University of Glasgow

Adam Smith Business School

WORKING PAPER SERIES

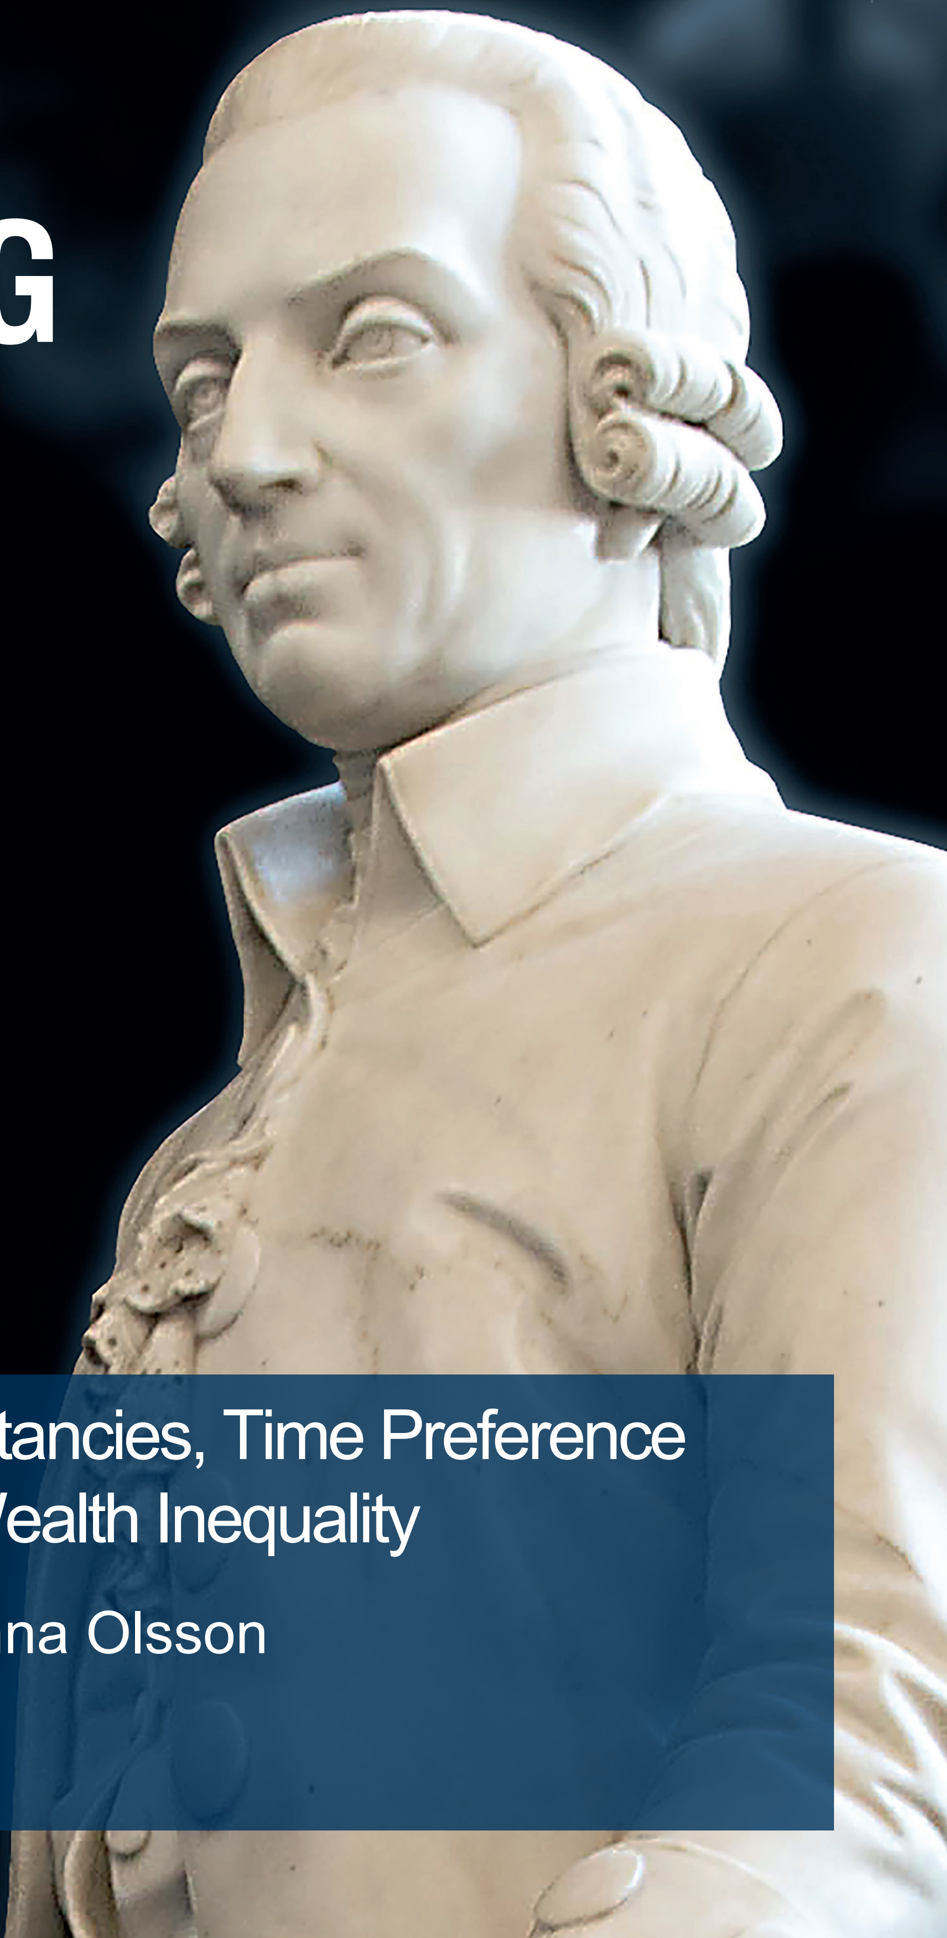

Subjective Life Expectancies, Time Preference Heterogeneity, and Wealth Inequality

Richard Foltyn and Jonna Olsson

Paper no. 2021-13
July 2021

Subjective Life Expectancies, Time Preference Heterogeneity, and Wealth Inequality*

Richard Foltyn† Jonna Olsson‡

July 2021

This paper explores how heterogeneity in life expectancy, objective (statistical) as well as subjective, affects savings behavior between healthy and unhealthy people. Using data from the Health and Retirement Study, we show that people in poor health not only have shorter actual lifespan, but are also more pessimistic about their remaining time of life. Using a standard overlapping-generations model, we show that differences in life expectancy can explain one third of the differences in accumulated wealth with an important part driven by pessimism among unhealthy people.

JEL Classification: D15, E21, G41, I14
Keywords: Life expectancy, preference heterogeneity, subjective beliefs, life cycle
Declaration of interest: none

1 Introduction

The determinants of the wealth distribution are of fundamental interest to economists. Standard consumption/savings theory predicts that people who place a larger weight on future states will be wealthier than people who are more impatient, all else equal. This paper explores one reason to put a higher weight on the future, namely higher survival probability, or, put differently, higher perceived probability to still be alive and thus experience the future states of the world.

*We are grateful to Timo Boppart, Per Krusell, Alexander Ludwig, Hannes Malmberg, Ayşegül Şahin, Paolo Sodini, Christian Stoltenberg, Magnus Åhl, Erik Öberg and participants in the IIES Macro Group, the Norface:TRISP 2019 conference, the Nordic Junior Macro Group, and the SED 2021 conference for helpful discussions and suggestions.

†richard.foltyn@glasgow.ac.uk, University of Glasgow.

‡jonna.olsson@ed.ac.uk, University of Edinburgh, School of Economics, 30-31 Buccleuch Place, Edinburgh, EH8 9JT, United Kingdom.

The difference in life expectancy in the population is large. Using data from the Health and Retirement Study, we show that an average 70-year-old American nonblack man in excellent health has a 75% chance of reaching his 80th birthday, while the corresponding probability for a nonblack man in poor health is below 40%. According to standard economic theory, the healthy man should save more for the future, given the higher probability of living a longer life. A strong positive correlation between life expectancy and wealth is well established empirically, but the causality and strength of different potential channels are less clear (see, e.g., Deaton 2002 and the references therein). Recent empirical work using plausibly exogenous variation in the timing of diagnosis of serious illness suggests a causal link from lower life expectancy to a decrease in savings and increase in consumption, potentially in the form of giving gifts (intervivo transfers) to children (Fernström 2019; Kvaerner 2020).

However, an individual's consumption/savings decision is not necessarily guided by the objective statistical life expectancy, but rather the individual's beliefs about survival. The first contribution of this paper is to document new facts about a *within-cohort steepness bias* in survival beliefs: people overestimate the health gradient of survival. It has previously been shown (e.g., Elder 2013, Ludwig and Zimper 2013, Heimer, Myrseth, and Schoenle 2019) that there is a systematic flatness bias over age: younger people tend to underestimate their survival probabilities, while older people are overly optimistic about their chances of a long life. We show that within a cohort, individuals in bad health not only have a shorter expected life span but are also relatively more pessimistic about their survival chances, while individuals in good health and thus with higher survival probability are more optimistic. This systematic bias exacerbates the survival expectancy heterogeneity in the population.

The differences in beliefs about survival translate into time preference heterogeneity in the population. Our second contribution is to quantify this heterogeneity and its implications for savings and wealth accumulation in an overlapping-generations model. With a stochastic health and survival process the effective discount rate (determined by the common discount factor β and the individual survival probability) varies depending on age and health and depends on the horizon. Over a one-year horizon, the effective discount rate for 50-year-olds varies between 2.0% (for an individual in best health) to more than 20% (for an individual in worst health). At longer horizons, the geometric averages of the expected discount rate becomes more similar for the two agents, but the difference is still 7.4 percentage points on a 10-year horizon. For 70-year-olds in worst vs. best health state, the difference at the 10-year horizon is more than 10 percentage points. This resulting time preference heterogeneity is in line with the magnitude of dispersion (Calvet et al. 2021) and the age gradient (Kureishi et al. 2021) of the time preference distribution found in other empirical studies.

To gauge the quantitative effect of survival heterogeneity and the implied time preference heterogeneity on savings behavior and wealth accumulation, we use an overlapping-generations general-equilibrium model with uninsurable idiosyncratic shocks. Agents face heterogeneous survival risk that depends on their age and current health state, and are subject to health shocks that follow a process estimated from data. Besides this uncertainty, we also include standard persistent and transitory shocks to labor productivity, and after agents reach a fixed retirement age, they are entitled to retirement benefits mimicking the U.S. social security system. We purposely use an otherwise standard off-the-shelf model of consumption/savings to establish a benchmark and focus on the survival heterogeneity savings channel in isolation.

We compare three scenarios, varying the model environment according to how agents form expectations about survival. The first scenario serves as our baseline: this is the standard model where all agents face the same survival risk and thus have the same effective discount factor, conditional on age. In the second scenario, people are perfectly informed about their true survival probability conditional on health and age. In the third scenario, agents believe and act according to their subjective survival beliefs.

The model simulations show that the survival expectation channel is important for understanding wealth accumulation. Not surprisingly, agents in bad health and thus with a shorter expected life span save less than their healthy counterparts, and the differences in savings rates are large. For example, for 60-year-olds in the poorest wealth decile, the total savings rates of an agent in the best and an agent in the worst health state differ by 14 percentage points when they are equipped with correct objective beliefs about survival. When we let them act according to the estimated subjective beliefs instead, the difference grows to almost 30 percentage points.

The differences in savings behavior give rise to large differences in accumulated wealth. In our model simulation with subjective beliefs about survival, the average 65- to 69-year-old in the best health state has 51% more wealth than the average agent in the worst health state in the same age group. This is approximately one third of the difference we see in the data comparing the net asset holdings between individuals in the best and worst health states at this age.

Despite large differences in savings rates and large within-cohort effects on wealth accumulation, the wealth Gini is virtually unaffected by the inclusion of survival heterogeneity in the model. The reason is that the savings behavior of the richest individuals in the model economy, healthy agents in their early 60s, is hardly affected. The main effect in this age group is that unhealthy individuals save less. Therefore, the effect of heterogeneity in survival on the main driver of the Gini coefficient, namely the savings behavior of the top wealthiest individuals in the economy, is negligible.

Finally, we evaluate the welfare loss of biased beliefs. We examine the magnitude of the consumption equivalent loss for a rational expectation individual who is forced to use the consumption policy functions of an individual who optimized under the subjective (erroneous) beliefs. The resulting consumption equivalent varies from approximately zero to up to -11% for some subgroups, with the largest cost found among the very old rich and healthy who are overly optimistic about their survival and therefore save "too much." However, there are not many very rich old people in our simulated economy. Generally, the average cost of the belief bias is very small for young agents, but approaching the age of 50 the cost increases, especially for agents in bad health who are pessimistic about their survival chances and therefore save less than an individual with objective beliefs. The average consumption equivalent loss among agents in their early 50s in the worst health state is 1.5%.

It is not evident if overly optimistic individuals would be better off learning about the true probabilities. With a broader concept of welfare, taking into account the anticipation utility from the belief about a long life (as suggested by Brunnermeier and Parker 2005), the overly optimistic individuals might very well be better off with their erroneous beliefs. From this perspective, the 50-year-olds in bad health are particularly troublesome: their bias is unambiguously welfare reducing since it leads to both non-optimal consumption/savings decisions and lower anticipation utility (the latter not being part of our quantification).

This paper speaks to three broad strands of literature. The first is macroeconomic studies pointing out the importance of heterogeneity in time preferences to explain wealth inequality (e.g., Krusell and Smith Jr. 1998, Hendricks 2007, Quadrini and Ríos-Rull 2015, Krueger, Mitman, and Perri 2016) and studies documenting time preference heterogeneity in the population (Epper et al. 2020; Calvet et al. 2021). Compared to these studies, we are concerned about time preference heterogeneity that is micro-founded by differences in life expectancy.

The second is the literature about the general impact of health (including life expectancy) on wealth (Smith 1999; Lee and Kim 2008; Coile and Milligan 2009; De Nardi, French, and Jones 2009; Kopecky and Koreshkova 2014; Capatina 2015; De Nardi, Pashchenko, and Porapakkarm 2017; Poterba, Venti, and Wise 2017; Margaris and Wallenius 2020 to name a few). These studies incorporate multiple links between health and economic outcomes, while we restrict ourselves to a very specific channel, namely survival heterogeneity. We argue that this channel is interesting in itself, and that it is necessary to understand it fully in order to include it in the broader assessment of the health-wealth gradient (we briefly discuss the challenge of combining a standard formulation of warm-glow bequest motive and survival shocks in section 5.5). In contrast to this strand, we include heterogeneity in *subjective* life expectancy and its impact on savings/consumption behavior.

