distributions for each calendar year, the distribution of the sum of the variables over the entire period, and the serial correlation matrix over the entire period are examined. We also compare the mortality rate in 2013.

B.2 Results from the assessment study

Shown are the results for 85-year old women in 2009. For household income, wealth and home care costs, the Q-Q plots (Figure 5) generally reveal a reasonable agreement between the simulated and observed lifecycles in terms of the marginal distributions per calendar, the sum over the entire period. Exception herein are the costs for nursing home care, in which there is substantial downwards bias for the simulated costs. We found that this was persistent for several choices of our algorithm settings. Upon further inspection, we found that the quality of the nearest neighbors is high (in terms of very small distances between observed and expected lifecycles), so the algorithm essentially works well. We suspect that the high degree of variance and skewness of the nursing home care costs (moreso than other variables) are a reason for this. As was documented previously by Wong et al. (2016), the performance might suffer when the actual underlying distribution is heavy tailed. Otherwise, we find that the serial correlations are also reasonably similar, even for the nursing home care costs(see Table 6). The simulated mortality rate also corresponds well with the observed mortality rate.

Table 6: Serial correlation in nursing home care costs for women of 85 in 2010. In the source data and in the simulated data.

Source data				
	2010	2011	2012	2013
2010	1	0.75	0.65	0.55
2011	0.75	1	0.79	0.68
2012	0.65	0.79	1	0.83
2013	0.55	0.68	0.83	1
Simulated data				
	2010	2011	2012	2013
2010	1	0.82	0.76	0.69
2011	0.82	1	0.87	0.77
2012	0.76	0.87	1	0.87
2013	0.69	0.77	0.87	1

