

Saving after Retirement and Preferences for Residual Wealth*

Giulio Fella, Martin B. Holm, and Thomas M. Pugh[†]

January 2024

Abstract

We use administrative data for Norway to estimate an incomplete-market life-cycle model of retired singles and couples with a bequest motive, health-dependent utility, and uncertain longevity and health. We allow the parameters of the bequest utility to differ between households with and without offspring. Our estimates imply a very strong utility of residual wealth (bequest motive), in line with the estimates by [Lockwood \(2018\)](#). The bequest motive accounts for approximately three-quarters of aggregate wealth at age 85. More surprisingly, we estimate similar utility of residual wealth for households with and without offspring. We interpret this as, *prima facie*, evidence that the utility of residual wealth represents forces beyond an altruistic bequest motive.

JEL: D11, D12, D14, E21

*The views expressed in this paper are solely those of the authors and not necessarily those of the Bank of Canada. Martin B. Holm acknowledges support from the ERC under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 851891).

[†]Giulio Fella: Queen Mary University of London, University of Bologna, CFM and IFS, g.fella@qmul.ac.uk. Martin B. Holm: University of Oslo, m.b.holm@econ.uio.no. Thomas M. Pugh: Bank of Canada, thomas.m.pugh89@gmail.com.

1 Introduction

The average retired household dissaves at a much lower rate than predicted by the standard life-cycle model.¹ Furthermore, richer retirees, both in terms of wealth and income, dissave at lower rates than poorer ones, to an extent that cannot be accounted for by their lower mortality rates.

The potential explanations for these patterns fall into two main categories (De Nardi, French and Jones, 2016b). The first emphasizes the risks of medical and long-term care expenses late in life and their interaction with differential mortality risks between richer and poorer retirees. The low dissaving would result from precautionary asset accumulation to self-insure against such risks (Palumbo, 1999; De Nardi, French and Jones, 2010; French and Jones, 2011; Kopeccky and Koreshkova, 2014; Nakajima and Telyukova, 2024). This saving motive is stronger for richer individuals because they are less likely to be covered by means-tested government-provided insurance of medical and long-term care expenses (e.g., Medicaid in the US, De Nardi, French and Jones, 2016a), and because they have longer expected lifetimes.

The second class of explanations emphasizes bequest motives. Namely, individuals may derive utility from accumulating or donating their wealth, possibly but not necessarily at death. This implies an extra incentive to accumulate wealth and therefore higher wealth levels at death over and above any accidental bequests stemming from the inability to annuitize longevity risk. To the extent that the utility from bequests is non-homothetic, in the sense that bequests are a luxury good, the bequest motive can also account for the differential saving patterns between poorer and richer retirees.

The two motives have similar implications for wealth accumulation late-in-life.² For this reason, a consensus on their relative importance has not yet been reached. As shown by Dynan, Skinner and Zeldes (2002) and elaborated in De Nardi et al. (2010), the problem is one of identification. Adding a luxury bequest motive to a model featuring medical and long-term care expense risk does not change saving patterns significantly. For this reason, a number of papers have instead focused on other data moments beyond saving to separately identify the two motives. For example, Ameriks, Caplin, Laufer and Van Nieuwerburgh (2011) and Ameriks, Briggs, Caplin, Shapiro and Tonetti (2020) use strategic survey questions to identify the bequest utility parameters. Similarly, Lockwood

¹Nakajima and Telyukova (2015) document that the pattern for the *median* retired wealthholder displays substantial variation across countries.

²Housing is a third saving motive often emphasized in explaining savings behavior after retirement (Nakajima and Telyukova, 2015, 2017, 2020; McGee, 2021; French, Jones and McGee, 2023). We abstract from housing in the current paper because we observe similar wealth profiles for homeowners and renters in our data.

(2018) exploits long-term care insurance decisions to separately identify the precautionary saving and bequest motives.

This paper follows a third approach. It estimates a life-cycle model of retirees' savings using population wealth data from Norway. Because medical and nursing home expenses are effectively fully insured in Norway, this leaves only longevity risk as a driver of precautionary saving. Moreover, our administrative dataset has four main advantages compared to prior studies. First, it has universal coverage and is highly accurate since wealth is third-party reported for tax purposes. As a result, we can document, and exploit for our estimation, life-cycle profiles of wealth at various percentiles of the distribution, for singles and couples, for households with and without offspring and conditional on health status. Second, because inheritances were taxed over the sample period we observe inter-vivos transfers from parents to offspring. These transfers are substantial and quantitatively relevant when estimating wealth age profiles for retirees. Third, we directly observe individuals moving into long-term care, allowing us to precisely account for how a deterioration in health and accession to nursing homes affect wealth accumulation. Finally, the relatively long panel dimension of the data allows us to precisely estimate the relevant wealth age profiles, dealing properly with time, age, and cohort effects.

Equipped with precisely estimated wealth profiles from age 70 to 90 within 19 groups, we build an incomplete-markets life-cycle model of retired singles and couples with a bequest motive and health-dependent utility. Specifically, our model includes state-dependent health and mortality risk, returns risk, and the important components of the tax and transfer system in Norway, including wealth and income taxes and user payments in long-term care. Importantly, and contrary to most of the literature, we allow the parameters of the bequest utility to differ between households with and without offspring.

Our main result is that our estimates imply a very strong utility of residual wealth (bequest motive), in line with the estimates by Lockwood (2018). The "bequest motive" accounts for three-quarters of aggregate wealth at age 85. More surprisingly, our estimate of the utility of residual wealth is very similar for households with offspring as those without. We interpret this as strong, *prima facie*, evidence that the utility of residual wealth is driven by forces beyond an actual bequest motive.

We also estimate a lower marginal utility of consumption for individuals who are in a nursing home. *Ceteris paribus*, it is about 90 percent of the marginal utility of individuals living autonomously. The intuition behind our findings is that households in a nursing home have a substantially steeper wealth-age profile compared to households living autonomously. Through the lens of the model, the higher rates of wealth growth require

the former group to have a higher marginal utility of residual wealth relative to current consumption.

Finally, we use the model to assess the wealth response of the elderly population to changes in wealth and estate taxes. The model implies an elasticity of aggregate saving with respect to the after-wealth-tax rate of about 4.8, which is in the ballpark of comparable empirical estimates (Brülhart, Gruber, Krapf and Schmidheiny, 2022; Jakobsen, Jakobsen, Kleven and Zucman, 2020). The response is mostly due to the (mechanical) impact of the change in the after-tax return of wealth accumulation with little change in saving. For the same reason, the model implies that the abolition of estate taxation would reduce tax revenue with hardly any effect on wealth accumulation.

Related literature. Our paper contributes to three strands of literature. The first one is to contribute to the identification of the role of the utility of residual wealth in accounting for the low rates of dissaving of the elderly relatively to the prediction of the life-cycle model (the retirement saving puzzle). De Nardi et al. (2016b) provide a recent survey of the literature on the puzzle.³ Our approach is similar to Christensen, Kallestrup-Lamb and Kennan (2022) in that we address this identification problem by focusing on a country where long-term care expense risk is negligible, thus allowing us to identify the bequest motive.⁴ However, while Christensen et al. (2022) restrict their focus to a very small and selected sample of single individuals not owning homes in Denmark, we estimate the bequest motive on *all* types of households in Norway. In common with Braun, Kopecky and Koreshkova (2016), Nakajima and Telyukova (2020) and De Nardi, French, Jones and McGee (2023) we model couples in addition to single households. The distinction is important due to the differences in wealth at retirement, health and life expectancy by marital status, sex and age in the data, as well as for the additional income risk associated with spousal death.⁵

The second goal of the paper is to understand the extent to which the utility of residual wealth captures, in a reduced form fashion, other forces beyond an actual bequest motive. We make progress in this direction by estimating separately the utility of residual wealth for households with and without offspring. This follows the early contribution of Hurd

³A related yet different literature attempts to use saving behavior also before retirement to identify bequest motives (Pashchenko and Porapakkarm, 2023).

⁴Similarly, Van Ooijen, Alessie and Kalwij (2015) presents wealth patterns by age in the Netherlands, a country where long-term care expense risk is also negligible, consistent with the patterns in our current paper.

⁵A less directly related paper, although it uses the same dataset, is Halvorsen, Hubmer, Ozkan and Salgado (2021), which estimates a life-cycle model over the whole working life augmented with a common to-all bequest motive. They use the model to assess the relative contribution of various forces, namely heterogeneity in lifetime income and returns, as well as a bequest motive, to wealth concentration.

et al. (1987) and Hurd (1989) who posited that differences in wealth accumulation between households with and without offspring provides a measure of the strength of the bequest motive. Contrary to us, they assume that only individuals with offspring have a bequest motive and, on the basis of similar rates of wealth growth for the two populations, concludes that the bequest motive is economically negligible.

Kopczuk and Lupton (2007) show that the findings in Hurd et al. (1987) and Hurd (1989) are due to the assumption of zero bequest motive among households without offspring. They instead assume a *common* strength of the bequest motive in the two populations but allow for its presence to reflect unobserved heterogeneity. They estimate the motive to be economically meaningful and its incidence higher for households with offspring, although the difference is only marginally statistically significant. Our approach assumes that the bequest motive affects both populations, but we allow the strength to differ. We find that the bequest motives for households with and without offspring are very similar, almost indistinguishable, suggesting that the utility from residual wealth is driven by forces beyond an altruistic bequest motive.

Our paper is also related to the literature that studies the interaction between health shocks and consumption and saving. Finkelstein, Luttmer and Notowidigdo (2013) argue that people who are in bad health have a lower marginal utility of consumption, on the basis of survey responses to questions about happiness. Blundell, Borella, Commault and Nardi (2020) find that consumption responds negatively to temporary health shocks with most of the effect being accounted for by reductions in the marginal utility of consumption rather than in available resources. Christensen et al. (2022) and Ameriks et al. (2020), which both estimate a life-cycle model with health-related fluctuations in the marginal utility of consumption, are the papers closest to ours. Differently from ours, they both restrict attention to single households. Foltyn and Olsson (2023) document systematic health-related biases in subjective survival probabilities and explore their implications for the health-wealth gradient in the elderly population. Similar to most of the aforementioned studies, we find, using a sample covering the whole population, that people in bad health behave in a way consistent with a lower marginal utility of consumption.

Roadmap. The paper is organized as follows. Section 2 describes the data and the institutional setting. Section 3 introduces our structural model. Section 4 discusses the estimation methodology, while Section 5 presents the parameter estimates. Section 6 uses the model to assess the contribution of the bequest motive on the wealth accumulation of the elderly, as well as the response of the latter to wealth and estate taxation. Section 7 concludes.

2 Data and institutional setting

This section is organized as follows. We first present the relevant institutional settings for our analysis. Next, we introduce the data sources, variable definitions, sample restrictions, and summary statistics applied to the administrative data from Norway. Lastly, we discuss the main calibration targets: wealth profiles by age for different household subsets.

2.1 Institutional setting

Our analysis relies on the argument that nursing-home or health-related expenditure risk are not important factors in explaining wealth savings after retirement in Norway. We here explain the relevant parts of the Norwegian nursing home and healthcare system.