The third is the literature concerned with subjective survival expectations (Hamermesh 1985; Smith, Taylor, and Sloan 2001; Hurd and McGarry 2002; Ludwig and Zimper 2013; Elder 2013; Gan et al. 2015; Groneck, Ludwig, and Zimper 2016; Heimer, Myrseth, and Schoenle 2019). Many studies have documented the existence of an age bias in subjective life expectancies, and a few of the papers within this group are concerned with the implications for the consumption/savings behavior. Some (Gan, Hurd, and McFadden 2007; Bissonnette, Hurd, and Michaud 2017; Grevenbrock et al. 2021) predict individual survival probabilities and compare with elicited beliefs, but none of these look at the implications for within-cohort savings behavior in a structural model with beliefs change in the event of health shocks, or analyze the implications for wealth inequality.

In the next section, we describe how we estimate the health and death process and give details about the systematic bias in survival expectations. Section three describes the model we use to quantify the importance of the heterogeneity in survival expectations. After that, we discuss the parametrization and then we present our results. The last section concludes.

2 Empirical evidence

2.1 Data

We use the Health and Retirement Study (HRS), a representative panel of elderly U.S. households, to investigate the evolution of health and longevity in the later stages of life. The survey includes questions about self-reported health and expectations about survival, and records the date of death, if applicable.

Our analysis is based on the survey years 1992–2014 taken from the HRS data compiled by RAND, version 2018 (V1) (Health and Retirement Study (2018)).[1,2] The first cohort included in the survey was between 51 and 61 years old in 1992, and thereafter new (older and younger) cohorts have been included. Many of the respondents died over the sample period, making it an ideal data set for studying survival.

2.2 The health-wealth gradient

The HRS asks participants to assess their health using one of the five categories *excellent*, *very good*, *good*, *fair*, or *poor*. Figure 1 shows net total wealth over the life cycle by self-reported health state.[3] The health–wealth gradient is well documented, but the underlying causal relationship is debated (Attanasio and Hoynes 2000; Deaton 2002; Duncan et al. 2002; Attanasio and Emmerson 2003; Hajat et al. 2010). One line of

[1] Up until RAND version O (covering waves until 2012), the survey was complemented with death dates taken directly from the National Death Index (NDI), but this data was later removed from the public files. Our analysis of death dates in the releases following version O shows that without the NDI data, death dates are sometimes recorded with considerable lag. Using the RAND 2018 (V1) files, but restricting the sample only up to the year 2012, produces almost identical results to the ones obtained with the original version O data that included the NDI death dates. However, for later years we suspect that not all death dates have been recorded yet, which we believe gives rise to the non-response patterns documented in section A.1. Based on these non-response patterns, we decided to only include waves up to and including the year 2014.

[2] The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

[3] Net total wealth is defined as sum of housing, other real estate, vehicles, businesses, IRA and Keogh accounts, stocks, checkings, and all other savings, net of mortgages and other debts.

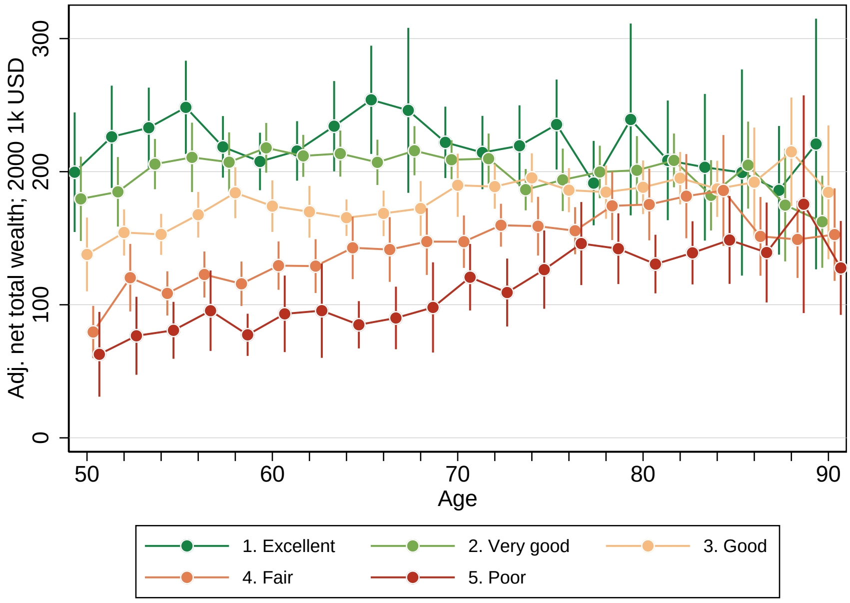

Figure 1: Average net total household wealth by self-reported health state (men). Pooled sample from HRS 1992–2014. Assets are adjusted for inflation, outliers, business cycles and growth and cohort effects. Error bars indicate 95% confidence intervals based on clustered standard errors (clusters are HRS PSUs).

argument is that low economic status leads to poor health. There could be many reasons: poor people have access to less or lower quality medical care, do not invest enough in preventive health measures, and/or have more health-deteriorating habits. However, there are also many arguments for the reversed causality: poor health has economic consequences in itself. First, poor health may restrict the individual's earnings potential by making it more costly to work and/or by lowering the wage. Second, poor health may lead to large medical expenditures. Third, poor health may lower the savings incentives due to a lower survival expectancy. This last channel is the focus of this paper.

There are a number of empirical studies that corroborate the existence of the life expectancy/savings channel and suggest a causal link. For instance, Heimer, Myrseth, and Schoenle (2019) demonstrate that greater survival optimism correlates with higher savings rates, not only controlling for standard demographic characteristics such education, marital status, and income, but also characteristics such as financial literacy and risk tolerance. Hurd, Smith, and Zissimopoulos (2004) show that individuals with very low subjective survival probabilities retire and claim social security benefits earlier.

Two recent empirical studies suggest a causal life expectancy link using the plausibly exogenous timing of the bad news. Kvaerner (2020) shows that news about a bad health shock increases the probability of an immediate intervivo transfer while Fernström (2019) finds that annual savings fall by 5 percent in response to a one standard deviation fall in survival probability.

2.3 Objective health and survival probabilities

In this paper, we examine the effect of heterogeneity in survival expectancy on savings behavior and its implications for wealth inequality through the lens of a structural model. Therefore, we need to formulate heterogeneity in survival expectations, both objective and subjective, in a way that can be used in a such a model.

For our quantitative model, we need a yearly Markov process for health transitions and survival as a function of the model's state variables. We estimate this Markov process as described in Foltyn and Olsson (2019). Conceptually, the method is a straightforward maximum likelihood estimator, where the probability of observing the transitions in the data is maximized.

To put structure on the Markov process, we follow Pijoan-Mas and Ríos-Rull (2014) and use a nested logit model, where survival and health transitions conditional on survival are modeled as functions of the current health state and age. The probability of survival follows the usual binary-outcome logit model while, conditional on survival, health transitions are modeled using multinomial logit. For example, the one-period-ahead probability of survival is given by

$$p^s_{t+1} = \frac{1}{1 + e^{-g(h_t, \boldsymbol{x}_t | \gamma)}} \tag{1}$$

where $g(\bullet)$ is a function of the current health state h_t and a vector $\boldsymbol{x}_t$ which contains any other variable of interest, in particular age, gender and race. Survival probabilities are governed by the parameter vector γ to be estimated. Transition probabilities for health conditional on survival are defined in an analogous manner.[4]

Estimation sample. We exclude all observations with missing age, race, gender or self-reported health, and those where we only have a single observation for the individual (since then we do not have any transition probability to estimate). We only consider individuals aged 50 or older.[5] Further, we restrict the sample to maximum age 99 at transition start (even though individuals can be older when we observe them in the end of a transition). We estimate the health and objective (statistical) survival process

[4]In Foltyn and Olsson (2019), we provide details about the estimation and also perform an extensive evaluation of the results. The estimated Markov process is shown to predict actual mortality very well, both short- and long-term.

[5]Each incoming HRS cohort is aged 51–56, but the survey contains younger individuals who are spouses of age-eligible respondents.

		All	Non-black		Black	
			Male	Female	Male	Female
N. of indiv.		34,194	12,744	15,459	2,423	3,568
N. obs.		219,647	81,314	103,904	13,269	21,160
Avg. obs./indiv.		6.4	6.4	6.7	5.5	5.9
Age	Min.	50.0	50.0	50.0	50.0	50.0
	Mean	66.7	65.9	67.6	64.6	66.0
	Max.	101.0	101.0	101.0	101.0	101.0

Table 1: Descriptive statistics for sample used to estimate objective health and survival transitions. Mean age is weighted using HRS sample weights.

separately for the subsamples of men/women and the nonblack/black population, since it is well known that the life expectancies for these subgroups follow very different trajectories. Table 1 shows the number of individuals and number of observations by subgroup.

Results. From these estimates, we construct a first-order Markov process defined on five health states and the absorbing state of death which governs the objective health and survival probabilities. This process can be used to calculate objective life expectancies conditional on age, health, race, and gender.

Figure 2 illustrates the dynamics of this health and survival process: the figure shows the evolution of probabilities of each health state and of being dead over a 30-year forecast horizon for a given initial health and age for nonblack men. As can be seen, the survival probabilities differ substantially depending on the initial health state: for a 70-year old in excellent health, the predicted probability of surviving an additional 10 years is approximately 75%, while the probability is just around 35% if instead starting out in poor health.

2.4 Expectation errors in survival probabilities

In the expectations survey module of the HRS, respondents are asked about the probability they assign to certain events. One of these questions is about the probability of surviving to a certain age, for example: "Using a number from 0 to 100, what do you think are the chances that you will live to be at least 100 years?"[6]

The exact target age depends on the respondent's age. For instance, in 1995, respondents below the age of 70 were asked about the probability of living until the age of 80, while respondents above the age of 85 were asked about the target age of 100.

[6]Before the respondent answers the questions about expectations, the interviewer discusses probabilities and verifies that the respondent understands the concept.

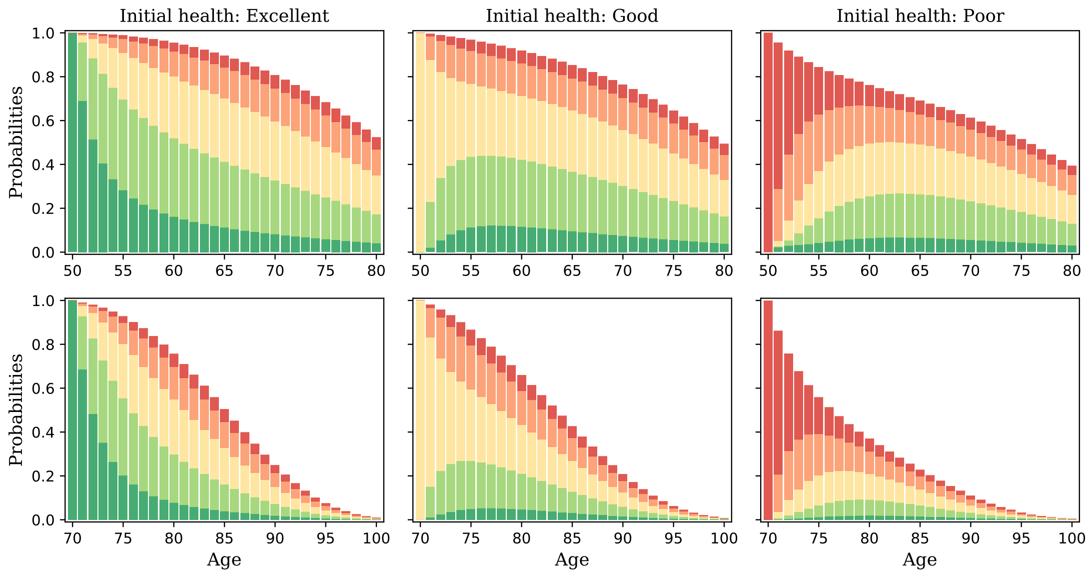

Figure 2: Predicted health and survival probabilities conditional on the initial health for nonblack men. The x-axis indicates years after initial age (upper row 50, lower row 70 years). The colors indicate probability per health state (dark green being the best health state, red the worst). The white area represents the probability of being dead.

In later surveys, individuals were asked about survival beliefs for up to two target ages.

Table 2 shows the number of individuals and observations (with an observation being one elicited survival belief) by subgroup. The tabulated distribution of target ages shows that the questions about survival to age 75 or 85 are the by far most common, hence we focus on these two in the main text.

Using these elicited beliefs, we compare the average probability that individuals of a certain age assign to survival until a given target age to the probability according to official lifetables. The results are shown in Figure 3. As can be seen, there is a systematic error along the age gradient: younger individuals on average tend to underestimate, while older individuals tend to overestimate their survival probability as compared to objective lifetable estimates.

This age-dependent error is a stylized fact in the literature on survival expectations (see, e.g., Ludwig and Zimper (2013), Groneck, Ludwig, and Zimper (2016), and Heimer, Myrseth, and Schoenle (2019) and the references therein). The pattern has been used to, e.g., improve the fit of the asset profile of the canonical life-cycle model with the data: due to underestimation early in life, young agents do not accumulate as much assets, while overestimation in later years dampens the rate at which assets are decumulated.

		All	Non-black		Black	
			Male	Female	Male	Female
N. of indiv.		31,784	11,761	14,537	2,195	3,291
N. obs.		251,073	91,505	119,954	14,796	24,818
Avg. obs./indiv.		7.9	7.8	8.3	6.7	7.5
Forecast horizon (years)	Min.	10.0	10.0	10.0	10.0	10.0
	Mean	19.3	19.4	19.1	20.1	19.8
	Max.	35.0	35.0	35.0	35.0	35.0
Target age distribution (in %)	Age 75	37.1	37.6	36.2	40.9	39.5
	Age 80	16.3	16.3	16.4	15.4	16.2
	Age 85	35.1	35.6	34.4	37.1	36.1
	Age 90	5.6	5.3	6.1	3.7	4.3
	Age 95	3.9	3.5	4.5	2.0	2.6
	Age 100	2.0	1.7	2.4	0.9	1.3

Table 2: Descriptive statistics for elicited subjective survival beliefs. Mean forecast horizon and distribution over target ages are weighted using HRS sample weights.