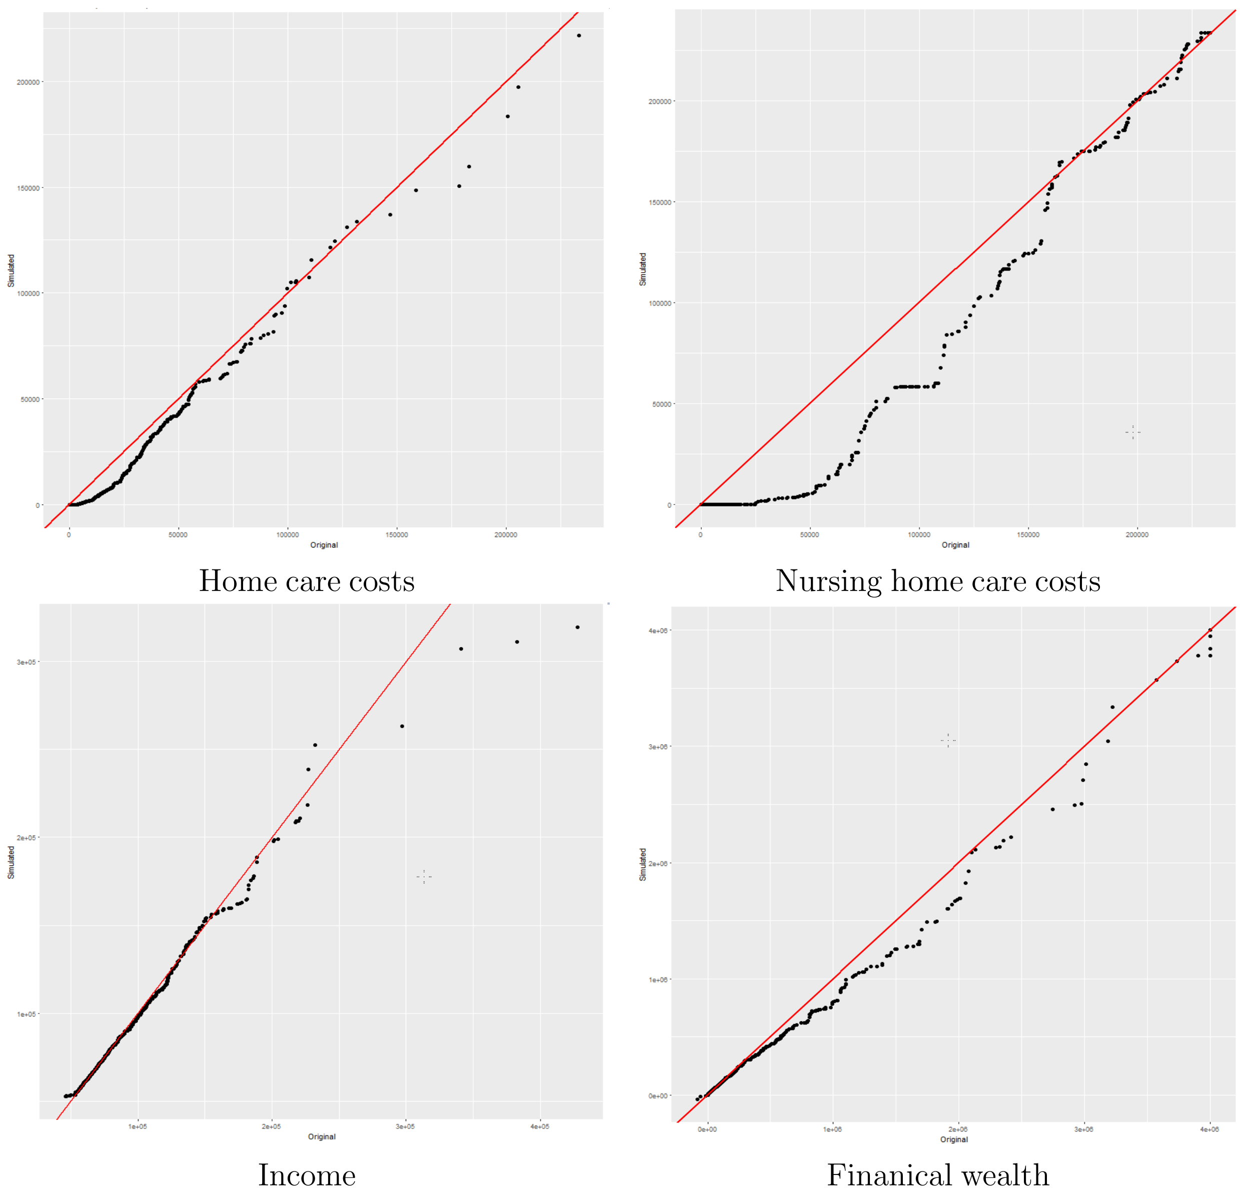

Figure 5: QQ plots of the source data versus the simulated data for women of 85 in 2010. LTC costs and wealth are conditional on having any costs or wealth.

C Numerical approach

The basic model

We solve the maximization prblem using adynamic programming approach developed by Koijen et al. (2010). The lifecycle optimization problem is divided into smaller yearly optimization problems, suing Bellman equations. In each period the optimization problem can be written as

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

The algorithm starts at the last time period T, and is then solved backwards recursively

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. We estimate this equation using a GLM with log-link, to ensure that the estimated values are strictly positive.

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.[9]

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.

Extensions

The inclusion of a bequest motive and state-dependent utility in the optimization procedure is relatively straightforward. The same is true for including a fixed pension income. The policy variants require an adaptation to the numerical approach. Wealth-dependent co-payments put an implicit tax on savings. Individuals have to include this tax when making decisions on current consumption. Specifically we adapt the Euler equation (14) by including the expected marginal implicit tax on wealth.

[9] As pointed out by Koijen et al. (2010) the endogenous grid method facilitates the use of this regression based approach. The optimal consumption can be derived analytically from the Euler equation (14) once the conditional expectation is known, so we do not need to determine this numerically. Else we would have to run for each gridpoint a (non-linear) regression at each iteration of the numerical optimization process, instead of just once for each gridpoint.

D Additional results

Table 7: CEC for individuals in the 1st, 3rd, or 5th income quintile (rows), and 1st, 5th, or 10th wealth quintile (columns). By policy variant. Main specification

	-116	13,822	301,120
1 Uniform premium			
13,707	12,846	12,423	27,019
21,572	20,222	18,935	34,300
46,200	43,320	39,577	56,010
2 Act. fair premium			
13,707	11,678	11,387	25,871
21,572	19,921	18,669	34,007
46,200	45,106	41,149	57,689
3 Inc dep. co-pay			
13,707	12,210	11,874	26,705
21,572	16,383	17,504	33,292
46,200	37,837	39,652	55,888
4 Inc and wealth dep co-pay.			
13,707	12,210	11,958	25,579
21,572	17,036	17,918	32,407
46,200	37,986	39,844	55,981
5 Flat-rate co-pay			
13,707	10,486	10,606	22,727
21,572	15,777	16,356	31,413
46,200	38,932	40,472	56,645