Nursing homes. The near totality of long-term care is publicly financed in Norway.⁶ In practice, individuals can choose among nursing homes, both private and public, all financed within the public system. Within the context of the publicly-funded system, many municipalities allow individuals to choose from a list of public and approved private nursing homes. In this manner, the private sector accounts for around 11.6 percent of long-term care provision (Førland, Ambugo, Døhl, Folkestad, Rostad and Sundsbø, 2020).⁷ There exists only a very limited number of private long-term alternatives where individuals pay themselves.⁸ These fully private alternatives are very expensive and relevant only for a selected few, which would belong to the top 1 percent wealthiest households which are excluded to our analysis below. The relevant long-term care option for the rest of the population is the publicly funded system which is financed through a progressive system of user fees detailed in Section 4.

Health expenses. Norway has an effectively Beveridgean public health care system where the government provides health care for its citizens financed by taxes. All approved treatments are covered by the public system, and there is limited user payment. According to the 2012 Norwegian survey of consumer expenditure, the share of total medical expenses out of personal income for individuals older than 70 years was around 5 percent. It was mostly accounted for by dental services and auxiliary equipment (3

⁶The Hospital Act of 1970 made local counties (municipalities from 1988) responsible for running nursing homes for those who need care, including long-term care for the elderly.

⁷This includes both institutions run by religious or humanitarian organizations and private companies.

⁸There is no official statistic on the extent of fully private nursing homes. However, occasionally, newspaper articles present luxury versions of nursing homes with prices of around NOK 1 million (around \$100,000) per year, see, e.g., tv2.no and a-magasinet.

percentage points) and out-of-pocket drug expenses (1 percentage point). These shares reflect the fact that, while there exist private alternatives—e.g., for relatively uncomplicated emergency room tasks or general practitioners—individuals tend to be treated in the public health care system whenever the treatment is expensive. This motivates us to abstract from health spending and health spending risk altogether in the model in Section 3.

2.2 Data

We use administrative data from Norway. The data contain information from tax records on the wealth and income of all individuals with tax residence in Norway. Because Norway levies a wealth tax, the tax data contains detailed household portfolio information by broad asset classes. Prior to 2014, there was also an inheritance tax in place in Norway. We therefore restrict our sample to the years from 2005 to 2013 because we use the inheritance tax data to measure the extent of inter-vivos transfers and transfers at death. The data are matched to population data with information on individual characteristics such as birth year, sex, marital status, and number of children. Finally, we combine the address registry and the population censuses to observe whether individuals move into long-term care. We additionally use the administrative data on deaths to compute expected mortality rates within subgroups of the population.

Variable definitions. Our model features households who consume and save out of their income and who differ by age, sex, health, marital status, and health. Beyond these traits, we use data on income and wealth to calibrate our model.


Permanent income. Permanent income at retirement captures in a parsimonious way households' ex-ante heterogeneity (e.g. [De Nardi et al., 2016a](#)). Households with different permanent income ranks not only receive different flow of retirement income but face different processes for health and mortality. We use gross pension income as our measure of *gross* permanent income. Pension income includes income from social security and private and public occupational pension.

Although there is some flexibility in when an individual receives a pension, from the age of 70, almost all individuals are retired, and there is limited time variation in annual pension income except for annual indexation. We therefore use pension income at age 70 within the cohort to define an individual's position in the gross permanent income distribution.

While permanent income, as captured by gross pension income, is invariant over the lifetime, the mapping from gross pension income to disposable income reflects two

additional sources of heterogeneity. The first is that the total tax burden depends on taxes other than that on pension income, such as capital income and wealth. The second is that we need to consider user payments for long-term care, as discussed in Section 4.

Wealth. The second main variable is household wealth. We define wealth as the sum of net worth (housing equity + stocks + stock funds + bonds + deposits + private business wealth + ownership of private loans - debt) and the present value of inter-vivos transfers. For most wealth variables, the data contains market values. A notable exception is housing equity, for which we use the estimated valuations from Fagereng, Holm and Torstensen (2020) because the tax values are imprecise prior to 2010. The other exception is private business wealth which is valued according to a criterion close to book value (see the discussion in Fagereng, Holm, Moll and Natvik, 2021). We sidestep this issue by restricting our sample to the bottom 99 per cent of wealth holders for whom the portfolio share of private business wealth is negligible.



Notes: The figure shows average wealth relative to wealth in t-1 for the two groups, computed as $w_{g,t+h}/w_{p,t-1} - 1$ for group = offspring or parent. Offspring wealth is defined as the sum of all household wealth of all households that includes any children of the parents.

Figure 1: Parent and offspring wealth around parent spousal death event

One issue that needs addressing is inter-vivos transfers. The average cumulated value of inter-vivos transfers is sizable and, conceptually, is akin to end-of-life bequests. One way to account for inter-vivos transfers is to allow for such transfers within the model.

However, this complicates the model significantly. Instead, we choose to keep the model relatively simple, by only allowing for end-of-life bequests, and account for inter-vivos transfers as part of these. To this effect, we carry over the observed inter-vivos transfers from the estate tax registry to the end of life, compounding them at the average real return in our sample in the years thereafter (1.72 percent).

A complication is that even after we correct for observed transfers from the estate tax registry, we estimate a large average wealth drop when a member of a couple dies. Figure 1 shows parent and offspring wealth around a parent death event where there is a surviving spouse. On average, the parent wealth drop is around 25 percent and offspring wealth increases by approximately 25 percent in the aftermath of the spousal death event.⁹ We interpret this as a transfer from parents to offspring and include it as part of inter-vivos transfers.¹⁰ Because we cannot precisely observe this wealth drop at the individual level, we impute inter-vivos transfers when a member of the couple dies as 25 percent of household wealth in the previous year.

Sample restrictions. We restrict attention to households whose head is between 70 and 90 years old in order to abstract from labor supply decisions. Among these households, we drop: (1) households with negative net wealth (2.4 percent) to enable estimation of log wealth profiles; (2) households in the top 1 percent of the wealth distribution within each cohort year (1 percent); (3) households for which the annual growth rate of wealth is greater than 100 percent in absolute value (2.7 percent); and (4) households whose gross pension income is below the minimum threshold (approximately \$10,000 per year) in the Norwegian social security system (0.08 percent of the population).¹¹ The final sample consists of 2,393,863 household-year observations between 2005 and 2013, approximately 265,000 households per year.

Descriptive statistics. Table 1 presents descriptive statistics for our sample of households. We highlight four observations. First, 63 percent of retired households between age 70 and

⁹There are tax incentives for transferring resources to offspring if the couple has children that are not all common. Specifically, an inheritance event is a transfer from an individual to another individual. In the case where your parent dies and there is a surviving step-parent, any transfers from your parent are counted as one inheritance event, while any remaining inheritance afterward is counted as an inheritance transfer from your step-parent. Hence, for tax reasons, one would want to make transfers upon parent death. In the case where an offspring is the offspring of both parents, all inheritance, also inheritance after parent death, from the parents to the children is counted as two inheritance events where half comes from each parent.

¹⁰This is a pattern of transfers similar to what has been documented for the U.S. (French, De Nardi, Jones, Baker and Doctor, 2006; Poterba, Venti and Wise, 2011; De Nardi et al., 2023; French et al., 2023).

¹¹Additionally, we exclude households residing in the counties of Troms and Finnmark (5 percent of the population) because of preferential tax treatment.

	Mean	S.D.	P10	P50	P90
<i>Panel A: Demographics</i>					
Age	78.69	5.74	71	78	87
Singles	0.63	0.48			
Male	0.38	0.49			
With children	0.86	0.35			
Homeowner	0.88	0.32			
In nursing home	0.05	0.22			
<i>Panel B: Wealth and Income (\$1'000)</i>					
Household net worth	827.9	674.2	186.7	668.0	1,628.0
Household net worth excl. inter-vivos transfers	737.8	628.8	118.2	594.0	1,477.9
Individual pension income	42.2	18.1	23.4	38.3	65.1

Notes: Descriptive statistics for our sample of individuals aged 70 to 90 in 2005 to 2013 expressed in 2011 U.S. dollars. The average PPP ratio exchange rate over the sample period is 11.35 Kr/\$.


Table 1: Descriptive statistics.

90 are single households. Second, the vast majority, 86 percent of retired households, have children. However, childless households still comprise a significant portion of our sample, allowing us to differentiate between households with and without offspring. Third, by comparing wealth with and without the net present value of inter-vivos transfers, we see that the inter-vivos transfers make up approximately 12 percent of total wealth. Inter-vivos transfers are disproportionately important for relatively low-wealth households. Fourth, home ownership is widespread: 88 percent of households in our sample own homes. It is even higher, at 96 percent, if we restrict attention to households at age 70 because individuals tend to move into nursing homes with age.


2.3 Main calibration targets: wealth by age

The main calibration target for our model is average wealth by age. A challenge in computing these age profiles is that one needs to distinguish between time, age, and cohort effects. Because time, age, and cohort dummies are perfectly collinear, identifying these effects relies on structural assumptions. We follow [Deaton and Paxson \(1994\)](#) and assume that the time effect is orthogonal to a linear time trend such that any linear time trend can be decomposed into either a cohort or an age effect.


To illustrate how the time-age-cohort decomposition operates, [Figure 2a](#) shows estimated age profiles of wealth and raw data for selected cohorts. In the raw data, some




(a) Average log wealth by age




(b) By wealth at 70



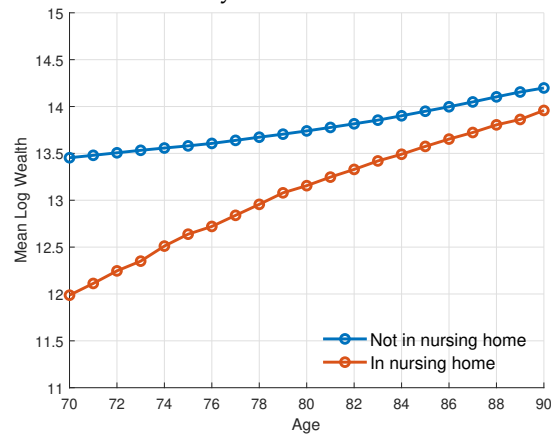
(c) By offspring



(d) By marital status



(e) By homeownership at 70



(f) By nursing home status

Notes: In Figures 2b, 2c, and 2e, the groups are defined at age 70, while the groups change over time in Figures 2d and 2f. The group ‘No offspring, bottom 20%’ in Figure 2c refers to households with no offspring that are in the bottom 20% of the wealth distribution within their cohort at age 70. Similarly, ‘No offspring, top 80%’ refers to households with no offspring that are in the top 80% of the wealth distribution at age 70.

Figure 2: Average log wealth by age.

movements are similar across all cohorts (all raw data profiles start in the same year in Figure 2a). These movements are interpreted as time effects and are taken out. Moreover, moving between cohorts, the wealth profiles shift. These shifts are interpreted as cohort effects and are also taken out. What is left is the average growth rate by age after taking out the time and cohort effects, the black circled line in Figure 2a, which starts from the initial value of a baseline cohort and grows with an averaged growth rate across cohorts.