Since our estimator conditions on health, gender and race, we can go one step further and document survival belief differences along these dimensions. Figure 4 shows these expectation errors among respondents who have answered the question about their perceived probability of survival until target ages 75 and 85. The first observation is the striking positive correlation between subjective self-reported survival probability and the predicted (objective) survival probability, which means that subjective beliefs are informative and not just random noise. This is in line with the consensus in the literature, which finds that subjective beliefs are highly correlated with objective survival probabilities and serve as predictors of mortality, and moreover that expectations are updated in the event of health shocks (Smith, Taylor, and Sloan 2001; Hurd and McGarry 2002; Gan, Hurd, and McFadden 2007).

The second observation is the systematic steepness bias in beliefs over the health gradient. As was shown in Figure 3, on average individuals underestimate their probability of survival until the age of 75. Looking at Figure 4(a), and focusing on nonblack males, it is clear that it is mainly individuals in bad health who are underestimating their survival probability, while individuals in excellent health are on average reasonably close to their objective survival probability.

Figure 4(b) shows expectation errors for target age 85. Again, individuals in bad health are more pessimistic than those in good health, even though on average the expectations are more positive for this target age. Figure 17 in the appendix section A.2 shows the corresponding graphs for target age 95, for which the average is a severe over-optimism. The figure shows that for those higher ages individuals in bad health have beliefs closer to their objective probability, while individuals in good or excellent health are severely overestimating their survival probability and driving up the average.

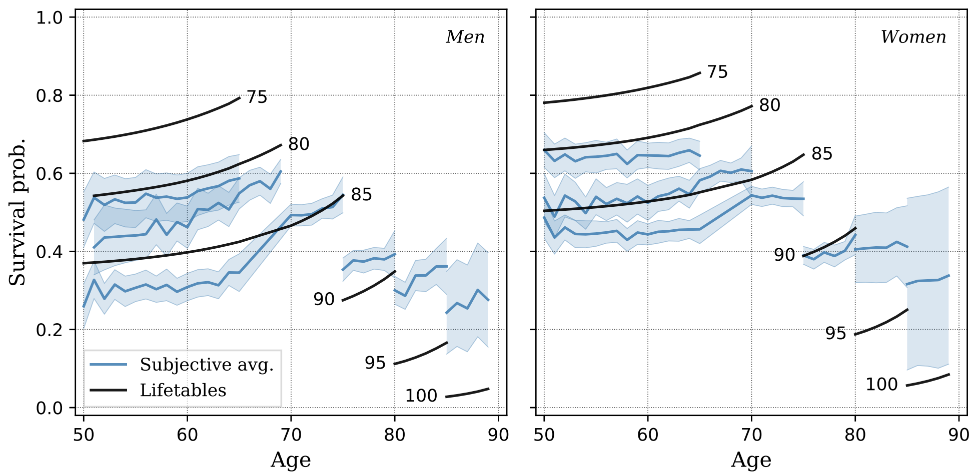

Figure 3: Objective vs. subjective survival probabilities by age and gender (controlling for year and cohort fixed effects). The number next to the black line indicates target age. The blue line shows (weighted) average expectation in the nonblack population. Shaded areas indicate 95% confidence intervals based on clustered standard errors (clusters are HRS PSUs).

Next, we impose some structure on the biased survival beliefs illustrated in these scatter plots. To this end, we define survival bias as the elicited subjective survival belief minus objective survival probability (thus pessimism is negative bias). For example, for some target age T we define the bias as

$$\Delta_{igt}^{T} \equiv p_{subj,igt}^{T} - \Pr\left(\text{ alive at } T \mid g, h_{it}, t \right)$$

for individual i in demographic group g (male/nonblack, female/nonblack, male/black, or female/black) when the individual is in health state h_{it}.

We first quantify the bias observed over the lifecycle by estimating the regression

$$\Delta_{igt}^{T} = \beta_{0g} + \beta_{1g} \cdot \text{age}_{it} + \beta_{2g} \cdot \text{age}_{it}^{2} + u_{it}$$

for each horizon T separately, where we allow the coefficients of the age polynomial to vary by race and gender. In Figure 5, we plot the predicted survival bias for each age and gender/race group for the target ages 75 and 85,

$$\widehat{\Delta}_{gt}^{T} = \widehat{\beta}_{0g} + \widehat{\beta}_{1g} \cdot \text{age}_{t} + \widehat{\beta}_{2g} \cdot \text{age}_{t}^{2} \qquad T \in \{75, 85\}$$

As the figure shows, men and black individuals are on average more optimistic than women and nonblack individuals. This confirms the findings by Bissonnette, Hurd, and Michaud (2017).

In a second step, we disaggregate the bias by initial health. To this end, we estimate the following regression separately for each demographic group,

$$\Delta_{ight}^{T} = \beta_{0gh} + \beta_{1gh} \cdot \text{age}_{it} + \beta_{2gh} \cdot \text{age}_{it}^{2} + \varepsilon_{it}$$

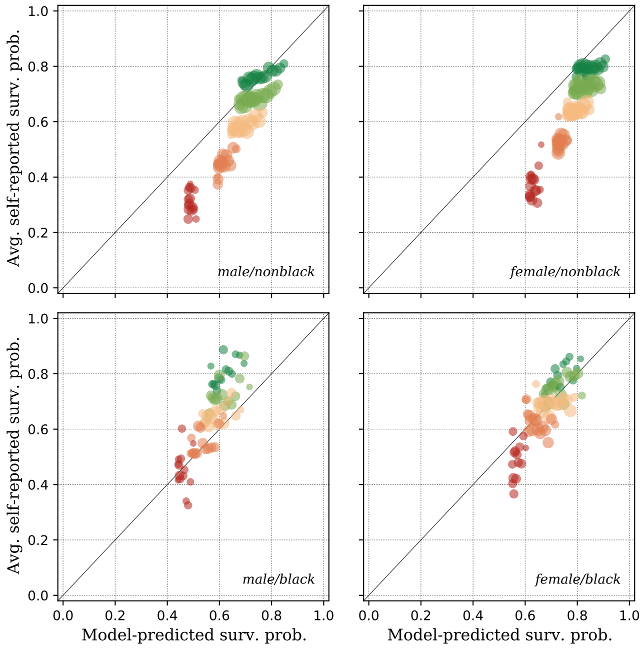

(a) Target age 75

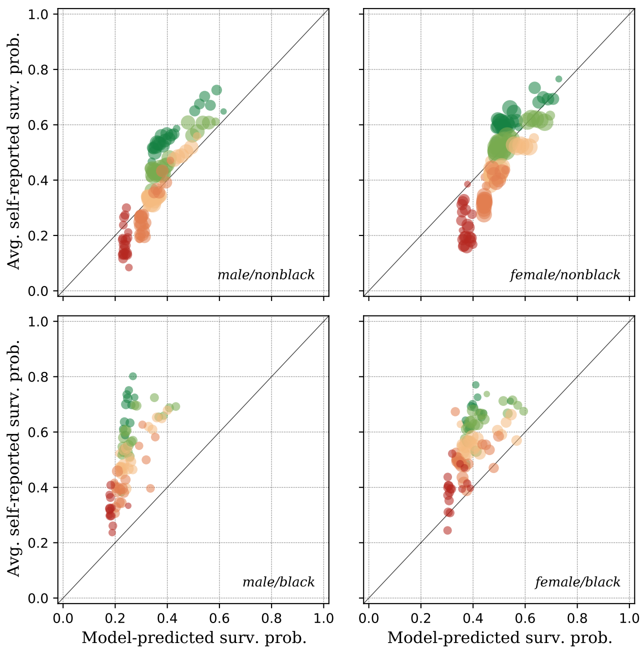

(b) Target age 85

Figure 4: Elicited beliefs about survival vs. estimated objective (statistical) survival probabilities. Each bubble represents the average for a gender/race/age/health group. The x-axis shows the model-predicted (objective) survival probability, the y-axis the average self-reported survival probability for that group. The color of the bubble indicates health state, with red being poor health and green being excellent health. The size indicates the number of observations in each age/health cell. We exclude cells with less than 20 observations.

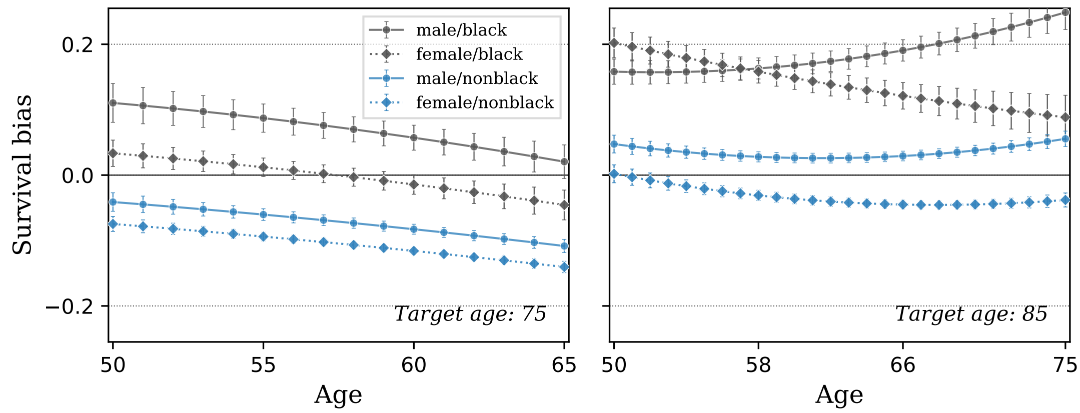

Figure 5: Survival bias. Error bars indicate 95% confidence intervals based on clustered standard errors (clusters are HRS PSUs).

where we interact the initial health state h with a quadratic polynomial in age. The predicted values for this exercise for target ages 75 and 85 are shown in Figure 6: individuals in bad health are generally more pessimistic than individuals in better health. The pattern is more pronounced for the nonblack population, for which we have more observations. Thus, Figure 6 confirms the steepness bias observed in the previous analysis: people in bad health, and thus with lower life expectancy, are more pessimistic about their survival probabilities than people in good health.

In a related paper, Grevenbrock et al. (2021) estimate survival based on several additional characteristics besides self-reported health, age, and gender, such as smoking and drinking behavior and chronic diseases. Grouping individuals based on their estimated objective survival probability, they find that individuals with low objective survival probability are optimists, while individuals with high objective probability underestimate their survival probabilities, in other words, the reverse pattern compared to what we find. There are two reasons for why our results are not directly comparable. First, they use a different estimation strategy for objective survival probabilities. Second, and more important, the grouping of individuals is different. We group based on age, gender, race, and self-reported health, which is the level of heterogeneity in our economic model (and thus, for our particular purposes, the most appropriate one) and show that there is an steepness bias along the health gradient within each subgroup.

More generally, the finding that individuals with high life expectancy are overly optimistic are also in line with evidence about forecast errors in other domains. For example, Rozsypal and Schlafmann (2017) document that people in the upper part of the income distribution overestimate their future income growth while the opposite is true for lower income households: they are too pessimistic and underestimate their future income growth.

To summarize, we stress two observations: first, subjective beliefs are informative and

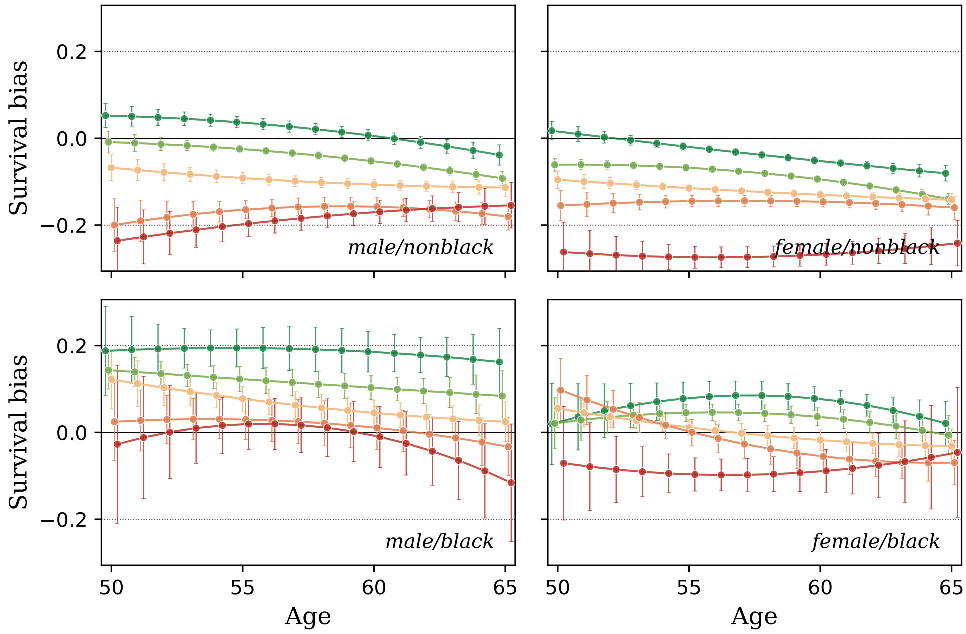

(a) Target age 75

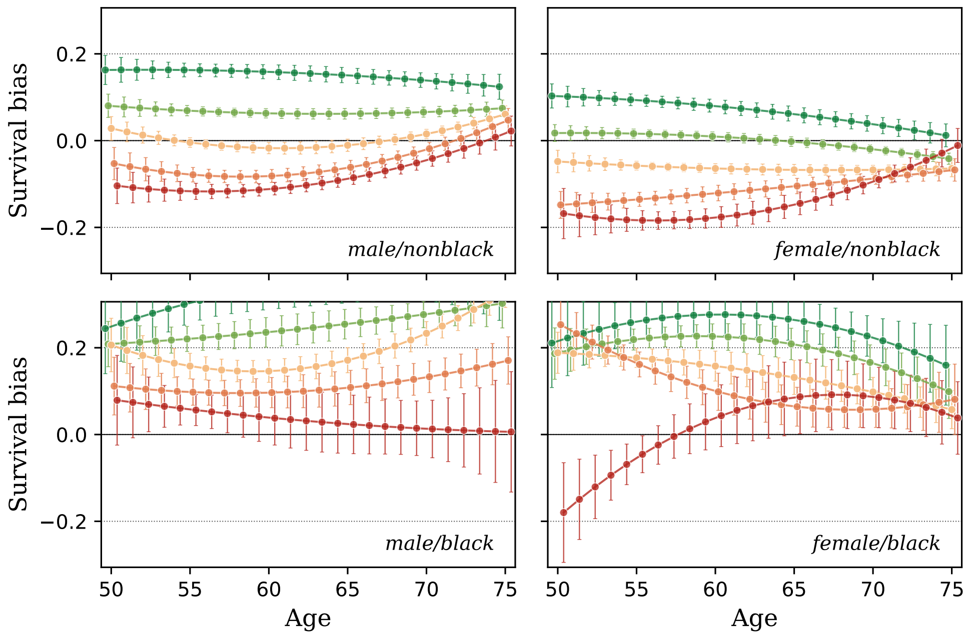

(b) Target age 85

Figure 6: Survival bias. Color indicates the health state: green is excellent while red is poor health. Error bars indicate 95% confidence intervals based on clustered standard errors (clusters are HRS PSUs).

correlated with objective probabilities. Second, subjective beliefs are biased. Subjective probabilities overestimate the health/survival gradient, with individuals in bad health underestimating their survival probability relative to individuals in good health. Hence, there is a systematic bias both along the age and health dimensions.