The rest of Figure 2 displays age profiles of wealth within groups. For our purpose of motivating and estimating a household model, we highlight four key facts in Figure 2. First, wealth is approximately flat or monotonically increasing in age, both in the aggregate (2a) and within all population subgroups. This holds by percentiles of initial wealth (2b), with or without offspring (2c), by marital status (2d), by initial homeownership (2e), and by health as measured by admission to a nursing home (2f). A standard life-cycle model without a bequest motive cannot explain these features (without large medical expenses and long-term care risks). Hence, the data point toward the importance of a bequest motive.

Second, the wealth profiles by age are steeper for poorer households. The effect is driven by the poorest groups, whose wealth at 70 is relatively low. The evidence points toward a bequest motive that is not necessarily of the luxury kind.

Third, the average age profile differs between households with and without offspring. Panel 2c reveals that childless households have a steeper wealth profile (purple vs blue line). The difference between households with and without children has been used previously to identify the role of altruistic bequest motives (Hurd et al., 1987; Hurd, 1989; De Nardi et al., 2023). When the wealth profile is steeper for those without children, it suggests a limited role for altruistic bequest motives, something we will return to when we estimate the model. However, part of the increase in wealth by age for those with no children is driven by a large group of households in this group with relatively low wealth. Indeed, when we split the group of households with no offspring into those in the top 80% of the wealth distribution at age 70 and those in the bottom 20%, the wealth profile of the no offspring population in the top 80% of the wealth distribution is almost indistinguishable from the wealth profile the population with offspring.

Fourth, wealth is increasing with age, both for homeowners and renters. This finding contrasts the results for the U.S. in Nakajima and Telyukova (2020), who find that wealth declines with age for renters but is flat for homeowners. Given the relative similarity in growth rates, once the level differences are controlled for, it motivates our choice of abstracting from housing equity in the model.

The reported profiles in Figure 2 are not corrected for survivor bias. Within a cohort, it

is well-documented that wealthier households tend to live longer (Smith, 1999; Attanasio and Hoynes, 2000), and that this composition effect biases the cross-sectional wealth profiles (De Nardi et al., 2010). Appendix A shows that the wealth profiles for differential mortality rates by age, permanent income, sex, and health are flatter than the unadjusted ones. Yet, this does not change the fact that they are at odds with a standard life-cycle model. Since our structural model will explicitly include differential mortality rates by age, permanent income, sex, and health, we target the unadjusted wealth profiles in Figure 2 and show that the individual wealth profiles in the model also match the untargeted adjusted wealth profiles in the data.

3 The model

3.1 Environment

Time is discrete and the horizon is finite. Households start at age 70 and live at most 105 years. They may be composed of a single individual, of gender $g = f, m$, or a couple, husband and wife. They are heterogeneous in wealth $a \in [0, \infty)$, the pension income $y_g \in \{y_1, y_2, \dots, y_N\}$ of each household member, age t , marital status (“paired” or single) $j \in \{p, s\}$, whether they don’t or do have offspring $o \in \{0, 1\}$ and health status $h_g \in \{0, 1, 2\}$ of each household member. The three health states are good health ($h_g = 0$), in long-term nursing home care, NHC in what follows, ($h_g = 1$) and death ($h_g = 2$).

Households derive utility from consumption and the bequests they leave at the end of their life. They are subject to two sources of uncertainty: health and death, and the realisation of the rate of return on their savings. In every period, they choose consumption expenditure c and savings a' , taking as given their members’ constant pension income and the one-period rate of return on their savings. They cannot borrow.

3.2 Preferences

Preferences are time separable and households discount the future with inter-temporal discount factor β . The utility from consumption is affected by an individual’s health status. Healthy individuals— $h_g = 0$ —value consumption according to a standard CRRA felicity function

$$u(c; h_g = 0) = \frac{(c/v_j)^{1-\sigma}}{1-\sigma}, \quad \sigma > 0, \quad (1)$$

where v_j is a consumption equivalence factor which depends on marital status j .

The corresponding function for individuals in NHC— $h_g = 1$ —is

$$u(c; h_g = 1) = \delta \frac{(c/v_j)^{1-\sigma}}{1-\sigma}. \quad (2)$$

where the difference between (2) and (1) is δ . To the extent that the scaling factor δ is different from one, consumers in nursing home value consumption differently from their healthy counterparts. We assume that couples maximize utility in a unitary fashion, which implies that their felicity function is the sum of the individual utilities.

Upon the death of a household's last surviving member, any outstanding wealth is distributed as a bequest. Agents value, net-of-tax, bequests b according to a warm-glow utility function

$$v_o(b) = \theta_o \frac{(\kappa_o + b)^{1-\sigma}}{1-\sigma}, \quad (3)$$

where θ is the bequest loading and κ is a Stone-Geary shifter. Both parameters depend on whether the household has children or not ($o = 0, 1$).

3.3 Uncertainty

Households are exposed to two types of uncertainty. First, at the beginning of each period they draw a, household-specific, i.i.d realisation of the one-period rate of return on their saving. This allows for the possibility that some of their wealth accumulation is driven by luck. Our timing assumption, together with the one-period maturity, implies that the one-period rate of return is known with certainty (i.e. risk-free) at the time of investment, although households face uncertainty over subsequent periods. The rate of return process follows a stationary, first-order Markov process with transition probabilities $\pi(r)$.

The second type of uncertainty concerns individual's health status. Individuals may be hit by a negative health shock and either, irreversibly, enter a nursing home or die according to a first-order Markov process with transition probabilities $\pi(h'_g|h_g, t, g, j)$ that depend on current health, age, gender and family composition as well as permanent income and wealth at age 70 (y_g and $a_{t=0}$). We assume that the transition to NCH, but not death, is perfectly correlated for members of a married couple. Health status affects both the household's preferences, as discussed above, and the mapping from individual pension income y_g to disposable income $\phi(y_g; h_g)$, in a way we describe below.

3.4 Government

The government taxes individual income, wealth and bequests and provides an income floor \underline{y} for individuals in NCH ($h_g = 1$). Disposable income $\phi(y_g; h_g)$ is a function not only of pre-tax income y_g but also of health status to capture the fact that in Norway individuals in NCH are taxed at a higher rate than individuals with the same income but not in NCH. Wealth above the exemption level is taxed at rate $\tau(a)$. Finally, intra-household wealth transfers upon the death of one spouse are tax-exempt, but bequests originating from the death of the last surviving member of a household are subject to taxation. After-tax bequests are related to end-of-life wealth a through the tax schedule $b(a)$.

3.5 The individual problem

Let $z_s = (a, r, y_g, a_{t=0}, t, g, h_g, o)$ and $z_p = (a, r, y_f, y_m, a_{t=0}, t, h_f, h_m, o)$ denote the vectors of states respectively for singles and couples. The household optimization problem in recursive form for singles alive in the previous period¹² can be written as

$$V^s(z_s) = (1 - \mathbf{1}_{h_g=2}) \{ \max_{c, a'} u(c; h_g) + \beta E_{r', h'_g} V^s(z'_s) \} + \mathbf{1}_{h_g=2} \{ v_o(b(a)) \}$$

$$\text{s.t. } a' = (1 + r)(1 - \tau(a))a + \phi(y_g, h_g) - c \geq 0.$$

The value function has two main sub-components. The first one, which applies if the individual is alive in the current period ($h_g \neq 2$) is the current utility plus the expected continuation value next period. The expectation is taken with respect to the interest rate and the individual health status next period. The second main sub-component is the warm glow bequest utility in case the individual has died at the beginning of the current period ($h_g = 2$). All households are subject to a borrowing constraint ($a' \geq 0$).

The recursive problem for couples is

$$V^p(z_p) = (1 - \mathbf{1}_d) \{ \max_{c, a'} u(c; h_f) + u(c; h_m) + \beta E_{r', h'_f, h'_m} V^j(z'_j) \} + \mathbf{1}_d \{ v_o(b(a)) \}$$

$$\text{s.t. } a' = (1 + r)(1 - \tau(a))a + \phi(y_f, h_f) + \phi(y_m, h_m) - c \geq 0,$$

where $\mathbf{1}_d$ equals 1 if both members of the couple die at the beginning of the current period— $h_f = h_m = 2$ —and zero otherwise.

If the couple is alive in the current period it pools income and maximizes the sum of the spouses' utility plus the continuation value. The expectation reflects the possibility

¹²Note this can still be written recursively as $a_{t=0}$ is exogenous for the household.

that either spouse may be hit by a health shock. In the case in which one spouse dies, the continuation value function is that of the surviving member being single.

4 Estimation methodology

We adopt a two-stage strategy to estimate the model. In the first stage, we estimate or calibrate a first set of parameters which can be identified outside the model. In the second stage, we estimate the remaining parameters by the method of simulated moments (MSM), taking as given the parameters estimated in the first stage.

4.1 External estimates

This section presents the external estimates: mortality rates and nursing home accession rates, preferences, the returns process, and the tax schedules.

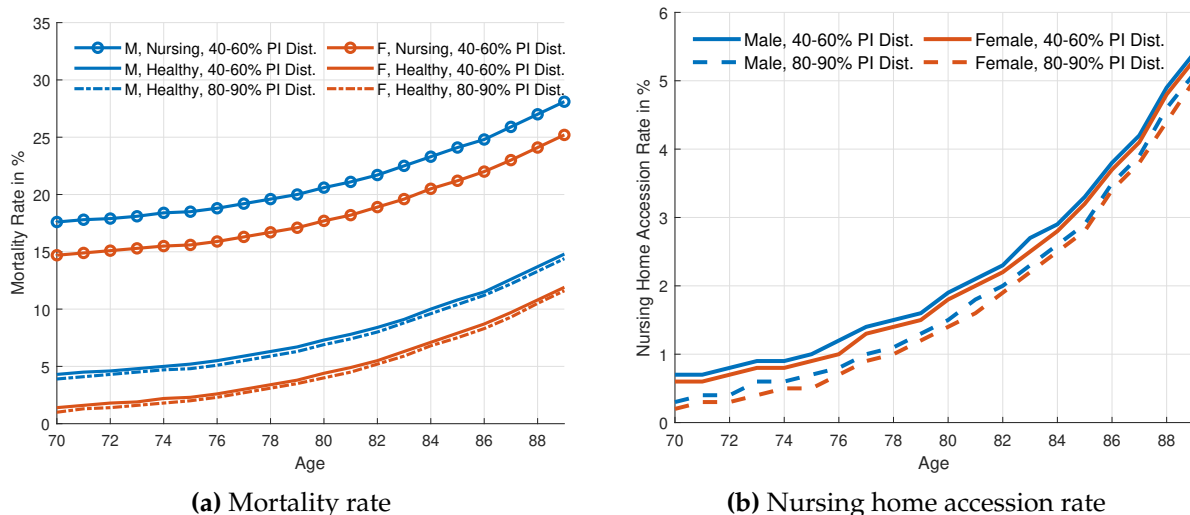


Figure 3: Mortality and nursing home accession rates for singles by age.

Demographics. We obtain demographic information from a number of sources described in Section 2. Our model starts out with the initial distribution of 70-year-old households by sex, health, children, wealth, and permanent income in Norway. We allow for three health states: healthy, nursing home, and dead. Health transitions follow an estimated Markov process where the transition probabilities depend on current health, age, sex, wealth at age 70 and permanent income.

Figure 3 displays a subset of these average rates. The mortality and nursing home accession rates are higher for males and poorer individuals, while being in long-term care raises the mortality rate. Men have, on average, about 3.5 percentage points higher mortality rates and around 0.05 percentage points higher nursing home accession rates than women. Higher permanent income is associated with reduced mortality and lower nursing home accession rates, but the effect is modest. Going from the median to the 80-90th percentiles in the permanent income distribution reduces mortality by around 0.3 percentage points and nursing accession rates by around 0.45 percentage points. Wealth at age 70 has a similar effect to permanent income, but is approximately twice as strong: going from the 40-60 percentile group to the 80-100 percentile group decreases mortality by approximately 0.9 percentage points. Mortality rates for individuals in long-term care are around 8 percentage points higher than those outside long-term care.

Preferences. We set the curvature of preferences over consumption and bequests σ to 2, a standard value (e.g. Kaplan and Violante, 2010), and the discount factor β to 0.9. It would be inappropriate to identify those parameters from post-retirement saving behavior alone. The chosen value for β is in line with Halvorsen et al. (2021) and Kopczuk and Lupton (2007).

As our model does not target consumption moments, we refer to the many external estimates for the value of the equivalence factor for couples ν , which controls the strength of economies of scale in household consumption. We choose a ν of 1.5, substantial economies of scale but well within the typical range—1.06-1.7—documented in the literature (e.g. Fernandez-Villaverde and Krueger, 2007, and references therein). This implies a couple living together has to spend, on average, \$1.5 to achieve the same level of utility as each member spending \$1 if living alone.

Returns to wealth. We estimate returns to wealth directly in our data set, defined as the sum of capital income and capital gains as a share of net wealth. In the appendix, we compare the data distribution with a normal distribution, provide evidence of low dependence of returns on wealth in our sample, and illustrate the low persistence of returns. The combined evidence suggests that assuming returns are normally distributed with a mean of 0.068 and a standard deviation of 0.086 is a reasonable approximation for our sample of elderly households for whom housing is the primary asset. Our assumption of i.i.d. normally distributed returns aligns with Fagereng, Guiso, Malacrino and Pistaferri (2020).

Taxes. Households are subject to five types of taxes: a tax on pension income, a capital income tax, a wealth tax, an inheritance tax, and a user payment for nursing homes.

We estimate the pension income tax schedule on pension income using data on taxes and pension income for individuals for whom pension income constitutes the bulk of their gross income (see Appendix B.2 for details). The estimated pension income tax schedule involves a lump-sum tax of \$1,400 and a constant marginal tax rate of 36 percent for pension income above \$25,000.

The capital income tax in our sample period is 28 percent on all capital returns in excess of the risk-free rate (equal to 1.72 percent).

The wealth tax is 1.1 percent for taxable wealth above \$125,000. Various asset classes have different discounts when computing taxable wealth in Norway, the most important being that housing is valued at 25 percent of its estimated market value. In the model, we approximate the wealth tax by estimating the wealth tax as a linear function of net wealth in the data (see Appendix B.2 for details). This results in the following wealth tax formula

$$\text{wealth tax} = \tau_w \max\{\text{average discount} \cdot \text{wealth} - \bar{\tau}, 0\},$$

with τ_w equal to 0.011, $\bar{\tau}$ equal to \$159,000, and the average discount equal to 0.35.

The inheritance tax is 6 percent for inheritance above \$81,500 and 10 percent above \$138,500. We do not model inter-vivos transfers, but, as explained in Section 2, we include inter-vivos transfers in net wealth to include it as part of the relevant end-of-life bequest motive. Hence, the inheritance tax is only applicable at the end of life and relevant for the bequest motive.

The user payment schedule for nursing homes is progressive and capped. The costs are 75 percent of net income above a minimum threshold of \$1,350 and 85 percent above \$14,600. The maximum annual cost is capped at \$137,800.

4.2 Internal estimates

We estimate the remaining five parameters: $(\theta_0, \theta_1, \kappa_0, \kappa_1, \delta)$. These are the bequest parameters (θ, κ) for households with and without offspring and the utility shifter (δ) for households in nursing home care. We use the method of simulated moments to estimate these, meaning that we find the parameter values that minimize a distance criterion between model-simulated life-cycle profiles and their data counterparts. We apply a global minimization algorithm that selects parameter values to minimize these deviations.¹³

¹³Our preferred method for robustness is to quasi-randomly sample the parameter space, then select a set of best-fitting points upon which to apply a root-finding algorithm to each and select the best-fit solution

The moments we target in our estimation are the age profiles, between age 70 and 90, of the logarithm of: (i) unconditional average wealth holdings, (ii) average wealth holdings for those between selected percentiles (0-5, 5-10, 10-20, 20-40, 40-60, 60-80, 80-90, 90-95, 95-99) of the wealth distribution at age 70, (iii) average wealth holdings for households with and without offspring, (iv) average wealth holdings of households with no offspring in top 80 percent and the bottom 20 percent of the wealth distribution, (v) average wealth holdings of couples and singles by gender, and (vi) average wealth holdings by nursing home status. In total, we are targeting 399 moments, 19 moments for each of the 21 years in the 70-90 age range. All parameters are estimated jointly.


5 Parameter estimates

This section presents the parameter estimates of the model and the model fit. We first show how the model fits the targeted moments. Next, we discuss the parameter estimates and their implications. Third, we present the model fit for some untargeted moments.


5.1 Model fit

Before discussing the parameter estimates, Figure 4 shows the model's fit with the targeted moments. Overall, the model fits well the age profiles of average wealth holdings both in the aggregate and within the selected groups. The model overestimates somewhat the growth rate for the wealthiest 5 percent, slightly underestimates the growth rates for the 20-40 percent wealth group and overestimates the initial growth rate for the bottom 5 percent. The model underestimates growth rate after age 80 for the poorest households without offspring, but matches very well all other moments. Overall, we argue that the model does a remarkable job, considering the small number of parameters (5) relative to the targeted moments (399).


resulting from the multiple local minimizers.




(a) Average log wealth by age



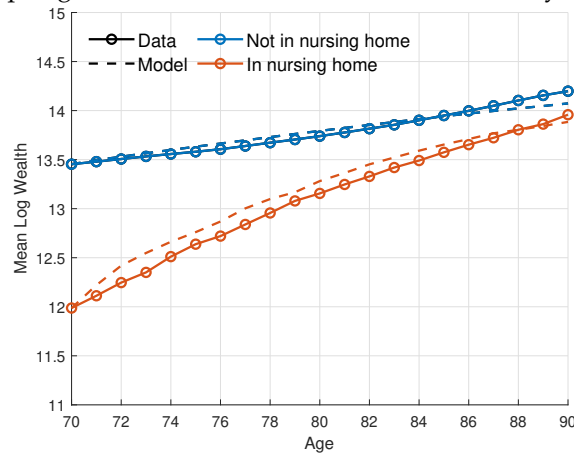
(b) By wealth at 70



(c) By offspring



(d) By marital status



(e) By health

Figure 4: Average log wealth by age. Data vs. model. Benchmark calibration.

Parameter	Description	Estimated value	Standard errors
θ_0	Bequest loading – no offspring	1,432,558	1,128,830
κ_0	Stone-Geary shifter – no offspring	471	191
θ_1	Bequest loading – with offspring	1,188,375	397,148
κ_1	Stone-Geary shifter – with offspring	392	707
δ	Felicity scaling – low wealth	0.90	0.18

Table 2: Parameter estimates.

5.2 Parameter estimates

Table 2 reports the parameter estimates in the baseline model. The values of (θ_i, κ_i) for $i = 0, 1$ capture the strength of the bequest motive. The first thing to notice is that the estimates imply a strong bequest motive not only for households with offspring but also for those without. This suggests that the bequest motive is capturing, in a reduced-form fashion, a reason to save which is not necessarily associated with a desire to bequeath.¹⁴

Our finding that the strength of the bequest motive is similar for households with and without children is in line with results in Hurd et al. (1987), Hurd (1989), and Kopczuk and Lupton (2007). For example, Hurd (1989) identifies the constant marginal utility of bequests by assuming that only households with children have a bequest motive. He estimates the (differential) motive to be economically negligible. Similarly, Kopczuk and Lupton (2007) also assume constant marginal utility of bequests but allow for unobserved heterogeneity in the presence of a bequest motive. They find it to be sizeable¹⁵ and its incidence to be higher among households with children. However, the difference in incidence is only marginally statistically significant. Compared to those two papers, we do not impose either (a) the absence of a bequest motive for households without offspring; or (b) we do not impose constant marginal utility of bequests.¹⁶

Figure 5 illustrates the strength of the bequest motive for both types of households, as well as those implied by the estimates in Kopczuk and Lupton (2007), De Nardi et al. (2010), Lockwood (2018), and Kvaerner (2023).¹⁷ It does so by plotting the share of cash at hand allocated to bequests that solves the one-period problem of a single who dies with certainty at the end of the period.

¹⁴While it is reasonable that individuals may derive utility from bequeathing to relatives other than their own children, it seems unlikely that this motive is similar to that towards one’s offspring.

¹⁵They find half of the bequeathed wealth is accounted for by an operative bequest motive.

¹⁶Constant marginal utility of bequests implies that consumption is independent of income and wealth for agents who leave positive bequests at the terminal date.

¹⁷As most of the literature, De Nardi et al. (2010) and Lockwood (2018) estimate the bequest motive on the whole population without distinguishing between households with and without children.




Figure 5: Estimated bequest motive. Comparing our result for households with and without offspring with [Kopczuk and Lupton \(2007\)](#), [De Nardi et al. \(2010\)](#), [Lockwood \(2018\)](#) and [Kvaerner \(2023\)](#).

$$\max_b \frac{(z - a')^{1-\sigma}}{1-\sigma} + \theta_j \frac{(\kappa_j + b(a'))^{1-\sigma}}{1-\sigma} \quad \text{s.t. } a' \geq 0. \quad (4)$$

The figure¹⁸ makes clear that households' propensity to leave bequests is basically the same, independently of whether the household has offspring. The estimates by [Lockwood \(2018\)](#) for the U.S. imply a bequest motive whose strength is very similar to ours. [Kvaerner \(2023\)](#) uses similar data for Norway focusing on singles with offspring. On the basis of the response of inter-vivos transfers from singles to offspring to negative life-expectancy news, [Kvaerner \(2023\)](#) estimates bequest parameters that imply a motive close, but slightly below, ours. On the other hand, [De Nardi et al. \(2010\)](#), and even more [Kopczuk and Lupton \(2007\)](#), estimate a substantially weaker bequest motive and with a substantially larger luxury-good feature.