2.5 Estimation of the subjective life expectancy process

In this section, we use the health transitions estimated in Foltyn and Olsson (2019) as a basis for estimating a different set of survival parameters that govern *subjective* survival beliefs. We take as given the parameters controlling health-to-health transitions conditional on survival since the HRS does not elicit any beliefs about future health states.

We take an agnostic approach as in *why* the erroneous beliefs arise. In the literature, various mechanisms have been proposed, such as likelihood insensitivity combined with Choquet expected utility maximization (Groneck, Ludwig, and Zimper 2016), overweighting the likelihood of rare events (Heimer, Myrseth, and Schoenle 2019), or age-dependent cognitive weakness and relative optimism (Grevenbrock et al. 2021).

As explained above, the underlying data for this exercise takes the following form: HRS respondents are asked at date t to state their beliefs about surviving to a certain target age $\bar{a}$ (for example 75 or 85), which we reinterpret as the probability of being alive in period T, with $T = t + (\bar{a} - a_t)$. Thus, an observation i is given by the tuple

$$(p_i, h_i, \boldsymbol{x}_i, T_i)$$

where h_i denotes current health state, $\boldsymbol{x}_i$ is a vector of covariates including age and p_i is the subjective survival belief. We treat multiple observations from one individual independently: say a respondent ℓ is surveyed on survival beliefs to horizons T^1 and T^2 in calendar years t_1 and t_2. This gives rise to the data

$$\begin{gathered}\left(p_{\ell t_1}, h_{\ell t_1}, \boldsymbol{x}_{\ell t_1}, T^1_{\ell t_1}\right)\\ \left(p_{\ell t_1}, h_{\ell t_1}, \boldsymbol{x}_{\ell t_1}, T^2_{\ell t_1}\right)\\ \left(p_{\ell t_2}, h_{\ell t_2}, \boldsymbol{x}_{\ell t_2}, T^1_{\ell t_2}\right)\\ (p_{\ell t_2}, h_{\ell t_2}, \boldsymbol{x}_{\ell t_2}, T^2_{\ell t_2})\end{gathered}$$

which we treat as four independent observations (except when bootstrapping confidence intervals, which we cluster at the individual level).

Assume that the i-th individual forms T-year-ahead survival beliefs based on the model

$$p^s_{it} = \phi_T\left(h_{it}, \boldsymbol{x}_{it}, \boldsymbol{z}_{it}\right) \tag{2}$$

where ϕ_T is an unknown nonlinear function that maps $(h, \boldsymbol{x}, \boldsymbol{z})$ into $[0, 1]$. The respondent's belief is allowed to depend on a vector of additional covariates $\boldsymbol{z}$ that are either unobserved or not included in our postulated model of survival.

In what follows, we partition the sample into groups indexed by g, such that each unique combination of (h, x, T) forms a separate group. Denote by Γ_g all individual/year observations that satisfy

$$\Gamma_g = \left\{ (i,t) \,\middle|\, h_{it} = h_g, \boldsymbol{x}_{it} = \boldsymbol{x}_g, T_{it} = T_g \right\}$$

i.e., all observations where the individuals are of the same age and have the same covariates, are in the same health state, and state their beliefs about survival to the same target age. Denote by $\overline{p}_g^s$ the (weighted) sample average of reported survival beliefs conditional on $(h_g, \boldsymbol{x}_g, T_g)$, i.e.,

$$\overline{p}_g^s = \frac{\sum_{(i,t)\in\Gamma_g} w_{it} \times \phi_T \left(h_g, \boldsymbol{x}_g, \boldsymbol{z}_{it} \right)}{\sum_{(i,t)\in\Gamma_g} w_{it}} \tag{3}$$

with w_{it} denoting sampling weights. Now consider the logit counterpart of (3), which we denote by

$$\widehat{p}_g^s = \Pr \left(\text{alive at } T_g \,\middle|\, h = h_g, \boldsymbol{x}_g, \boldsymbol{\nu} \right)$$

i.e., the predicted probability of being alive for group g which is parametrized by the vector $\boldsymbol{\nu}$. The observed sample moment for each group can then be written as

$$\overline{p}_g^s = \widehat{p}_g^s + u_g$$

where u_g is the deviation from the group mean not explained by our model. Our aim is to minimize these group-specific residuals using the least-squares objective function

$$J(\boldsymbol{\nu}) = \frac{1}{W} \sum_{g=1}^{N_G} W_g \left(\overline{p}_g^s - \widehat{p}^s \left(h_g, \boldsymbol{x}_g, T_g \,\middle|\, \boldsymbol{\nu} \right) \right)^2 \tag{4}$$

where $W_g = \sum_{(i,t)\in\Gamma_g} w_{it}$ is the sum of weights in group g. The estimated vector $\widehat{\boldsymbol{\nu}}$ is hence the arg min of $J(\boldsymbol{\nu})$.

Estimation sample. We use all target ages from Table 2 for the estimation of the subjective life expectancy process. In the main paper, we present the results for nonblack men, as these are later incorporated into our quantitative model.

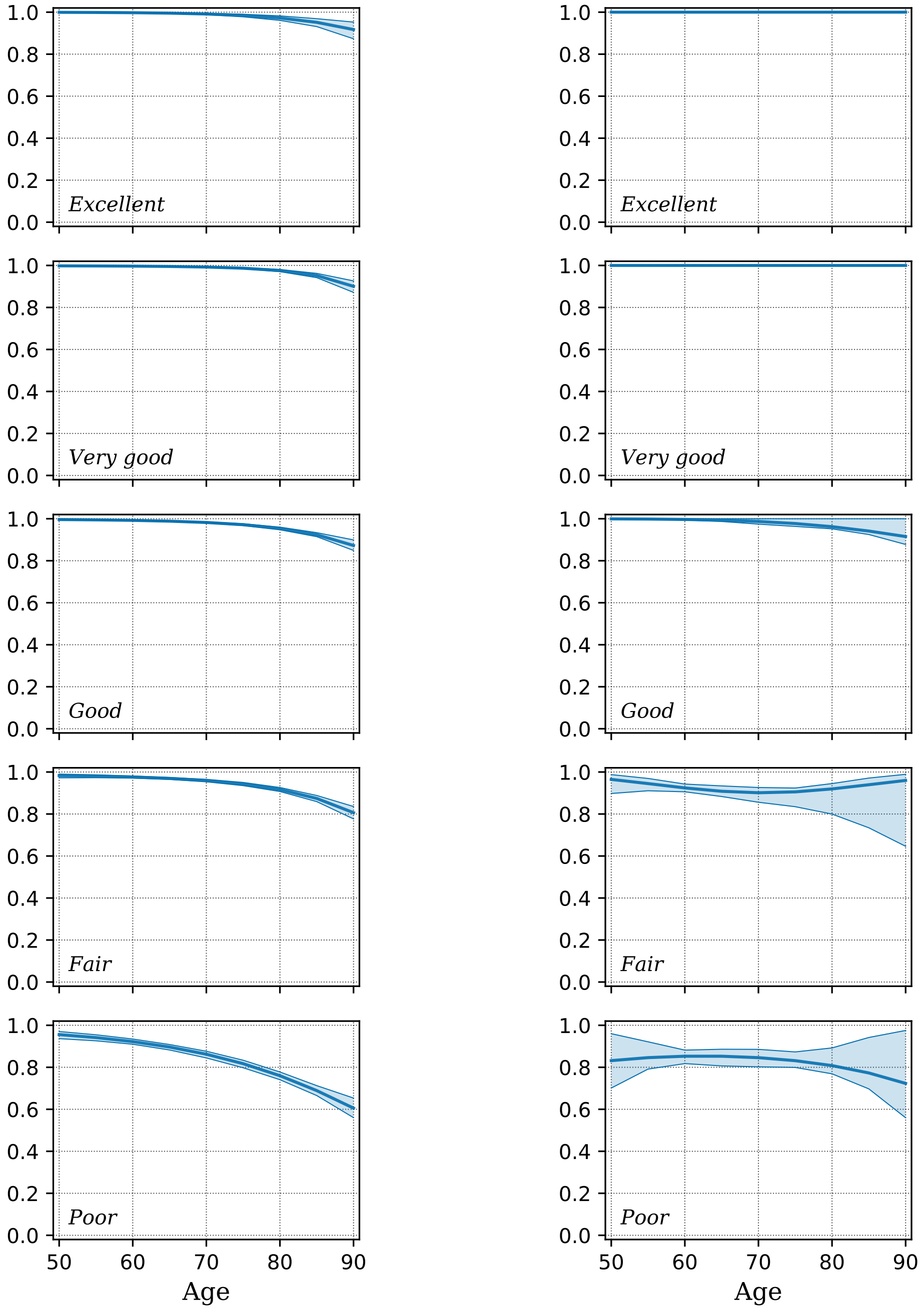

(a) Objective probabilities

(b) Subjective probabilities

Figure 7: One-year survival probabilities by health state (model estimates). Shaded areas indicate bootstrapped 95% confidence intervals. For each bootstrapped sample we re-estimate the objective health process.

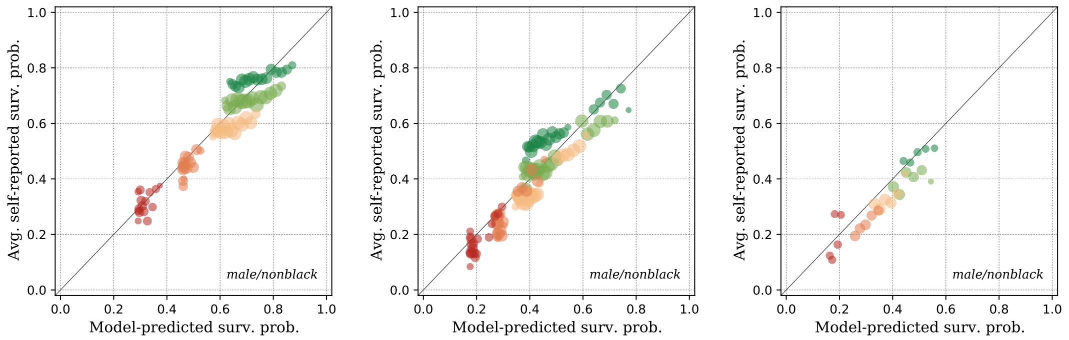

(a) Target age 75 **(b)** Target age 85 **(c)** Target age 95

Figure 8: Elicited beliefs about survival vs. estimated *subjective* survival probabilities. Each bubble represents the average for an age/health group. The x-axis shows predicted survival probability according to the subjective model and the y-axis the average self-reported survival probability for that group. The color of each bubble indicates health state, with red being poor health and green being excellent health.

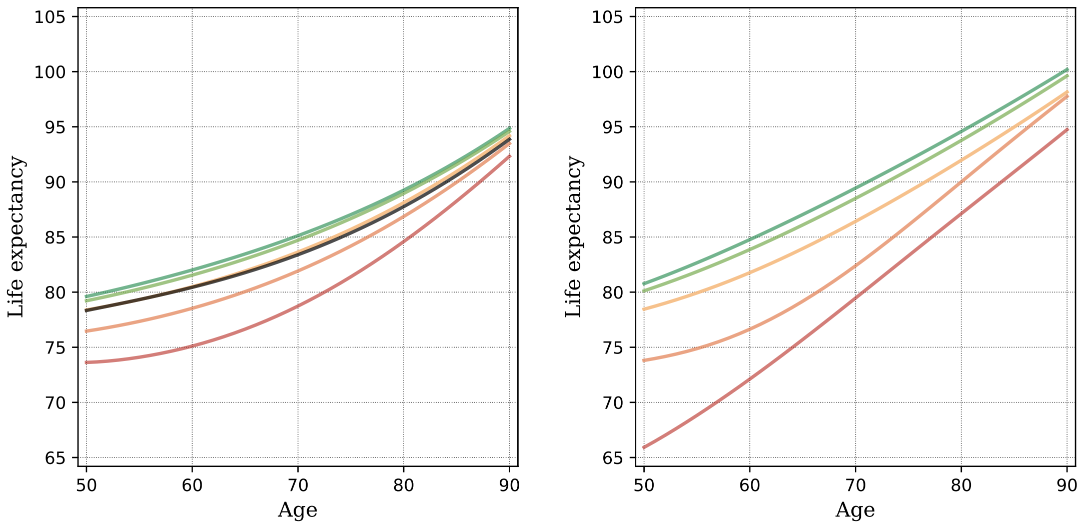

(a) Objective probabilities **(b)** Subjective probabilities

Figure 9: Life expectancy by age and health for nonblack men. Color indicates the health state: green is excellent while red is poor health. On the left, the black line indicates the weighted population average.

Results. The estimated subjective survival beliefs for nonblack men are shown in Figure 7 on the right, juxtaposing the objective survival probabilities estimated in Foltyn and Olsson (2019) on the left. As can be seen, the subjective belief about survival while in health state *excellent* or *very good* is almost 100% for all ages. This does not mean that individuals in those health states believe that they will live forever, rather that they believe that death is necessarily preceded by a deterioration in health.

In Figure 8, we plot the model-predicted subjective survival against elicited beliefs. As can be seen, the estimated model for subjective beliefs captures the main picture, since the dots, each representing an age/health/target-age group, line up closely along the 45-degree line.

Figure 9 summarizes the results, showing the life expectancy by age and health state using the objective and the subjective survival process. At all ages, the difference in life expectancy between the best and the worst health state is larger when using subjective life expectancies. The divergence between objective and subjective life expectancies is particularly large for individuals in bad health states, who substantially underestimate survival at younger ages. Conversely, individuals in all health states overestimate their chances of survival late in life.

3 Model

In this section, we describe the overlapping-generations model we use to quantify the implications of survival heterogeneity. Time is discrete and every time period is assumed to be one year. Agents derive utility from consumption and face three types of idiosyncratic risks: shocks to persistent productivity, transitory productivity shocks and shocks to health and survival. Agents can only save in a riskless bond, and they face an exogenous borrowing constraint.

3.1 The agent's problem

There is a unit mass of individuals distributed across N_t cohorts according to the ergodic distribution implied by the transition matrix of survival probabilities. An individual of age $t \in \{1, \dots, N_t\}$ and health $h \in \{1, \dots, N_h\}$ has a one-period survival probability to age $t+1$ given by π^s_{th}, with $\pi^s_{th} = 0$ for $t = N_t$ regardless of health state.

Individuals are assumed to be working for the first $T_R - 1$ years of their life and exogenously retire in the period when they attain age T_R. While working, they are hit by persistent and transitory labor productivity shocks. During retirement, individuals receive social security retirement benefits which depend on their last persistent labor productivity in working age.

Bequests are distributed to new-born individuals in a dynastic way: one young individual receives the bequests from one dying individual.

Retired agents. A retired individual in state (a, p, h, t), where a is cash-at-hand, p is the (fixed) persistent component of labor earnings, h is the current health state, and t is the age, maximizes utility according to

$$V_R\left(a, p, h, t\right) = \max_{c,\, b'} \left\{ u(c) + \beta \pi^s_{th} \mathbf{E}\left[\, V_R\left(\boldsymbol{x}'\right) \, \middle| \, h, t \right] \right\} \tag{5}$$

subject to the constraints

$$\begin{aligned} a &\geq c + b' \,, \quad c \geq 0 \,, \quad b' \geq 0 \\ a' &= Rb' + \iota'_R \\ \iota'_R &= y'_R w - T_y\left(y'_R w\right) \\ y'_R &= \omega_{T_R - 1} p_R(p) \bar{\epsilon} \end{aligned} \tag{6}$$

where $\boldsymbol{x}' = (a', p, h', t+1)$ is the continuation state conditional on survival. Next-period after-tax retirement income is denoted by ι'_R and depends on the non-linear tax schedule $T_y\left(\bullet\right)$. $p_R(\bullet)$ is a function mimicking the regressive replacement rate of the U.S. social security system, w is the economy-wide wage rate, $\bar{\epsilon}$ is the average of the transitory earnings shocks hitting the working-age population and ω_{T_R-1} is the value of the deterministic age profile of earnings just prior to retirement. The gross return on savings is given by $R = 1 + r - \delta_k$.

Working-age agents. We denote the value and policy functions of working-age agents using the subscript W.

A working-age individual draws a persistent and a transitory labor shock component that together with a deterministic age profile of earnings, ω_t, pin down his labor productivity. The persistent component p takes on the values from the set $\mathcal{P}$, while the transitory shock realizations ϵ are drawn from $\mathcal{E}$.

Working-age individuals with $t < T_R - 1$ who continue working next period solve

$$V_W\left(\boldsymbol{x}\right) = \max_{c,\, b'} \left\{ u(c) + \beta \pi^s_{th} \mathbf{E}\left[\, V_W\left(\boldsymbol{x}'\right) \, \middle| \, p, h, t \right] \right\}$$

$$\begin{aligned} \text{s.t.} \quad a &\geq c + b' \,, \quad c \geq 0 \,, \quad b' \geq 0 \\ a' &= Rb' + \iota' \\ \iota' &= \left[y' - T_{ss}\left(y'\right)\right] w - T_y\left(\left[y' - T_{ss}\left(y'\right)\right] w\right) \\ y' &= \omega_{t+1} p' \epsilon' \end{aligned} \tag{7}$$

where $\boldsymbol{x}' = (a', p', h', t+1)$. We denote the earnings of working individuals, net of income taxes $T_y(\bullet)$ and payroll taxes $T_{ss}(\bullet)$, as ι'.

In the final period of their working life, i.e., when $t = T_R - 1$, individuals solve

$$V_W(a, p, h, t) = \max_{c,\, b'} \left\{ u(c) + \beta \pi^s_{th} \mathbf{E} \left[V_R(a', p, h', t+1) \,\middle|\, h, t \right] \right\}$$

subject to

$$\begin{aligned} a &\geq c + b' \,, \quad c \geq 0 \,, \quad b' \geq 0 \\ a' &= Rb' + \iota'_R \\ \iota'_R &= y'_R w - T_y(y'_R w) \\ y'_R &= \omega_{T_R - 1} p_R(p) \bar{\epsilon} \end{aligned}$$

which is identical to the retired individual's problem.

3.2 Technology

The production side of the model is standard. Competitive firms employ labor and capital hired from households to produce a homogeneous final good, which is used for both consumption and investment. The aggregate production function is assumed to be Cobb-Douglas:

$$F(K, L) = K^{\alpha_k} L^{1-\alpha_k}.$$

Capital depreciates at the rate δ_k.

3.3 Government

We assume that the government runs a PAYGO social security system that has to balance in each period, and that remaining (wasteful) government expenditures have to be fully financed by income taxes. We first describe the social security system and thereafter the general government budget.

3.3.1 Social security system

We use a stylized version of the actual retirement income formula used in the U.S. social security system. It captures the main features, such as a regressive replacement rate based on pre-retirement income and a cap for maximum benefits. In the model, we define retirement benefits to be a product of the economy-wide wage rate, the life-cycle profile wage component from the last year before retiring, the average transitory component, and a function that mimics the regressive replacement rate of the U.S. social security system:

$$\iota_R(p) = w \times y_R(p) = w \times \omega_{T_R-1}\bar{\epsilon} p_R(p)$$

where the replacement function $p_R(\cdot)$ is given by

$$p_R(p) = \begin{cases} \rho_1 p & \text{if } p \leq \widetilde{p}_1 \\ \rho_1 \widetilde{p}_1 + \rho_2 \left(p - \widetilde{p}_1\right) & \text{if } \widetilde{p}_1 < p \leq \widetilde{p}_2 \\ \rho_1 \widetilde{p}_1 + \rho_2 \left(\widetilde{p}_2 - \widetilde{p}_1\right) + \rho_3 \left(\min\left\{\widetilde{p}_{max}, p\right\} - \widetilde{p}_2\right) & \text{else} \end{cases}$$

where $\widetilde{p}_1$ and $\widetilde{p}_2$ are bend points and $\widetilde{p}_{max}$ the contribution and benefit base (CBB) in the social security income formula, expressed in terms of the individual's permanent labor state. In the appendix, section B.1.1, the transformation between actual bend points and CBB expressed in USD ($b_1^{\$}$, $b_2^{\$}$, and $e_{max}^{\$}$) and the model counterparts are described in detail, as well as the derivation of total government expenditures on retirement, G_{ss}.

The government expenditures on retirement are financed by a payroll tax. The payroll tax function is defined as

$$T_{ss}\left(y\right) = \tau_{ss} \times \min\{y_{max}, y\}$$

where y_{max} expresses maximum taxable earnings in terms of labor productivity, i.e.,

$$y_{max} = \left(e_{max}^{\$} / e_{med}^{\$}\right) y_{med}$$

where $e_{max}^{\$}$ is the contribution and benefit base expressed in USD as above, and $e_{med}^{\$}$ are the median earnings in the reference year.

The derivation for total payroll taxes raised in each period, $\overline{T}_{ss}$, can be found in the appendix, section B.1.2. To balance the social security system, we need to find τ_{ss} such that $G_{ss} = \overline{T}_{ss}$.

3.3.2 Government budget

We assume that the government raises labor income taxes to finance non-discretionary expenditures that amount to a constant fraction g of output.

Income taxes. For income taxes, we adopt the same tax function as in Heathcote, Storesletten, and Violante (2017), which is defined as

$$T_y(\iota) = \iota - \lambda \iota^{1-\tau} \tag{8}$$

where ι is either earnings (net of payroll taxes) or retirement income. We assume that the progressivity parameter τ is fixed, and we pin down λ such that the government budget is balanced in each period.

Total income taxes raised by the government amount to the sum of income taxes raised from the working and from the retired individuals:

$$T_{inc} = T_W + T_R$$

The derivation of the total income tax raised by the government in each period can be found in the appendix, section B.2.

Government budget balance. The non-discretionary expenditures amount to a constant fraction g of output: $G = gY$. The government budget is balanced by solving for a λ in (8) such that $G = T_{inc}$ holds.

The equilibrium definition of the model is standard and can be found in the appendix, section B.3.

4 Calibration

4.1 Preferences

We assume log preferences, $u(c) = \log c$.[7] The value of β, the discount factor, is set to target a capital-to-output ratio of 3.0, which gives us a β of 0.981.[8]

4.2 Externally calibrated parameters

Demographics. Agents are assumed to enter the economy at age 20, and retire at the age of 65, which corresponds to setting $T_R = 46$. The maximum age an agent can reach is 99, and hence we let $N_t = 80$.

[7] In the Appendix section C.4 we show results from an alternative calibration with Epstein-Zin-Weil preferences.

[8] For this calibration, we use the model where everyone faces the same average survival probability conditional on age.

Age-dependent wage profile and idiosyncratic earnings risk. We assume that the log-labor earnings of an individual in the labor force follow a process with transitory and persistent shocks:

$$\log y_t = \log \omega_t + \log p_t + \log \epsilon_t$$

where ω_t is the age profile part, p_t is the persistent component and ϵ_t is the transitory component of earnings. The persistent component is assumed to follow an AR(1) process,

$$\log p_t = \rho \log p_{t-1} + \eta_t$$

with persistence ρ and innovation $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$. The transitory shock is given by $\log \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

Hence, the stochastic part of the wage process is characterized by the parameters $(\rho, \sigma_\eta^2, \sigma_\epsilon^2)$ which we set to $(0.9695, 0.0384, 0.0522)$, following Krueger, Mitman, and Perri (2016).[9]

We use the Rouwenhorst procedure to discretize the persistent part of the process into an nine-state Markov chain, and we discretize the transitory shock into three states.

We choose the deterministic age profile of earnings estimated for high-school graduates by Cocco, Gomes, and Maenhout (2005), and renormalize it such that the average labor productivity is unity.

Remaining externally calibrated parameters. The remaining parameters that are set externally are listed in Table 3. The bend points and the contribution and benefit base are reported in U.S. dollars to facilitate the interpretation. Details for transforming them into values used in the model are found in the appendix, section B.1.1.

4.3 Health and survival process

We use the processes for health transitions and survival probabilities described in section 2.3 and section 2.5 for nonblack men. Agents enter the model at the age of 20, but the health and survival processes we estimated starts at the age of 50. Therefore, we make the assumption that everyone is born in the excellent health state. Thereafter, we use the health transition matrix for the age of 50 to roll forward the population, assuming certain survival. At the age of 50, agents start facing a positive probability of death according to our estimated process. The resulting cohort sizes and distribution of health states are shown in Figure 18 in the appendix section B.4.

While the model is solved with five health states, in what follows we report results only for the worst, middle and best health states to reduce visual clutter.

[9] Krueger, Mitman, and Perri (2016) remove the age effect before estimating this process and hence, we can use this stochastic process on top of the age-dependent profile.

Parameter	Description	Value	Source
Production technology parameters			
α_k	Capital share	36%	Krueger, Mitman, and Perri (2016)
δ_k	Depreciation rate	9.6%	Krueger, Mitman, and Perri (2016)
Social security			
ρ_1	Replacement rate bracket 1	90%	2013 SS rules
ρ_2	Replacement rate bracket 2	32%	2013 SS rules
ρ_3	Replacement rate bracket 3	15%	2013 SS rules
$b_1^{\$}$	Bendpoint 1 (in USD)	9,492	2013 SS rules
$b_2^{\$}$	Bendpoint 2 (in USD)	57,216	2013 SS rules
$e_{max}^{\$}$	CBB (in USD)	113,700	2013 SS rules
Government budget			
g	Gov. spending (share of GDP)	6%	Brinca et al. 2016
τ	Tax progressivity	0.137	Brinca et al. 2016

Table 3: Calibrated parameters

5 Results

We solve the model under three distinct assumptions about survival expectations:

1. *No survival heterogeneity (NSH):* Each cohort faces the same survival expectations.
2. *Objective survival heterogeneity (OSH):* Each cohort has heterogeneous objective survival expectations that depend on health.
3. *Subjective survival heterogeneity (SSH):* Each cohort has heterogeneous subjective survival expectations that depend on health.

The first scenario serves as our baseline: this is the standard model where all agents face the same survival risk and hence, all agents have the same effective discount factor conditional on age. For this scenario, we eliminate health heterogeneity and use the average survival rates.

In the second scenario, we use the objective process for health transitions and survival probabilities described in section 2.3. Hence, in this model, people are perfectly informed about their true survival probability conditional on health and age.

In the third scenario, agents believe and act according to the subjective survival process estimated in section 2.5. However, this subjective process does not correspond to the true survival process, which we use when simulating the model.

5.1 Effective discount rates

The effective discount rate (determined by the common discount factor β and the individual survival probability) for the agents in the model is time-varying and depends on the horizon. With a bad health shock, the discount rate immediately rises, since it

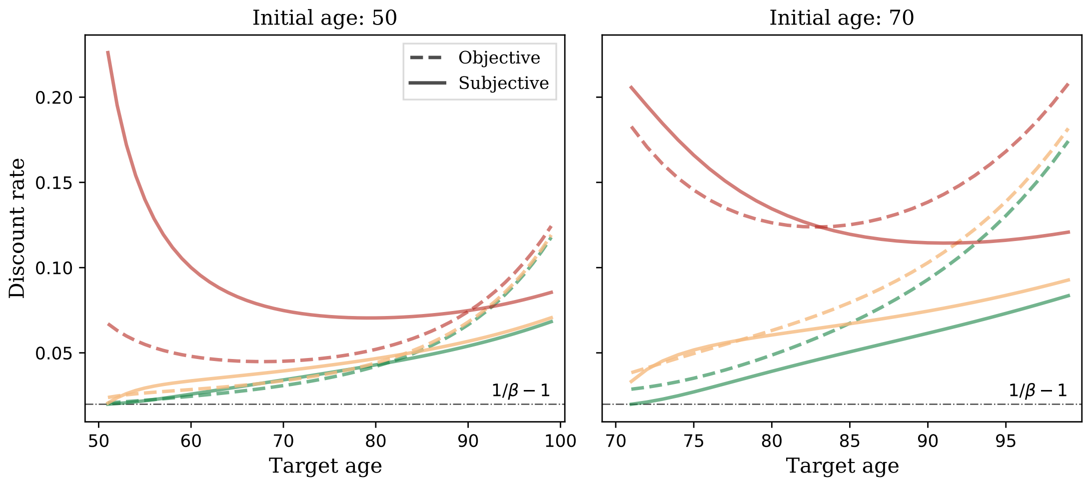

Figure 10: Effective discount rate (geometric average) as a function of current age (50 or 70) and forecast horizon (x-axis). The worst, middle, and best health state are shown, with color indicating the state: green is excellent while red is poor health. The thin black line indicates the discount rate net of any survival effect ($1/\beta - 1$).

implies a shorter expected life span while the opposite happens in the event of a good health shock. Figure 10 plots the geometric average of the expected effective discount rates, given the stochastic health and survival process, for different time horizons. For example, for an individual of age t with health h, the figure plots the mean discount rate ϱ to target age T implicitly defined as

$$\beta^{T-t} \cdot \Pr\left(\text{alive at } T \mid t, h\right) = \left(\frac{1}{1+\varrho}\right)^{T-t}$$

As can be seen, the effective discount rate varies substantially in the population. Using the subjective probabilities, the one-year horizon discount rate for a 50-year-old agent in good health is observationally equivalent to $1/\beta - 1 = 2.01\%$ (since the one-year ahead survival for this agent is perceived to be almost certain), while the one-year horizon discount rate for an equally old agent in bad health is above 20%, more than ten times as large. At longer horizons, the geometric averages of the expected discount rate becomes more similar for the two agents, but the difference is still 7.4 percentage points on a 10-year horizon. For 70-year-old agents in worst vs. best health state, the difference at the 10-year horizon is more than 10 percentage points.