The utility loading δ for individuals in a nursing home is estimated to be 0.90, which implies a lower utility from consumption for households in long-term care. This is identified by the differential wealth growth rates of households in and outside nursing homes (see [Figure 4e](#)), accounting for the selection into nursing home accession. In line

¹⁸The figure assumes $b(a') = a'$ for comparability with the cited papers.

with our finding, [Finkelstein et al. \(2013\)](#) argue that people in bad health have a lower marginal utility of consumption on the basis of survey responses to questions about happiness. [Blundell et al. \(2020\)](#) find that consumption responds negatively to temporary health shocks, with most of the effect being accounted for by reductions in the marginal utility of consumption rather than in available resources. [Christensen et al. \(2022\)](#) is the closest paper to ours. They estimate a life cycle model with health-dependent utility on retired, non-home-owning singles using Danish registry data. Their estimates of the ratio of the marginal utility of a given consumption flow in the bad-health relative to the good-health state, the counterpart of our δ , range between 0.18 and 0.25.

[Ameriks et al. \(2020\)](#) also build a life cycle model with state-dependent utility and estimate preference parameters using purpose-designed survey data. Their estimates imply a substantially *larger* relative marginal utility of expenditure in bad health states. One possible way to reconcile the two findings is that the survey questions that identify the utility shifter in [Ameriks et al. \(2020\)](#) ask about how individuals would allocate wealth: (a) between the state of the world in which they are healthy and that in which they need long-term care; (b) between spending on self or bequest when needing long-term care in the last year of their life. Both questions identify the relative utility of expenditure on consumption *plus* long-term care, which may be realistically driven by a large relative utility of spending on long-term care when in need of it rather than a higher relative utility of consumption net of long-term care costs. On the other hand, the risk of large, out-of-pocket medical and care expenditures is negligible both in Denmark and in Norway, which implies that the estimates in [Christensen et al. \(2022\)](#) and our paper capture the relative utility of consumption, net of long-term care costs.

5.3 Untargeted moments

In Section 5.1, we discussed how the model fits the targeted moments in the calibration. In this section, we show how the model fits one set of moments not explicitly targeted at the estimation stage: the mortality-adjusted wealth profile.

As discussed in more detail in Appendix A.2, there are several alternative ways of estimating wealth profiles. In the body of the paper, we present average wealth profiles by age, correcting for time and cohort effects. However, because wealthy individuals tend to live longer, these average wealth profiles grow more than what one individual should expect, due to selection on mortality by wealth ([Attanasio and Hoynes, 2000](#)); namely, the sample of households at older ages disproportionately include households that were richer when young. The model deals with such selection by using estimated mortality rates that depend on state variables. This approach controls for the selection on wealth




Figure 6: Mortality-adjusted mean log wealth.

only to the extent that the estimated mortality rates precisely reflect the selection in the data.

To verify that this is the case, Figure 6 presents the mortality-adjusted wealth profiles in the data, following the approach in [Attanasio and Hoynes \(2000\)](#), together with the mortality-adjusted wealth profiles in the model, using the same approach. In the data, we adjust for mortality by estimating how the mortality rates depend on wealth and age. We then use the inverse of the implied cumulative survival probabilities as weights when we compute the wealth profiles. Similarly, in the model, we use the employed mortality rates to compute survival probabilities and the inverse of these survival probabilities as weights to compute wealth profiles. The two approaches differ because we do not allow the mortality rates in the model to depend on *current* wealth because wealth is an endogenous state variable, and it would affect behavior. Instead, mortality rates in the model depend on initial wealth at age 70, permanent income, health, age, and marital status. Figure 6 shows that the resulting mortality-adjusted wealth profiles in the model and data are similar. They first increase until the mid-80s before they start decreasing. We take this as reassuring evidence that the mortality rates we use in the model estimation reasonably approximate the mortality rates in the data.

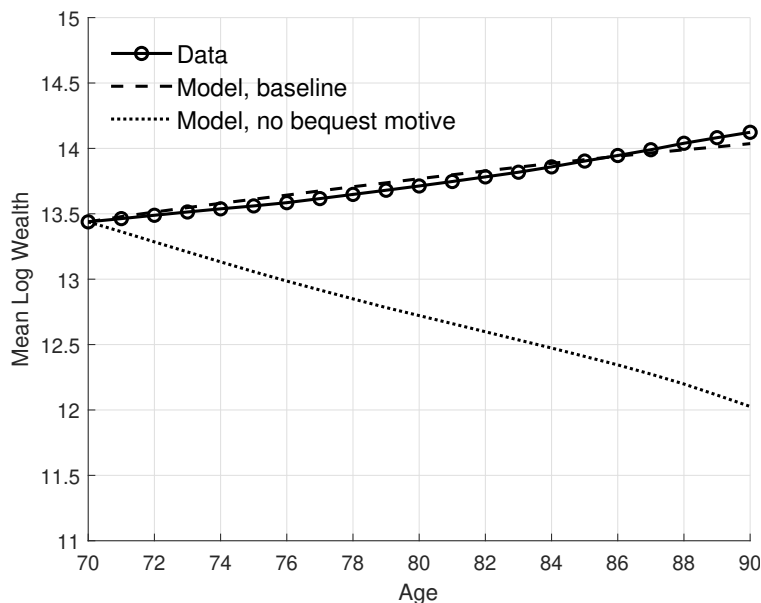


Figure 7: Average log wealth by age, baseline calibration vs. calibration with no bequest motive.

6 Model implications

In this section we use the model to understand the role of the bequest motive in accounting for the life-cycle wealth patterns of retirees and their responsiveness to tax policies.

6.1 Drivers of saving

To assess the importance of the bequest motive, we re-estimate the model by setting to zero the parameters of the utility from leaving bequests; $\theta_0 = \theta_1 = 0$. The resulting economy features only accidental bequests due to the lack of markets to unitize survival risk. Figure 7 compares the the age profiles of average wealth in the two calibrations (Figure C.1 in Appendix C reports all the moments in Figure 4) for the economy without bequest motive. In the aggregate, in the absence of a bequest motive, average wealth at age 85 is less than 25 percent (approximately \$300,000 against \$1,300,000) of its data counterpart. Generally, the model without a bequest motive implies strongly declining wealth-age profiles, inconsistent with the flat or upward-sloping profiles in the data.

In terms of the associated parameter estimates (reported in the second row of Table 2), the restricted model has only one free parameter: the utility shifter δ for households in nursing homes. The estimated value is at the upper bound (10) of the respective range. Intuitively, the MSM algorithm aims to fit the age profile of wealth holdings by couples by increasing marginal utility when in a nursing home by adjusting δ to induce self-

sufficient individuals to postpone consumption to when they will be in a nursing home (see Figure C.1e).

6.2 Savings response to tax changes

In this section we use the model to study the response of wealth to changes in both wealth and estate taxes for the population of retirees in the benchmark economy. To be precise, we take the full sample of those aged 70 and above in the initial period and simulate this group forward using the parameters and transition matrices from the model, with the policy rules implied by the new tax policy. This is the same modelling strategy taken in [Jakobsen et al. \(2020\)](#) to match their difference-in-difference empirical analysis of the Danish wealth tax reform in 1989.

Change in the wealth-tax rate. The first tax change we consider is an increase of 1 percentage point in the marginal rate of tax on wealth above the exemption threshold.¹⁹ Note that approximately 70% of the simulated population in the 2013 baseline year have wealth above the threshold.

Panel (a) in Figure 8 plots the percentage change in wealth for the whole simulated population. Wealth falls by 0.64% (\$6,354 per household) in the first year and is 6.26% lower by the tenth year. This corresponds to an average elasticity with respect to the net-of-tax rate on wealth of approximately 0.64 and 6.26 respectively.²⁰ A large proportion of the effect of the tax increase is ‘mechanical,’ namely due to the change in the after-tax rate of return on wealth accumulation at unchanged saving function. By contrast, the ‘behavioral’ effect is the change in wealth growth due to the change in the saving function induced by the lower after-tax rate of return. In the first year, the behavioral effect is 16% of the total wealth response, but by the tenth year, this proportion has dropped to 10%.

The change in aggregate wealth discussed above is the relevant quantity from the perspective of assessing the implications of changes in wealth taxation for the supply of capital in the economy. To compare the tax elasticity in the model to the available empirical estimates, though, we need to look at changes in total *taxable* wealth, namely wealth above the exemption threshold. This is because most difference-in-difference estimates exploit differences in post-reform saving behavior between (treated) agents whose wealth is above the exemption threshold (thus subject to tax) after the reform and a (control) group

¹⁹Since 35% of a household’s wealth is taxed, this requires a raise in the headline tax rate from 1.10% to 3.86%.

²⁰Given a pre-reform marginal tax rate of 1 per cent, the elasticity is basically equal to the percentage change in wealth.

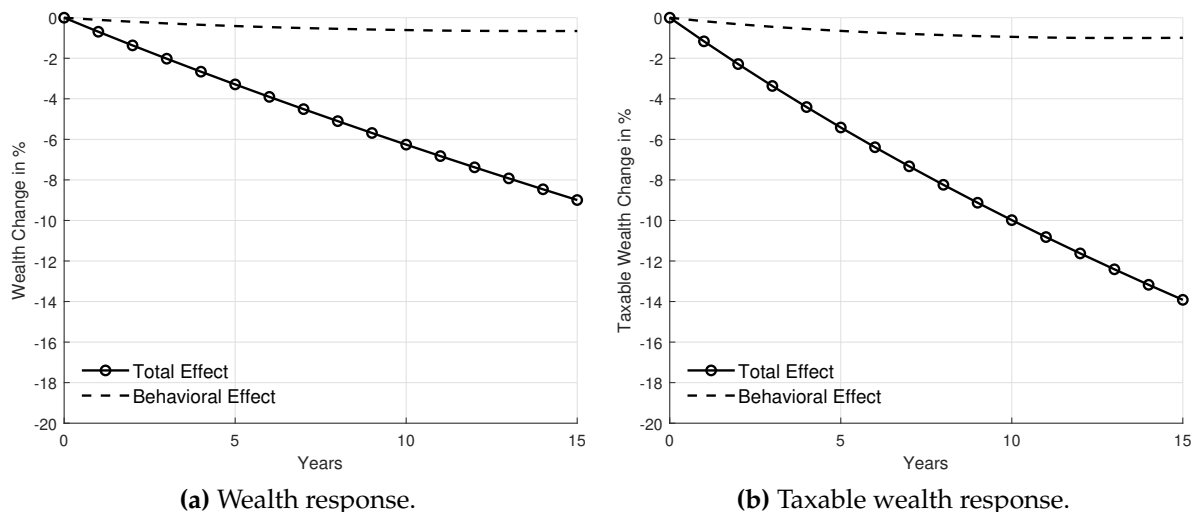


Figure 8: The responses of aggregate wealth and wealth above the wealth-tax threshold to a one percentage point wealth tax increase.

of agents with wealth below the threshold. For this reason, panel (b) in Figure 8 reports the change in total taxable wealth, namely wealth above the exemption threshold, over time. This corresponds to the treatment-on-the-treated effect in difference-in-difference studies. The average reduction in wealth above the threshold is -1.17% (\$6,480 per affected household) in the first year, 15% of which is due to the behavioral effect. By year ten, taxable wealth has fallen by 9.98%, of which 9.4% is the behavioral effect.

The relevant empirical benchmark for the wealth tax elasticities in the model are empirical estimates that (a) are not local to kinks in the tax schedule, unlike those relying on bunching techniques; and (b) attempt to isolate the pure saving response from changes in wealth due to tax avoidance.²¹ Two studies that best meet these two requirements are [Jakobsen et al. \(2020\)](#), for Denmark, and [Brülhart et al. \(2022\)](#), for Switzerland.²² Both follow a difference-in-difference approach. [Jakobsen et al. \(2020\)](#), whose setup is closest to our experiment and whose approach we have tried to reproduce here, estimate an (eight-year) average elasticity of taxable wealth to the net-of-tax rate of 5.1 and 8.9 for moderately wealthy couples.²³ The two values are obtained by comparing the response of (treated) couples benefiting from a doubling of the couple exemption threshold to, respectively,

²¹See [Saez and Zucman \(2019\)](#) and [Advani and Tarrant \(2021\)](#) for comprehensive surveys of the, burgeoning, empirical literature on the topic and the issues involved.

²²Similarly to Norway, wealth taxation in Denmark is broad-based and relies on third-party reporting. For this reason, estimated elasticities are less likely to be biased upward by strategic asset reallocation and avoidance. This is not the case in Switzerland, but [Brülhart et al. \(2022\)](#) make an effort to identify the pure saving response net of the response due to multiple channels of tax avoidance.

²³Formally, couples with wealth between percentiles 97.6 and 99.3 of the household wealth distribution. This is the most comparable group to our sample which excludes the top 1% of the wealth distribution.




Figure 9: Wealth response by age to a one percentage point increase in the tax rate.

that of two alternative control groups: (a) couples below the post-reform threshold and (b) singles with the same average, pre-reform household wealth as the treated couples.²⁴ The average elasticity over the first eight years in our model is 4.8. The corresponding estimate by Brülhart et al. (2022) is 6.9.

So the wealth tax elasticity in the model is in the ballpark, if at the low (in absolute value) end, of comparable empirical estimates. This is reassuring given that it is an untargeted moment in the model estimation. Furthermore, part of the difference can be accounted for by the higher average age in our sample. We estimate our model on individuals aged more than 70 compared to more than 60 in Jakobsen et al. (2020) and the whole population in Brülhart et al. (2022). Empirical estimates of wealth elasticities are larger for younger individuals (see the Online Appendix of Jakobsen et al., 2020; Brülhart et al., 2022). Figure 9 shows that this is true also in the model, as one would expect given that the tax change affects younger households over a longer horizon.

Removal of the estate tax. The fact that the model economy displays a wealth-tax elasticity broadly in line with empirical estimates gives us confidence in using it to evaluate the implications of estate taxation for wealth accumulation. Until 2014, inherited wealth in Norway was subject to progressive taxation as in our baseline economy described in Section 4.1. The estate tax was eliminated in 2014. Since our baseline model is estimated

²⁴Therefore, the average per-capita wealth of singles in the second control group is double that of the treated couples.




Figure 10: Aggregate wealth and bequest responses to abolishing the bequest tax.

on the pre-2014 period, it seems quite natural to use it to evaluate the implications of the 2014 reform.

Figure 10a plots the response of aggregate wealth among those alive. Aggregate wealth *falls* by a negligible 0.27 percent (\$2,562 per household) after 15 years, in response to the elimination of the estate tax. This corresponds to an aggregate wealth elasticity with respect to the net-of-estate-tax rate of about -0.03, negative but very close to zero.

To understand the sign and magnitude of the response of wealth to estate taxation, it is helpful to inspect the (interior) optimum of two-period problem in equation (4) under the simplifying assumption that the estate tax is proportional; i.e. $b(a') = (1 - \tau)a'$. The optimal pre-tax bequest satisfies

$$a' = \frac{z - \theta_j^{-\frac{1}{\sigma}}(1 - \tau)^{-\frac{1}{\sigma}}\kappa_j}{1 + \theta_j^{-\frac{1}{\sigma}}(1 - \tau)^{\frac{\sigma-1}{\sigma}}}. \quad (5)$$

If preferences are homothetic— $\kappa_j = 0$ —the pre-tax bequest a' is decreasing in the net-of-tax rate $1 - \tau$ if and only if $\sigma > 1$ —the positive income effect on consumption more than offsets the negative substitution effect.²⁵ This is indeed the case in our parameterization, which features $\sigma = 2$. When $\kappa_j > 0$, preferences are non-homothetic and bequests are a luxury good. An increase in $1 - \tau$ reduces the numerator on the right hand side of equation (5), reinforcing the substitution effect. Our parameter estimates, though, imply that the income effect more than offsets the other two and pre-tax bequests respond negatively to

²⁵Similarly, Ring (2023) shows that the elasticity of the intertemporal substitution drives the sign of the response of wealth to the wealth tax rate in an economy without a bequest motive.

an increase in the net-of-tax rate.²⁶

Equation (5) also implies that the magnitude of the response of pre-tax bequests to $1 - \tau$ is decreasing in θ_j , the bequest loading. Given our estimates for θ_j , the coefficients of $(1 - \tau)$ in the above expression are very close to zero and so is the response of a' to estate taxes. Effectively, the economy behaves as if the average retired household is satiated in consumption. Figure 10b illustrates this by plotting the change in pre- and post-tax bequeathed wealth. Total *pre-tax, bequeathed* wealth falls by \$1,870 per household (or 0.18% of aggregate bequests) for the whole cohort, but the *post-tax, bequeathed* amount would rise by \$89,535 per household (9.69% of aggregate bequests). As a result, government revenue falls by \$91,405 per household relative to the baseline.

There is very little empirical work on the response of wealth to estate taxation. For the U.S., [Kopczuk and Slemrod \(2001\)](#) estimate an elasticity of *reported* bequests to (one minus) the tax rate of 0.16, while [Joulfaian \(2006\)](#) a value of 0.09. [Goupille-Lebret and Infante \(2018\)](#) find elasticities between 0.25 and 0.35. These numbers are not trivial, given estate tax rates of 0.4-0.5 in France and the U.S. The elasticity implied by our experiment is negative, reflecting the optimal response, but an order of magnitude smaller in absolute value. It has to be kept in mind, though, while empirical estimates may reflect changes in estate planning or reporting, our model captures the pure behavioral, saving response.

7 Conclusion

In this paper, we use rich, high-quality administrative data to study late-in-life saving, inter-vivos transfers and the role of a bequest motive in that saving. We document the presence of widespread, large and positive saving post-retirement in an institutional environment with little to no medical and nursing-home expenditure risk. Importantly, this is the case across households that differ by health, couple status, existence of offspring, age, sex, wealth, and income. The poorest households, and childless households, increase wealth the fastest. These patterns are unlike many documented, e.g. for the U.S.A., which often exclude inter-vivos transfers.

To account for these facts, we build and estimate a heterogeneous agent model with bequest motives, realistic income, morbidity and mortality processes, and idiosyncratic returns. We find that the model matches the data well and that strong and non-homothetic bequest motives are important for explaining saving patterns in the population. The

²⁶Note that, if anything, our chosen value for the CRRA coefficient is conservative in the light of the meta study by [Havránek \(2015\)](#) and the findings in [Best, Cloyne, Ilzetzki and Kleven \(2019\)](#). Also, although the *sign* of the wealth response does depend on σ , quantitatively the size of its absolute value is small if θ_j is large.

estimated model indicates that around three-quarters of bequeathed wealth is attributable to the bequest motive. It is particularly striking that childless households have a bequest motive quantitatively comparable to households with offspring. This suggests that the utility from residual wealth in the model is also capturing other positive saving forces beyond altruism towards descendants.

We use the model to assess the wealth response of the elderly to wealth and estate taxes. Reassuringly, the model implies an elasticity of aggregate saving with respect to the after-wealth-tax rate broadly in line with existing comparable empirical estimates (Brülhart et al., 2022; Jakobsen et al., 2020). The response is mostly due to the (mechanical) impact of the change in the after-tax return of wealth accumulation with little change in saving. For the same reason, the abolition of estate taxation would reduce tax revenue with hardly any effect on wealth accumulation.

References

- Advani, A. and Tarrant, H. (2021), 'Behavioural responses to a wealth tax', *Fiscal Studies* **42**(3-4), 509–537.
- Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D. and Tonetti, C. (2020), 'Long-term care utility and late-in-life saving', *Journal of Political Economy* **128**(6), 2375–2451.
- 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', *Journal of Finance* **66**(2), 519–561.
- Attanasio, O. P. and Hoynes, H. W. (2000), 'Differential mortality and wealth accumulation', *Journal of Human Resources* **35**(1), 1–29.
- Best, M. C., Cloyne, J. S., Ilzetzki, E. and Kleven, H. J. (2019), 'Estimating the elasticity of intertemporal substitution using mortgage notches', *The Review of Economic Studies* **87**(2), 656–690.
- Blundell, R., Borella, M., Commault, J. and Nardi, M. D. (2020), Why does consumption fluctuate in old age and how should the government insure it? NBER Working Paper 27348.
- Braun, R. A., Kopecky, K. A. and Koreshkova, T. (2016), 'Old, sick, alone, and poor: A welfare analysis of old-age social insurance programmes', *The Review of Economic Studies* **84**(2), 580–612.
- Brühlhart, M., Gruber, J., Krapf, M. and Schmidheiny, K. (2022), 'Behavioral responses to wealth taxes: evidence from Switzerland', *American Economic Journal: Economic Policy* **14**(4), 111–150.
- Christensen, B. J., Kallestrup-Lamb, M. and Kennan, J. (2022), Consumption and saving after retirement. NBER Working Paper 29826.
- 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. (2016a), 'Medicaid insurance in old age', *American Economic Review* **106**(11), 3480–3520.
- De Nardi, M., French, E. and Jones, J. B. (2016b), 'Saving after retirement: a survey', *Annual Review of Economics* **8**, 177–204.
- De Nardi, M., French, E., Jones, J. B. and McGee, R. (2023), 'Why do couples and singles save during retirement? household heterogeneity and its aggregate implications', *NBER Working Paper No. 28828* .
- Deaton, A. and Paxson, C. (1994), 'Intertemporal choice and inequality', *Journal of Political Economy* **102**(3), 437–467.