The magnitude of these differences is consistent with the findings by Calvet et al. (2021) who estimate the cross-sectional distribution of time preference rate differences based on Swedish micro data, assuming a common survival probability conditional on age for all agents. Their estimation of the standard deviation of the time preference rate is 6.0 percentage points around a mean of 6.2 percent.

These findings are also consistent with a downward sloping discount rate along the age gradient, as found by Kureishi et al. (2021). Netting out the average objective (statistical) survival probability at all ages, the belief bias in survival probability give rise to a downward sloping residual discount rate, since older people are on average more optimistic about their survival compared to the young (see Figure 3).

5.2 Effect on savings behavior

The main effect from introducing health and survival heterogeneity is, not surprising, that agents in bad health save less than their healthy counterparts as they expect to live a shorter life. The same but opposite effect is present for individuals in good health. This is true regardless of whether agents assess their life expectancy using objective survival probabilities or subjective beliefs. However, the difference in savings between agents in bad and good health increases if we let the agents act according to their subjective beliefs due to the within-cohort steepness bias we documented in section 2.5: individuals in bad health underestimate their survival probability more (or, for some ages, do not overestimate to the same extent) as compared to agents in good health.

These patterns are illustrated in Figure 11. For selected health states, the bars in darker color show the differences in total savings rates for the model with objective heterogeneity compared to the baseline model with no heterogeneity in survival. The lighter color shows the additional effect of adding subjective beliefs. Here, we define *total savings rate* as the fraction of current cash-at-hand, i.e., beginning-of-period assets plus current income, which the individual saves for the next period. The horizontal axis represents the cash-at-hand percentiles by age from the baseline model with no heterogeneity.[10]

As can be seen, for 50-year-olds the difference is most pronounced for individuals in bad health, as indicated by the red bars. With objective survival heterogeneity, individuals in the 10th cash-at-hand percentile in poor health have a savings rate that is 9 percentage points lower than individuals of the same wealth in the model with no heterogeneity. Adding subjective beliefs, the difference is magnified: the total savings rate of an individual in the 10th percentile is now 22 percentage points lower than for an individual with the same wealth with average survival expectations.

The reason for the bigger impact of the subjective belief on the individuals in worst health can be understood by looking at Figure 4(a) for nonblack men. At lower ages, the subjective and the objective survival expectations are similar for individuals in the better

[10]We plot results by cash-at-hand percentiles (as opposed to the cash-at-hand level) since the graphs then show the difference at the wealth levels that matter, i.e., where there is a positive mass of agents in equilibrium. A large difference in policy functions for very high wealth levels for the 80-year-olds is not that informative, since extremely few 80-year-old will be that rich in equilibrium anyhow. We plot standard policy functions for savings in the appendix section section C.1.

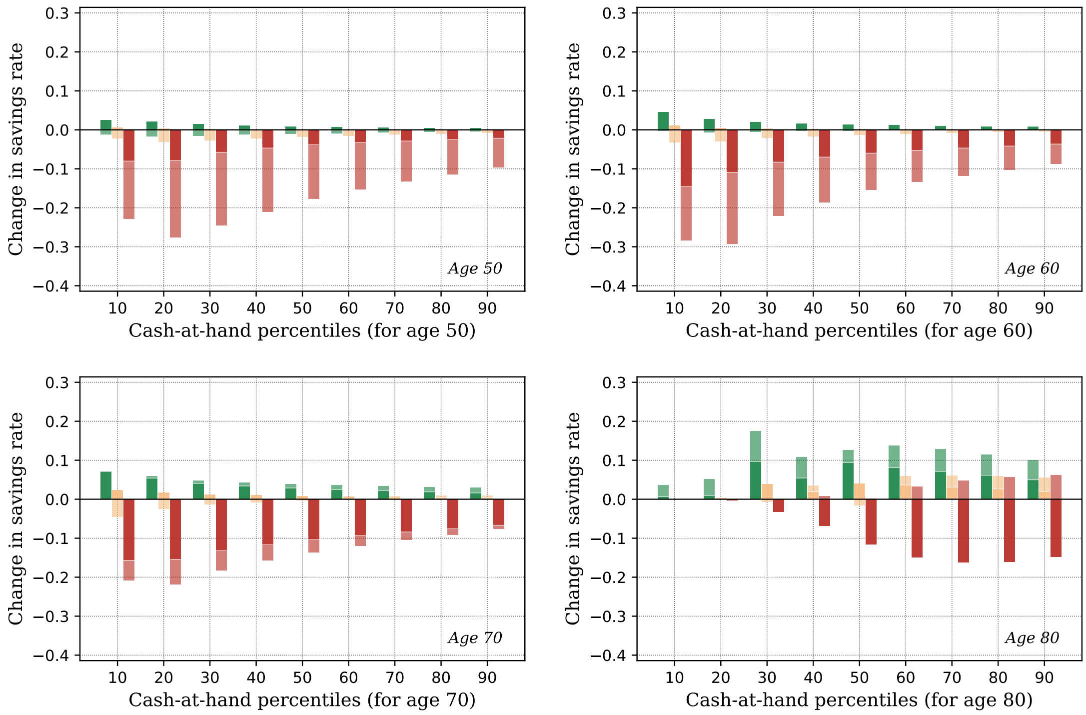

Figure 11: Absolute difference in savings rate (defined as the fraction of current cash-at-hand the agent saves for the next period). The darker color shows the model with objective survival heterogeneity compared to the baseline model with no survival heterogeneity. The lighter color shows the additional effect of adding subjective survival beliefs. The x-axis depicts cash-at-hand percentiles by age (equilibrium values for the baseline model). The color indicates health state: green is excellent while red is poor health.

health states and therefore the additional effect of adding subjective beliefs is small. However, individuals in worse health are severely underestimating their longevity and thus the additional effect on the total savings rate is larger.

The same pattern can be seen in the graphs for 60-year-olds and 70-year-olds, even though for the latter, the additional effect from adding the subjective beliefs is slightly weaker if we continue to focus on the individuals in poor health. The reason for this can be found in Figure 4(b): at the age of 70, the subjective survival beliefs of the individuals in bad health are now closer to the objective probabilities.

Moreover, at the age of 70, the effect of adding subjective beliefs starts showing up for individuals in good health: there is an additional effect of the subjective beliefs on the total savings rate for healthy individuals, as indicated by the lighter green parts of the green bars. As can be seen from the graph of 80-year-olds, at higher ages the general optimism dominates. The addition of subjective beliefs at this age increases the savings rates for individuals both in bad and good health.

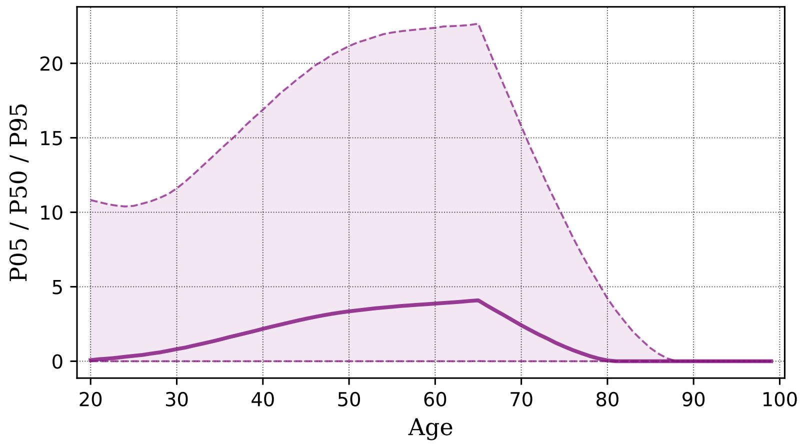

(a) No survival heterogeneity (NSH)

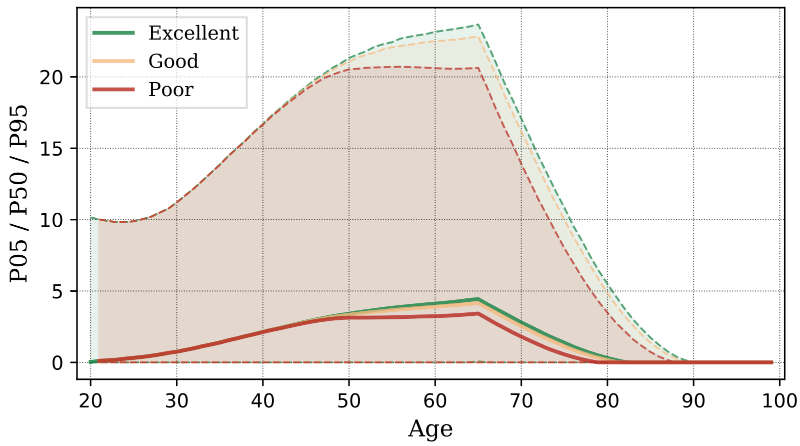

(b) Objective survival heterogeneity (OSH)

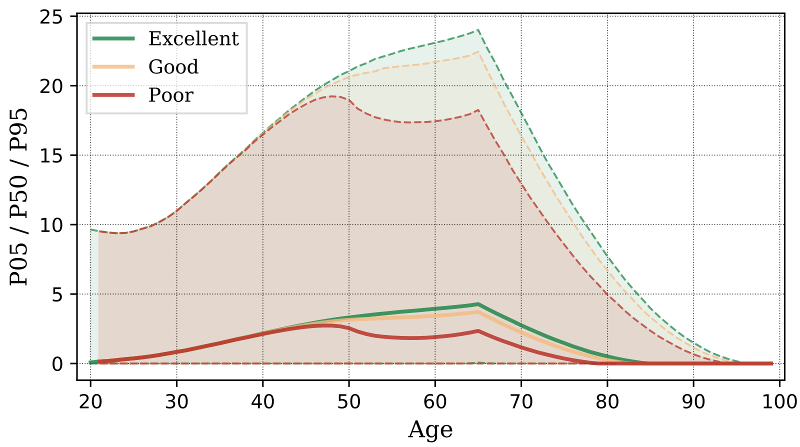

(c) Subjective survival heterogeneity (SSH)

Figure 12: Life-cycle profiles for wealth. The worst, middle, and best health state are shown, with color indicating the state: green is excellent while red is poor health.

	NSH: no heterogeneity	OSH: objective exp.	SSH: subjective exp.
Interest rate (%)	2.37	2.39	2.57
Wealth Gini	0.67	0.67	0.67

Table 4: Resulting aggregate variables from the three scenarios.

5.3 Effect on wealth accumulation

Figure 12 shows the resulting life-cycle profiles for the three different scenarios: no survival heterogeneity (NSH), objective survival heterogeneity (OSH), and subjective survival heterogeneity (SSH). In all graphs, we average out the productivity dimension.

As expected, and as all three graphs show, the life cycle profile for wealth peaks at the age of 64, which is the last year before retirement. When the agents enter retirement, they start drawing down their wealth and the average individual who survives until the age of 90 has drawn down all of his savings (remember that all individuals receive retirement benefits, so they are not risking zero consumption).

The profile for the OSH model illustrates that individuals in bad health have less wealth than individuals in good or excellent health. It should be noted that the mass of individuals in bad health is not static, but consists of individuals who have been in bad health for many periods, as well as individuals who just recently draw a bad health shock. Moreover, the profile for the SSH model shows that the difference in wealth between individuals in good and bad health at the age of 64 increases when the agents act according to their subjective beliefs. Note also that the asset holdings at very high ages increase slightly, since old individuals are, on average, over-optimistic about their lifespan.

In the SSH model, the average 65- to 69-year-old in the best health state has 51% more wealth than the average agent in the worst health state in the same age group. This is approximately one third of the difference we see in the data comparing the best and worst health states at this age.

5.3.1 General equilibrium effects and economy-wide inequality

The general equilibrium effects of introducing survival heterogeneity are small. Table 4 shows two key statistics from the model. The reason why the interest rate increases slightly in the OSH model (with objective survival heterogeneity) compared to the NSH model (with no survival heterogeneity) is the on average lower incentive to save among those in the ages 50 to 70. As Figure 11 shows, the demand for savings is substantially reduced at those ages among the individuals in bad health as compared to the baseline without heterogeneity. This effect is not fully counteracted by the slight increase in demand for savings by those in excellent health. Therefore, the interest rate needs to rise slightly to induce enough savings in the aggregate.

When agents act according to their subjective survival beliefs (SSH model), the interest rate increases slightly more to 2.57%. The reason is that in the age group around 55–65, which is when people save the most, agents on average have a downward bias in expected longevity. Hence, their incentives to save are reduced further and consequently, the interest rate increases.

The last thing to note is that the wealth Gini is virtually unchanged between the three models. However, if we look at the within-cohort inequality, there are some differences in wealth holdings. Within the oldest age groups, the Gini falls in the model with subjective survival beliefs. For instance, within the group of 80- to 85-year-olds, the top 1% richest in the SSH model are 70% wealthier than the top 1% in the OSH model. However, there is also a smaller fraction within this age group having virtually zero assets (due to the general overoptimism about survival among the oldest agents), thus the net results is that the Gini is 5pp lower in the model with subjective survival beliefs than in the baseline OSH model.

To conclude, the within-cohort inequality in the very oldest cohorts is affected by survival heterogeneity, as could be expected after seeing the large differences in the policy functions. However, the richest individuals in the economy are found in the age group 60–64, as we saw in the life-cycle profiles. For this age group, the "extra savings" by healthy individuals due to a longer expected life are very small, as we saw in Figure 11 – the main effect in this age group is that unhealthy individuals save less. Therefore, the effect of heterogeneity in survival on the top wealthiest individuals in the economy is very small, and consequently the effect on overall inequality is negligible.