- Dynan, K. E., Skinner, J. and Zeldes, S. P. (2002), 'The importance of bequests and life-cycle saving in capital accumulation: a new answer', *American Economic Review* **92**(2), 274–278.
- Fagereng, A., Guiso, L., Malacrino, D. and Pistaferri, L. (2020), 'Heterogeneity and persistence in returns to wealth', *Econometrica* **88**(1), 115–170.
- Fagereng, A., Holm, M. B., Moll, B. and Natvik, G. (2021), Saving behavior across the wealth distribution: The importance of capital gains, NBER Working Papers 26588.
- Fagereng, A., Holm, M. B. and Torstensen, K. N. (2020), 'Housing wealth in Norway, 1993–2015', *Journal of Economic and Social Measurement* **45**(1), 65–81.
- Fernandez-Villaverde, J. and Krueger, D. (2007), 'Consumption over the life cycle: Facts from consumer expenditure survey data', *Review of Economics and Statistics* **89**, 552–565.
- Finkelstein, A., Luttmer, E. F. P. 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**, 221–258.
- Foltyn, R. and Olsson, J. (2023), Subjective life expectancies, time preference heterogeneity, and wealth inequality. Mimeo.
- Førland, O., Ambugo, E. A., Døhl, Ø., Folkestad, B., Rostad, H. M. and Sundsbø, A. O. (2020), 'Variasjon i kvalitet i omsorgstjenestene', *Senter for omsorgsforskning rapportserie* .
- French, E., De Nardi, M., Jones, J. B., Baker, O. and Doctor, P. (2006), 'Right before the end: Asset decumulation at the end of life', *Economic Perspectives* **30**(3), 2–13.
- French, E. and Jones, J. B. (2011), 'The effects of health insurance and self-insurance on retirement behavior', *Econometrica* **79**, 693–732.
- French, E., Jones, J. B. and McGee, R. (2023), 'Why do retired households draw down their wealth so slowly?', *Journal of Economic Perspectives* **37**(4), 91–114.
- Goupille-Lebret, J. and Infante, J. (2018), 'Behavioral responses to inheritance tax: evidence from notches in France', *Journal of Public Economics* **168**, 21–34.
- Halvorsen, E., Hubmer, J., Ozkan, S. and Salgado, S. (2021), Why are the wealthiest so wealthy? an empirical-quantitative investigation of lifecycle wealth dynamics. Mimeo.
- Havránek, T. (2015), 'Measuring intertemporal substitution: The importance of method choices and selective reporting', *Journal of the European Economic Association* **13**(6), 1180–1204.
- Hurd, M. D. (1989), 'Mortality risk and bequests', *Econometrica* **57**(4), 779–813.
- Hurd, M. D. et al. (1987), 'Savings of the elderly and desired bequests', *American Economic Review* **77**(3), 298–312.

- Jakobsen, K., Jakobsen, K., Kleven, H. and Zucman, G. (2020), 'Wealth taxation and wealth accumulation: Theory and evidence from Denmark', *Quarterly Journal of Economics* **135**(1), 329–388.
- Joulfaian, D. (2006), 'The behavioral response of wealth accumulation to estate taxation: Time series evidence', *National Tax Journal* **59**(2), 253–268.
- Kaplan, G. and Violante, G. L. (2010), 'How much consumption insurance beyond self-insurance?', *American Economic Journal: Macroeconomics* **2**, 53–87.
- Kopczuk, W. and Lupton, J. P. (2007), 'To leave or not to leave: the distribution of bequest motives', *Review of Economic Studies* **74**(1), 207–235.
- Kopczuk, W. and Slemrod, J. (2001), The impact of the estate tax on wealth accumulation and avoidance behavior, in W. Gale, J. Hines and J. Slemrod, eds, 'Rethinking Estate and Gift Taxation', North-Holland.
- Kopecky, K. and Koreshkova, T. (2014), 'The impact of medical and nursing home expenses and social insurance policies on savings and inequality', *American Economic Journal: Macroeconomics* **6**, 29–72.
- Kvaerner, J. S. (2023), 'How large are bequest motives? Estimates based on health shocks', *The Review of Financial Studies* **36**(8), 3382–3422.
- Lockwood, L. (2018), 'Incidental bequests: bequest motives and the choice to self-insure late-life risks', *American Economic Review* **108**(9), 2513–2550.
- McGee, R. (2021), 'Old age savings and house price shocks', *The University of Western Ontario Research Report No. 2021-4* .
- Nakajima, M. and Telyukova, I. A. (2015), Housing in retirement across countries, in J. E. Stiglitz and M. Guzman, eds, 'Contemporary Issues in Microeconomics', Palgrave MacMillan.
- Nakajima, M. and Telyukova, I. A. (2017), 'Reverse mortgage loans: a quantitative analysis', *Journal of Finance* **72**(2), 911–950.
- Nakajima, M. and Telyukova, I. A. (2020), 'Home equity in retirement', *International Economic Review* **61**(2), 573–616.
- Nakajima, M. and Telyukova, I. A. (2024), 'Medical expenses and saving in retirement: The case of u.s. and sweden', *American Economic Journal: Macroeconomics* . Forthcoming.
- Palumbo, M. G. (1999), 'Uncertain medical expenses and precautionary saving near the end of the life cycle', *Review of Economic Studies* **66**, 395–421.
- Pashchenko, S. and Porapakarm, P. (2023), 'Saving motives over the life-cycle', *Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3585755>* .

- Poterba, J., Venti, S. and Wise, D. (2011), 'The composition and drawdown of wealth in retirement', *Journal of Economic Perspectives* **25**(4), 95–118.
- Ring, M. A. K. (2023), Wealth taxation and household saving: Evidence from assessment discontinuities in Norway. mimeo.
- Saez, E. and Zucman, G. (2019), 'Progressive wealth taxation', *Brookings Papers on Economic Activity* pp. 437–533.
- Smith, J. P. (1999), 'Healthy bodies and thick wallets: the dual relation between health and economic status', *Journal of Economic perspectives* **13**(2), 145–166.
- Van Ooijen, R., Alessie, R. and Kalwij, A. (2015), 'Saving behavior and portfolio choice after retirement', *De Economist* **163**(3), 353–404.

Online Appendix

A Appendix to Section 2

A.1 Institutional setting


Pensions. The pension system in Norway builds on three pillars: the social security system, occupational pensions, and private pension savings. The social security system provides defined benefit pensions to all Norwegian citizens. While working, individuals accumulate pension rights above a minimum as a function of gross income. These pension rights are converted to pension claims after retirement. Annual pension payments are indexed to average wage growth and adjusted for average life expectancy at retirement. After retirement, there is no uncertainty about future pension payments except for annual wage growth indexation. This social security pension system is the main pension income of most households in our sample of data.

Occupational pensions vary depending on employers. In the public sector, individuals accumulate a state occupational pension. In our sample, a former public sector employee with full pension contributions (at least 30 years of public employment) receives 66 percent of their last salary. Private occupational pensions vary, depending on the employer and type of contract.


The inheritance law. During our sample period, the inheritance law states that when an individual in a couple (either married or registered partner with joint children) dies, the nearest relatives inherit their wealth. The nearest relatives are defined in the following order: (i) offspring, (ii) parents, (iii) children of parent (sibling), (iv) grandparents, and (v) grandchildren of grandparents (cousins). Unless there is a will, the wealth is distributed equally among individuals in the first group that contains individuals. For example, if the deceased individual has children, the wealth is distributed equally among the children. And if the deceased has no (living) offspring or parent, the wealth is distributed equally among the siblings. In the sample of deceased individuals with no offspring in our data, 7.3 percent of inheritances go to siblings, while 4.3 percent of inheritances go to cousins. The rest goes to individuals or entities outside the family.

One complication is that a surviving spouse of the deceased is entitled to a minimum of the wealth.²⁷ This minimum threshold introduces a kinked scheme in which all wealth


²⁷This minimum was computed as the maximum of one-quarter of the wealth and four times the minimum in the social security system (approximately \$40,000). This minimum was approximately \$60,000 when the




(a) Mean log wealth




(b) Mean log wealth ex. inter-vivos transfers




(c) Mean log wealth, mortality adjusted



(d) Mean log wealth ex. inter-vivos transfers, mortality adjusted



(e) Median log wealth



(f) Median log wealth ex. inter-vivos transfers

Figure A.1: Log wealth by age. Raw data and estimated age profiles.

closest relative was not an offspring.


goes to the spouse if wealth is sufficiently low while it goes to other relatives if wealth is large. Another complication is that the law allows the surviving spouse to keep the estate undivided, that is, postponing dealing with the inheritance of the deceased until the surviving spouse's death. However, if the deceased has offspring that are not the surviving spouse's offspring, the children of the deceased have to consent that the estate stays undivided.

A.2 Adjusted wealth profiles


A major challenge in understanding the wealth accumulation of the elderly is constructing relevant wealth profiles by age. There are two issues. First, we are interested in age profiles, which are contaminated by cohort and time effects in cross-sectional data. For example, when wealth grows over time, the cross-sectional age profile of wealth may be downward sloping partly because older households belong to older cohorts with lower wealth. Second, within a cohort, wealthier households tend to live longer (Attanasio and Hoynes, 2000). This composition effect biases the cross-sectional wealth profiles even within a cohort.

Figure A.1 illustrates how some central assumptions matter for estimating wealth profiles by age. First, Figure A.1a is the same as Figure 2a in the body of the paper and illustrates how our benchmark time-age-cohort decomposition works. Second, Figure A.1b displays the estimated wealth profiles when we exclude inter-vivos transfers from the wealth definition. The main difference is that the growth rate of wealth is lower compared with the benchmark case, where we include inter-vivos transfers. Wealth still, however, monotonically increases with age. Third, Figure A.1c shows how a mortality adjustment affects the wealth profiles. We follow Attanasio and Hoynes (2000) and weigh the estimated wealth profile by the inverse of the cumulative survival probability. Because the survival probability is positively correlated with wealth, the mortality adjustment tends to reduce the growth rate of wealth. As a result, wealth is first increasing with age, then decreasing. Fourth, Figure A.1d shows the estimated wealth profile when we both exclude inter-vivos transfers and mortality adjust. Wealth now decreases for age for all age groups.


When comparing our estimated wealth profiles with the model, they correspond to different objects. Because the model includes mortality rates that are correlated with permanent income, and thus wealth, the population-averaged wealth profiles from the model should be compared with Figure A.1a as in the body of the paper. An alternative approach, however, is to use Figure A.1c. In a model, Figure A.1c corresponds to the expected wealth profile of an individual, more closely mapping into the *savings policy*




(a) Mean wealth




(b) Mean wealth ex. inter-vivos transfers




(c) Mean wealth, mortality adjusted



(d) Mean wealth ex. inter-vivos transfers, mortality adjusted



(e) Median wealth



(f) Median wealth ex. inter-vivos transfers

Figure A.2: Wealth by age. Raw data and estimated age profile.

function. In Section 5.3, we therefore compare the saving policy functions in the model with the wealth profile in Figure A.1c as an untargeted moment.

An alternative is to estimate wealth profiles for medians rather than mean log wealth. Figure A.1e and A.1f display the estimated median wealth profiles and raw data from some selected cohorts. The median wealth profiles are very similar to the mean wealth profiles. We therefore use mean wealth profiles because they more naturally map into wealth profiles in simulated data.

Another alternative is to estimate wealth profiles in levels rather than in logs. Figure A.2 presents raw data and estimated age profiles of mean wealth. The estimated wealth profiles are still increasing in wealth. Estimated mean wealth increases from around \$800,000 to \$1,400,000, an increase of approximately 50 percent, similar to what we find when we estimate wealth profiles using log wealth. The patterns as we exclude inter-vivos transfers, adjust for mortality, or compute median wealth profiles are also very similar.

B Appendix to Section 4

B.1 Returns to Net Wealth

This appendix contains details on how we compute returns to net wealth. We define returns as the sum of capital income and capital gains as a share of net wealth, where net wealth excludes inter-vivos transfers relative to its definition in the paper. Capital income includes capital income and housing service flows. Capital gains covers both realized and unrealized capital gains. When we divide by wealth, we take the average of wealth at Dec 31 the previous and current year. We focus on individuals with positive net wealth. In addition, because the returns are volatile for some individuals, for example due to large changes in net wealth, we trim at the top and bottom 5 percent (-25 percent and + 37 percent) of the returns distribution.

Figure B.1 shows the cross-sectional distribution of returns to net wealth. In our sample, average returns is 0.068 with a standard deviation of 0.086. In the model, we assume that returns are i.i.d. normal with the same mean and standard deviation. Figure B.1 includes the normal distribution on top of the data distribution. While the data distribution is more kurtotic than the normal distribution, the normal distribution seems to be a reasonable approximation of the returns distribution for our purpose.

A further assumption we make is that expected returns are independent of wealth. Figure B.2 plots the mean and standard deviation of returns against percentiles of the net wealth distribution. In our sample of older households, mean returns is independent

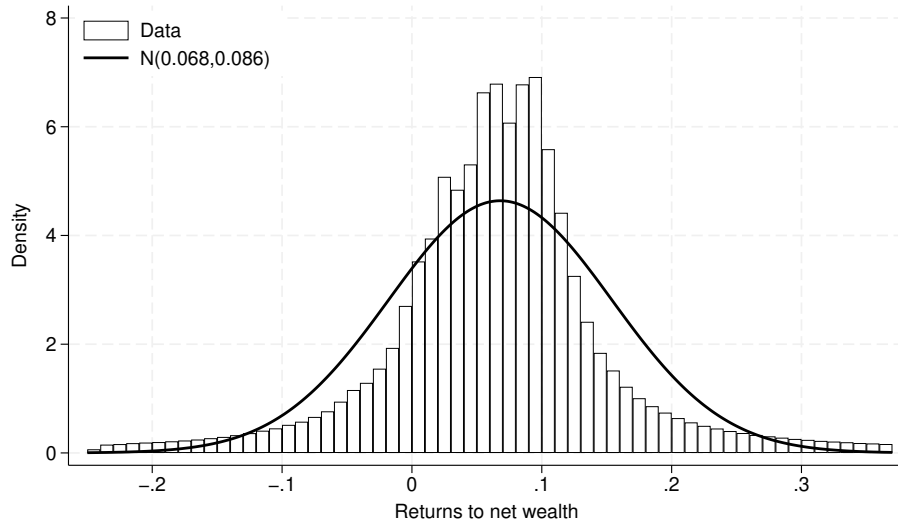


Figure B.1: The distribution of returns to net wealth.

of wealth above a threshold. Similarly, the standard deviation of returns can also be reasonably well approximated as independent of wealth.




Figure B.2: Mean and standard deviation of returns to net wealth by wealth percentile.

The third assumption is to say that returns are i.i.d., i.e., there is no persistence in returns among individuals. Table B.1 tests this assumption by running a regression of current returns on previous period returns. If returns are to a large extent individual-specific, current returns should be predicted by past returns when we include no controls. Instead, the coefficient on past returns is negative, suggesting mean-reversion of returns. This result also holds when we add individual- and time-fixed effects.

Dependent variable: returns to net wealth _t				
	(1)	(2)	(3)	(4)
Returns to net wealth _{t-1}	-0.043 (0.001)	-0.119 (0.001)	-0.178 (0.001)	-0.243 (0.001)
Time-fixed effects	No	Yes	No	Yes
Individual-fixed effects	No	No	Yes	Yes
Observations	3,818,237	3,818,237	3,818,237	3,818,237

Notes: The table shows the coefficients of a regression of returns on lagged returns. Standard errors are clustered at the individual level.

Table B.1: Returns to net wealth.

B.2 Approximated Income and Wealth Tax Schedules

Income tax. We set the income tax schedule in the model to match Norway by estimating an approximate tax schedule in the data.²⁸ For each individual, we observe both their pension income and their total taxes paid that year. These observable taxes include all income taxes, including taxes on capital income, capital gains realizations, and other similar posts. To estimate the tax schedule, we focus on a subset of the population that has less than \$1,000 in combined capital and earned income. Among these individuals, pension income is their primary and almost only source of income. Hence, for these individuals, the taxes we observe reflect almost exclusively income taxes on pension income.

Figure B.3 displays the taxes and pension income in data together with the approximated tax schedule. The tax schedule is kinked in which taxes when income is below \$25,000 is approximately constant at \$1,400. From \$25,000, the marginal tax rate is constant at 36 percent. Figure B.3 shows that our approximated tax schedule matches the data well in our sample of individuals where income primarily comes from pension income.

²⁸Alternatively, we could directly impose the tax laws in Norway. However, because the tax system includes a set of deductions for variables which we have no information, we instead estimate a simplified tax schedule.

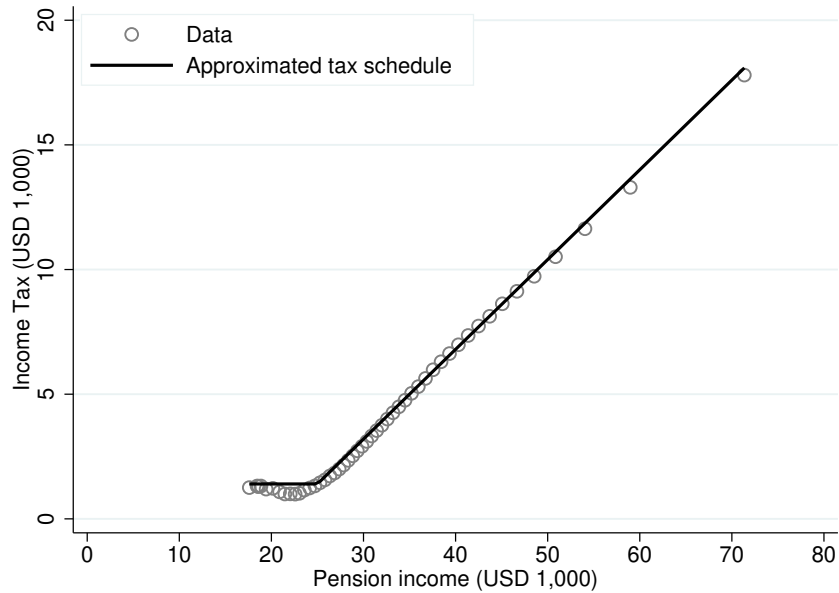


Figure B.3: Pension income and taxes.

Wealth tax. For the wealth tax, we also estimate an approximate tax schedule from the data. During our sample period, the wealth tax changed from year to year. We use the tax schedule from 2010, which is in the middle of our main sample years, to approximate a tax schedule. In 2010, the wealth tax was 1.1 percent on taxable wealth above a threshold of NOK 700,000 (approximately \$125,000) for singles and NOK 1,400,000 (approximately \$250,000) for couples. Taxable wealth is wealth after imposing a discount where housing wealth is valued at 25 percent of the market price. We approximate the wealth tax as a function of the market value of net wealth and housing, appropriately discounted. We first compute the wealth tax directly in the data from a measure of taxable wealth, after calculating an estimate for the empirically observed discount for total wealth, which is 65%. The functional form we assume is

$$\text{wealth tax} = \tau_w \max\{0.35 \cdot \text{net wealth} - \bar{\tau}, 0\},$$

where τ_w is an approximate marginal tax above the threshold and $\bar{\tau}$ is an approximate threshold. Figure B.4 displays the wealth and the wealth tax in data together with the approximated tax schedule. We compute $\bar{\tau}$ to be \$159,000 and τ_w to be 0.011 by minimizing the squared distance between the data and the approximated tax schedule. Figure B.4 shows that our approximated tax schedule matches the data reasonably well.

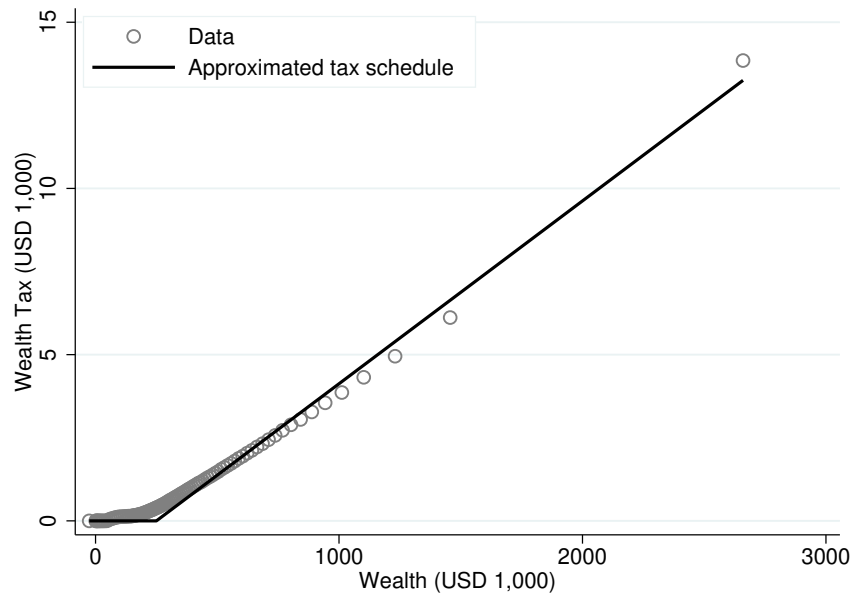




Figure B.4: Wealth and wealth taxes.

C Appendix to Section 6


C.1 Estimation results with no bequest motive




(a) Average log wealth by age




(b) By wealth at 70



(c) By offspring



(d) By marital status



(e) By health

Figure C.1: Average log wealth by age. Data vs. model. Calibration with no bequest motive.

C.2 Change in the wealth tax threshold




Figure C.2: Response of aggregate wealth to an increase of the wealth tax threshold from \$159,000 to \$318,000 by wealth relative to the thresholds at time 0.

An alternative wealth tax change is to modify the threshold at which the tax is paid. The estimated limit covers around two-thirds of the population in 2013. We consider doubling the limit from \$159,000 to \$318,000, leaving the tax to cover only 37% of the 2013 population. In Figure C.2 we group the households by their position relative to the wealth tax thresholds in 2013: those below the old (\$159,000) threshold, those above the old threshold but below the new threshold (\$159,000-318,000) and those above the new (\$318,000) threshold.²⁹ Those below the old threshold are not affected immediately or to a large degree, as they benefit from the threshold increase only if their wealth rises above the old threshold, and relatively few do. Their mean log wealth changes by only 0.003% after 20 years relative to the benchmark simulation with the original tax threshold. Those above the old threshold benefit from the tax relief on the previously taxed wealth. This forms an overall large, positive wealth response, most of which is mechanical, and for which they compensate with a negative behavioural response to the change. Those

²⁹Due to wealth growth as households age, and faster death of older, poorer cohorts, for clarity we group households by tax threshold status at time 0.

already above the new threshold gain a very small amount initially (since they receive the full benefit of the tax reprieve on \$318,000) but the response is very similar. This response, despite halving the taxable population, has a relatively small effect compared to the change in the marginal rate of tax. This is due to the latter's proportional effect across the wealth distribution compared to the small impact on the wealthy of a threshold change. Importantly, both tax changes demonstrate that the behavioural effect of the wealth tax changes is very small compared to the mechanical effect, i.e. given the bequest motive estimated from the data, wealth taxes have limited distortionary effects on saving.