Appendix section C.2 gives a full breakdown of Gini by 5-year age groups.

5.4 Welfare cost of the belief bias

Imagine your friend is more optimistic about her survival probabilities than you are, since you happen to know her true survival probabilities. Otherwise you share the same preferences and information. In your opinion, what is your friend's welfare loss due to her erroneous beliefs and therefore, in your opinion, non-optimal consumption/savings choices? This is the type of welfare cost we calculate in this section. Note that your friend is optimizing her behavior given her beliefs, so without changes to perceived probabilities, in her view there is no scope for improvement (and she would certainly not perceive it as welfare improving if you tried to force her to follow your well-meant advice).

Another way to phrase the question is the following: if a person with objective beliefs about survival were forced to use the consumption policy functions of an individual optimizing under the subjective beliefs, what is the consumption equivalent loss of welfare? Formally, consider the value function of an individual with objective beliefs about survival probabilities,

$$V_o(\boldsymbol{x}_t) = \sum_{\tau=t}^{T} \beta^{\tau-t} \sum_{x_\tau | x_t} \Pr\left(\boldsymbol{x}_\tau \mid \boldsymbol{x}_t \right) u\big(C_o(\boldsymbol{x}_\tau)\big)$$

at state $\boldsymbol{x}_t = (a, p, h, t)$, with $C_o(\bullet)$ denoting the optimal consumption policy function for the objective individual.

Next, consider the expected utility from letting this person adopt the subjective individual's policy function, denoted $C_s(\bullet)$:

$$EU_o\left(\boldsymbol{x}_t | C_s\right) = \sum_{\tau=t}^{T} \beta^{\tau-t} \sum_{x_\tau | x_t} \Pr\left(\boldsymbol{x}_\tau \mid \boldsymbol{x}_t \right) u\big(C_s(\boldsymbol{x}_\tau)\big)$$

By definition, $V_o \geq EU_o$. We compute the value Δ such that

$$EU_o\left(\boldsymbol{x}_t \mid C_s\right) = EU_o\left(\boldsymbol{x}_t \mid C_o, \Delta\right) = \sum_{\tau=t}^{T} \beta^{\tau-t} \sum_{x_\tau | x_t} \Pr\left(\boldsymbol{x}_\tau \mid \boldsymbol{x}_t \right) u\big(C_o(\boldsymbol{x}_\tau)(1+\Delta)\big)$$

for each (a, p, h, t). Thus, Δ is the consumption equivalent variation that makes an agent with objective beliefs indifferent to being "forced" to use the subjective belief policy functions.

The resulting $\Delta(a, p, h, t)$ varies from approximately zero to -11% for some groups. The largest costs are found among very rich and very old agents, who constitute a small subset of the population (the vast majority of old agents are relatively poor). In the appendix, section C.3, we show an overview of the full distribution.

In Figure 13, we average out the productivity and asset dimensions and show the consumption equivalent variation conditional on age and health state. As can be seen, the cost of the belief bias is very small for young agents. Approaching the age of 50, the cost increases, especially for agents in bad health: these agents are pessimistic about their survival chances and therefore save less for their future than an agent with objective beliefs. The average cost of the belief bias for agents in bad health at the age of 50 is equivalent to 1.5% of annual consumption. In older ages, the cost of the subjective beliefs is driven by the opposite effect: here agents are too optimistic about their survival, which leads to oversaving and too little current consumption, compared to what an agent with objective beliefs would have chosen, given the same wealth, retirement income, health state, and age.

Going back to the example of your friend with overly optimistic beliefs about her survival, it is moreover not clear if she would be better off from learning about the true probabilities. With a broader view of welfare in line with what is suggested by Brunnermeier and Parker (2005), there are two effects of being too optimistic. Being overly optimistic about survival has the advantage of higher anticipation utility (assuming an

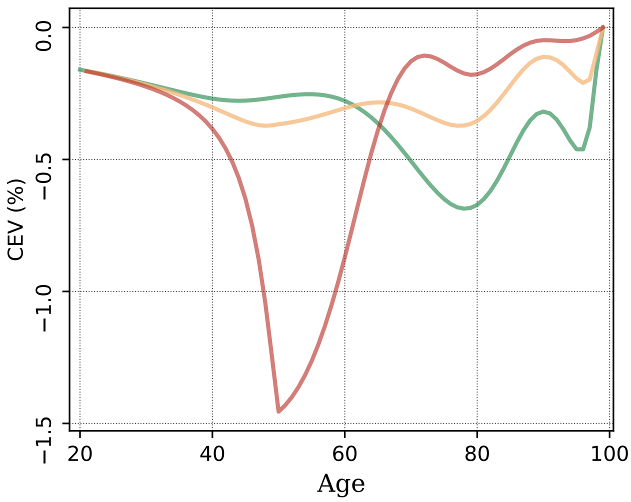

Figure 13: Weighted average consumption equivalent variation (CEV, Δ) of the welfare loss from using subjective belief consumption functions, conditional on age/health. See text for full description. The worst, middle, and best health state are shown, with color indicating the state: green is excellent while red is poor health.

intrinsic value of life that is high enough so that living is preferred to being dead) and the disadvantage of making decisions that turn out to be sub-optimal ex post. With this view of welfare, the cost of the belief bias for the 50-year-olds in bad health is more problematic than the cost for the 80- and 90-year-olds. In the latter case, the agents at least get a higher anticipation utility from their overly optimistic expectations (assuming a sufficiently high intrinsic value of life so that survival is preferred). The 50-year-olds in bad health have a bias that is unambiguously welfare reducing: it leads to both non-optimal consumption/savings allocation *and* lower anticipatory utility (the latter not being included in our quantification above).

5.5 Determinants of savings not included in our benchmark model

There are many drivers of savings in old age that could vary across health but are not included in our model: medical shocks, the existence of (employer-tied) health or life insurance, human capital investment, endogenous retirement decisions, an income-health gradient, private pensions, housing and other illiquid assets, to name a few (see for instance De Nardi, French, and Jones (2010), Capatina (2015) or De Nardi, Pashchenko, and Porapakkarm (2017) for studies taking a broader perspective including several channels). We acknowledge that it is beyond the scope of this paper to fully characterize the consumption/savings decision over the lifecycle. Rather, our aim is to quantify the channel of heterogeneity in (subjective) life expectancy and its effect on savings in isolation.

To capture the asset profile over the lifecycle it is tempting to introduce a bequest motive of the warm-glow De Nardi (2004) type, the most commonly used formulation in the macroeconomic literature. The effect of combining survival heterogeneity with such a bequest motive is however counter-intuitive and perhaps also unexpected: agents in poor health save *more* than their otherwise identical healthy counterparts. The reason is as follows: since agents in poor health are more likely to die soon, they put an increased weight on bequest utility, thus raising their incentives to save. Hence, there are two effects from lower life expectancy that work in opposite directions: a shorter expected life span makes agents want to save less for their own consumption, but a stronger bequest motive induces them to save more. The net effect varies depending on the calibration of bequest parameters, but the second mechanism is always present with a bequest formulation of this type: a shorter life span makes agents want to save more to leave bequests, and if a rich individual receives a bad health shock, there is an incentive to save more due to the bequest motive. This creates a health-wealth gradient and a change in savings behavior in the event of bad health shocks that are challenging to square with the data. One would ideally develop an extension of the warm-glow bequest motive, but this is beyond the scope of this paper.

Appendix section D illustrates the issue further and shows results from a model where we include a standard warm-glow bequest motive.

6 Conclusions

This paper explores how heterogeneity in life expectancy, objective (statistical) as well as subjective, affects savings behavior between healthy and unhealthy people. Using HRS data, we show that there exists a *within-cohort steepness bias* in survival beliefs: individuals in bad health not only have a shorter expected life span, but are also relatively more pessimistic about their survival chances, while individuals in good health and thus with higher survival probability are more optimistic. This systematic bias exacerbates the survival expectancy heterogeneity in the population.

The differences in beliefs about survival translate into time preference heterogeneity and consequently savings behavior. We show that differences in life expectancy can explain one third of the differences in accumulated wealth between healthy and unhealthy people, mostly due to pessimism among unhealthy people.

This paper ties into a strand of current research investigating preference heterogeneity and its importance for individual choices and aggregate outcomes. We provide an intuitively plausible and micro-founded source of heterogeneity: the probability of surviving to future states of the world. Our quantification of this channel shows that life expectancy heterogeneity is important and should be included in the list of potential sources of heterogeneity that we need to consider in our analyses. Investigating the importance of the steepness bias for within-cohort differences in terms of portfolio allocations, demand for financial products, or retirement behavior is left for future research.

References

Attanasio, Orazio P, and Carl Emmerson. 2003. Mortality, health status, and wealth. *Journal of the European Economic Association* 1 (4): 821–850.

Attanasio, Orazio P, and Hilary Williamson Hoynes. 2000. Differential mortality and wealth accumulation. *Journal of Human Resources:* 1–29.

Bissonnette, Luc, Michael D. Hurd, and Pierre-Carl Michaud. 2017. Individual survival curves comparing subjective and observed mortality risks. *Health Economics* 26 (12): e285–e303.

Blundell, Richard, Margherita Borella, Jeanne Commault, and Mariacristina De Nardi. 2020. *Why does consumption fluctuate in old age and how should the government insure it?* Technical report. National Bureau of Economic Research.

Brinca, Pedro, Hans A. Holter, Per Krusell, and Laurence Malafry. 2016. Fiscal multipliers in the 21st century. *Journal of Monetary Economics* 77:53–69.

Brunnermeier, Markus K., and Jonathan A. Parker. 2005. Optimal expectations. *American Economic Review* 95 (4): 1092–1118.

Calvet, Laurent E, John Y Campbell, Francisco Gomes, and Paolo Sodini. 2021. *The cross-section of household preferences.* Technical report. National Bureau of Economic Research.

Capatina, Elena. 2015. Life-cycle effects of health risk. *Journal of Monetary Economics* 74:67–88.

Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout. 2005. Consumption and portfolio choice over the life cycle. *The Review of Financial Studies* 18 (2): 491–533.

Coile, Courtney, and Kevin Milligan. 2009. How household portfolios evolve after retirement: the effect of aging and health shocks. *Review of Income and Wealth* 55 (2): 226–248.

De Nardi, Mariacristina. 2004. Wealth inequality and intergenerational links. *The Review of Economic Studies* 71 (3): 743–768.

De Nardi, Mariacristina, Eric French, and John B Jones. 2009. Life expectancy and old age savings. *The American Economic Review Papers and Proceedings* 99 (2): 110–115.

———. 2010. Why do the elderly save? The role of medical expenses. *Journal of Political Economy* 118 (1): 39–75.

De Nardi, Mariacristina, Svetlana Pashchenko, and Ponpoje Porapakkarm. 2017. *The lifetime costs of bad health.* Technical report. National Bureau of Economic Research.

Deaton, Angus. 2002. Policy implications of the gradient of health and wealth. *Health affairs* 21 (2): 13–30.

Duncan, Greg J, Mary C Daly, Peggy McDonough, and David R Williams. 2002. Optimal indicators of socioeconomic status for health research. *American Journal of Public Health* 92 (7): 1151–1157.

Elder, Todd E. 2013. The predictive validity of subjective mortality expectations: evidence from the Health and Retirement Study. *Demography* 50 (2): 569–589.

Epper, Thomas, Ernst Fehr, Helga Fehr-Duda, Claus Thustrup Kreiner, David Dreyer Lassen, Søren Leth-Petersen, and Gregers Nytoft Rasmussen. 2020. Time discounting and wealth inequality. *American Economic Review* 110 (4): 1177–1205.

Epstein, Larry G. 1988. Risk aversion and asset prices. *Journal of Monetary Economics* 22 (2): 179–192.

Epstein, Larry G., and Stanley E. Zin. 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57 (4): 937–969.

Fernström, Clara. 2019. *Dying and dissaving.* Technical report. Stockholm University.

Foltyn, Richard, and Jonna Olsson. 2019. *Health dynamics and heterogeneous life expectancies.* Technical report. Working paper, Institute for International Economic Studies.

Gale, William G, and John Karl Scholz. 1994. Intergenerational transfers and the accumulation of wealth. *Journal of Economic Perspectives* 8 (4): 145–160.

Gan, Li, Guan Gong, Michael D. Hurd, and Daniel McFadden. 2015. Subjective mortality risk and bequests. *Journal of Econometrics* 188 (2): 514–525.

Gan, Li, Michael D. Hurd, and Daniel McFadden. 2007. *12. Individual subjective survival curves.* University of Chicago Press.

Grevenbrock, Nils, Max Groneck, Alexander Ludwig, and Alexander Zimper. 2021. Cognition, optimism, and the formation of age-dependent survival beliefs. *International Economic Review* 62 (2): 887–918.

Groneck, Max, Alexander Ludwig, and Alexander Zimper. 2016. A life-cycle model with ambiguous survival beliefs. *Journal of Economic Theory* 162:137–180.

Hajat, Anjum, Jay S. Kaufman, Kathryn M. Rose, Arjumand Siddiqi, and James C. Thomas. 2010. Long-term effects of wealth on mortality and self-rated health status. *American Journal of Epidemiology* 173 (2): 192–200.

Hamermesh, Daniel S. 1985. Expectations, life expectancy, and economic behavior. *The Quarterly Journal of Economics* 100 (2): 389–408.

Health and Retirement Study. 2018. *RAND HRS longitudinal file 2018 (V1) public use dataset.* Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, MI.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2017. Optimal tax progressivity: an analytical framework. *The Quarterly Journal of Economics* 132 (4): 1693–1754.

Heimer, Rawley Z., Kristian Ove R. Myrseth, and Raphael S. Schoenle. 2019. Yolo: mortality beliefs and household finance puzzles. *Journal of Finance* 74 (6): 2957–2996.

Hendricks, Lutz. 2007. How important is discount rate heterogeneity for wealth inequality? *Journal of Economic Dynamics and Control* 31 (9): 3042–3068.

Hurd, Michael D., and Kathleen McGarry. 2002. The predictive validity of subjective probabilities of survival. *Economic Journal* 112 (482): 966–985.

Hurd, Michael D., James P. Smith, and Julie M. Zissimopoulos. 2004. The effects of subjective survival on retirement and social security claiming. *Journal of Applied Econometrics* 19 (6): 761–775.

Kaplan, Greg, and Giovanni Violante. 2014. A model of the consumption response to fiscal stimulus payments. *Econometrica* 82 (4): 1199–1239.

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

Krueger, D., K. Mitman, and F. Perri. 2016. Chapter 11 - Macroeconomics and household heterogeneity, edited by John B. Taylor and Harald Uhlig, 2:843–921. Handbook of Macroeconomics. Elsevier.

Krusell, Per, and Anthony A. Smith Jr. 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106 (5): 867–896.

Kureishi, Wataru, Hannah Paule-Paludkiewicz, Hitoshi Tsujiyama, and Midori Wakabayashi. 2021. Time preferences over the life cycle and household saving puzzles. *CESifo Working Paper.*

Kvaerner, Jens. 2020. Intergenerational altruism: estimates based on news about expected mortality. *Available at SSRN 2985465.*

Lee, Jinkook, and Hyungsoo Kim. 2008. A longitudinal analysis of the impact of health shocks on the wealth of elders. *Journal of Population Economics* 21 (1): 217–230.

Ludwig, Alexander, and Alexander Zimper. 2013. A parsimonious model of subjective life expectancy. *Theory and Decision* 75 (4): 519–541.

Margaris, Panagiotis, and Johanna Wallenius. 2020. Can wealth buy health? A model of pecuniary and non-pecuniary investments in health. *CEPR Discussion Paper No. DP14796.*

Pijoan-Mas, Josep, and José-Víctor Ríos-Rull. 2014. Heterogeneity in expected longevities. *Demography* 51 (6): 2075–2102.

Poterba, James M., Steven F. Venti, and David A. Wise. 2017. The asset cost of poor health. *The Journal of the Economics of Ageing* 9:172–184.

Quadrini, Vincenzo, and José-Víctor Ríos-Rull. 2015. Inequality in macroeconomics. In *Handbook of income distribution*, 2:1229–1302. Elsevier.

Rozsypal, Filip, and Kathrin Schlafmann. 2017. Overpersistence bias in individual income expectations and its aggregate implications. *CEPR Discussion Paper No. DP12028.*

Smith, James P. 1999. Healthy bodies and thick wallets: the dual relation between health and economic status. *Journal of Economic Perspectives* 13 (2): 145–166.

Smith, V. Kerry, Donald H. Taylor, and Frank A. Sloan. 2001. Longevity expectations and death: Can people predict their own demise? *The American Economic Review* 91 (4): 1126–1134.

Weil, Philippe. 1990. Nonexpected utility in macroeconomics. *The Quarterly Journal of Economics* 105 (1): 29–42.

A Appendix: Data details

A.1 Data used for the analysis

The RAND HRS Longitudinal File 2018 (V1) includes 14 waves administered over the years 1992–2018. The first cohort included in the survey was between 51 and 61 years old in 1992, and thereafter new (older and younger) cohorts have been added, as illustrated in Figure 14. Figure 15 shows the number of respondents with positive sampling weights by wave and cohort. At the time of this writing, the sampling weights for wave 14 are not yet publicly available.

Figure 16 shows the fraction of respondents in each wave who are marked as non-respondents in all subsequent waves but do not have a death date on record. For example, in wave 11 (administered in 2012), approximately four percent of participants did not respond to any of the later waves 12–14. Since no death date is recorded for these individuals, we cannot use these observations to estimate survival probabilities. However, since death dates are sometimes recorded with considerable lag, we suspect that some of these individuals are already deceased, but their death dates will be updated only in future waves.

So as to not bias our survival probability estimates, we opt to drop the last two waves from the estimation sample, since these exhibit unusually high non-response rates compared to the historical averages.

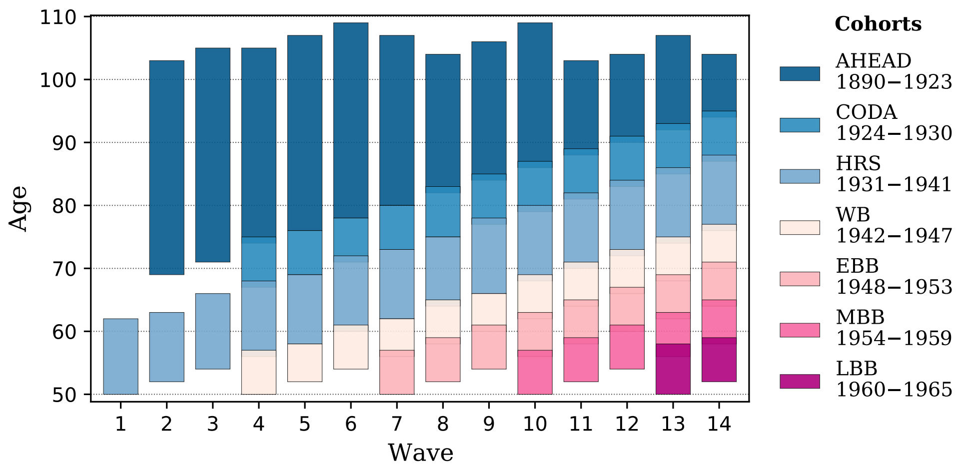

Figure 14: Longitudinal survey design of the HRS. The y-axis shows respondents' age by cohort and wave, ignoring spouses who are not age eligible. The legend lists all birth cohorts included in the HRS (using their "official" names) as well as their birth years.

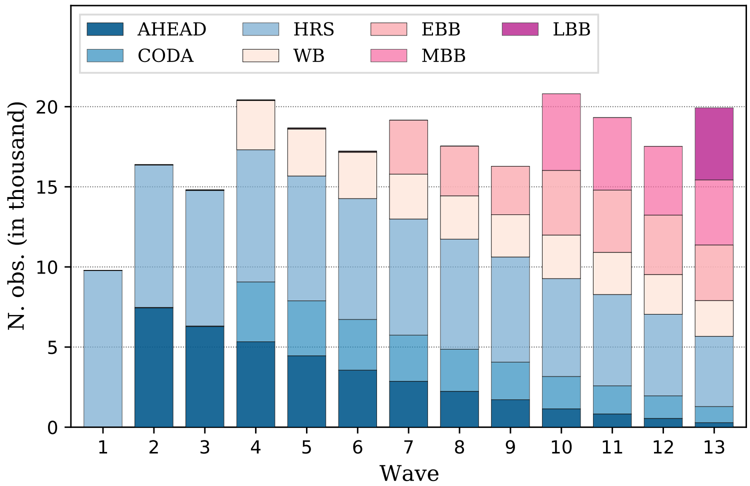

Figure 15: Number of observations by wave and cohort. Only observations with positive weight are included. No sampling weights are currently available for wave 14.

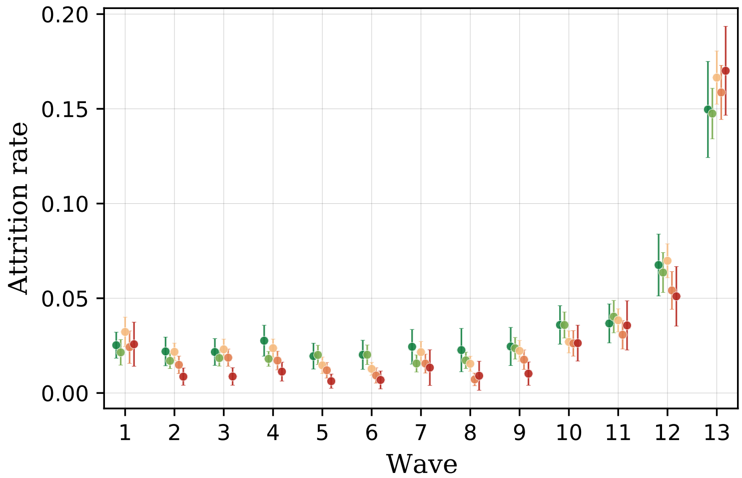

Figure 16: HRS attrition by wave: fraction of participants who do not respond to any of the subsequent survey waves. Color indicates the health state: green is excellent while red is poor health. Error bars indicate 95% confidence intervals based on clustered standard errors (clusters are HRS PSUs).

A.2 Expectation errors in survival probabilities: additional graphs

Figure 17 shows the elicited beliefs about survival vs. estimated objective (statistical) survival probabilities for target age 95. Target ages 75 and 85 are shown in Figure 4 in the main text.

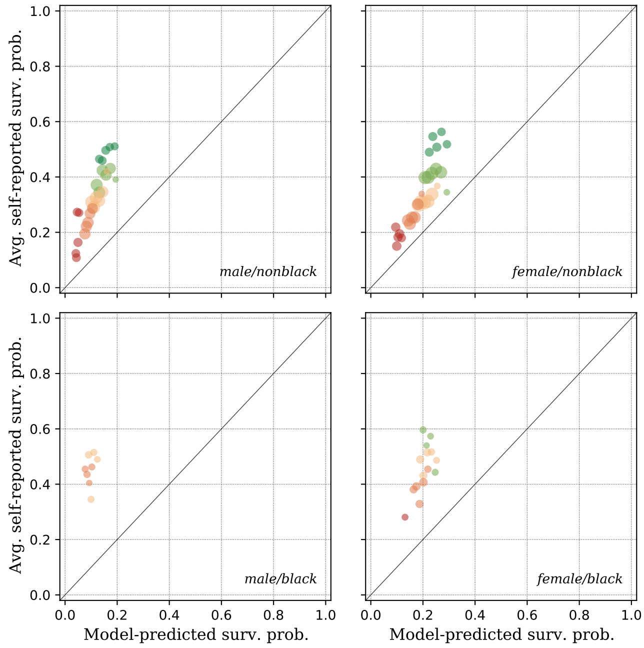

Figure 17: Elicited beliefs about survival vs. estimated objective (statistical) survival probabilities for *target age 95*. Each bubble represents the average for a gender/race/age/health group. The x-axis shows the model-predicted (objective) survival probability, the y-axis the average self-reported survival probability for that group. The color of the bubble indicates health state, with red being poor health and green being excellent health. The size indicates the number of observations in each age/health cell. We exclude cells with less than 20 observations.

B Appendix: Model details

B.1 Social Security system

B.1.1 Retirement benefits

First, consider the following stylized version of the actual retirement income formula used in the U.S. social security system, where $\bar{e}$ is an (annualized) measure of historical average monthly earnings, $b_1^{\$}$ and $b_2^{\$}$ are bend points in USD for some reference year, and $e_{max}^{\$}$ is the contribution and benefit base (CBB), i.e., the maximum earnings subject to payroll taxes. Retirement income $\iota_R^{\$}$ measured in USD is approximately given by

$$\iota_R^{\$}(\bar{e}) = \begin{cases} \rho_1 \bar{e} & \text{if } \bar{e} \leq b_1^{\$} \\ \rho_1 b_1^{\$} + \rho_2 \left(\bar{e} - b_1^{\$}\right) & \text{if } b_1^{\$} < \bar{e} \leq b_2^{\$} \\ \rho_1 b_1^{\$} + \rho_2 \left(b_2^{\$} - b_1^{\$}\right) + \rho_3 \left(\min\left\{e_{\max}^{\$}, \bar{e}\right\} - b_2^{\$}\right) & \text{else} \end{cases} \tag{9}$$

where ρ_1, ρ_2 and ρ_3 are decreasing replacement rates applied to earnings ranges bracketed by the bend points $b_i^{\$}$.

In the model, we define retirement income ι_R as the product of the following components:

$$\iota_R(p) = w \times y_R(p) = w \times \omega_{T_R-1} \bar{\epsilon} p_R(p) \tag{10}$$

where we construct the function $p_R(\bullet)$ to mimic (9) but define it in terms of persistent labor productivity p.

To this end, denote by $e_{med}^{\$}$ the median earnings in dollars in the reference year. To express the bend points in terms of the persistent labor productivity, we implicitly define a model bend point $\widetilde{p}_i$ corresponding to $b_i^{\$}$ by the relationship

$$\frac{b_i^{\$}}{e_{med}^{\$}} = \frac{w \times \omega_{T_R-1} \widetilde{p}_i \bar{\epsilon}}{w \times y_{med}} \qquad i \in \{1, 2\}$$

We normalize the dollar bend point $b_i^{\$}$ by dollar median earnings $e_{med}^{\$}$ and its model counterpart by the median earnings in the model. Solving for $\widetilde{p}_i$, we obtain

$$\widetilde{p}_i = \frac{\left(b_i^{\$} / e_{med}^{\$}\right) y_{med}}{\omega_{T_R-1} \bar{\epsilon}}.$$

Analogously, the CBB in terms of persistent labor productivity is

$$\widetilde{p}_{max} = \frac{\left(e_{max}^{\$} / e_{med}^{\$}\right) y_{med}}{\omega_{T_R-1} \bar{\epsilon}}.$$

By factoring out the common term $\omega_{T_R-1}\bar{\epsilon}w$ that is independent of a retired individual's idiosyncratic state vector, we can write the replacement formula purely in terms of the permanent labor state as follows:

$$p_R(p) = \begin{cases} \rho_1 p & \text{if } p \leq \widetilde{p}_1 \\ \rho_1 \widetilde{p}_1 + \rho_2 \left(p - \widetilde{p}_1\right) & \text{if } \widetilde{p}_1 < p \leq \widetilde{p}_2 \\ \rho_1 \widetilde{p}_1 + \rho_2 \left(\widetilde{p}_2 - \widetilde{p}_1\right) + \rho_3 \Big(\min\Big\{\widetilde{p}_{max}, p\Big\} - \widetilde{p}_2\Big) & \text{else} \end{cases}$$

Retirement income can then be computed as a function of p and the equilibrium wage w according to (10).

B.1.2 Social security budget balance

Government expenditures on retirement benefits in each period are given by

$$\begin{aligned} G_{ss} &= \sum_{t=T_R}^{N_t} \sum_p \mu_t(t) \mu_p(p) \iota_R(p) \\ &= w\omega_{T_R-1}\bar{\epsilon} \cdot \bar{p}_R \Pi_R \end{aligned} \tag{11}$$

which is a weighted sum over the retirement incomes received by all retired cohorts, with weights μ_t and μ_p denoting the PMFs of the ergodic distribution over age and persistent labor productivity, respectively. We denote the mass of retired individuals by

$$\Pi_R = \sum_{t=T_R}^{N_t} \mu_t(t)$$

and the average permanent component of retirement income as

$$\bar{p}_R = \sum_p \mu_p(p) p_R(p).$$

The payroll taxes raised each period are

$$\bar{T}_{ss} = \sum_{t=1}^{T_R-1} \sum_p \sum_\epsilon \mu_t(t) \mu_p(p) \mu_\epsilon(\epsilon) T_{ss}\left(y_e\right) w \tag{12}$$

where $\mu_\epsilon(\epsilon)$ is the ergodic distribution over transitory labor shocks. The payroll tax function is defined as

$$T_{ss}\left(y\right) = \tau_{ss} \times \min\{y_{max}, y\}$$

where y_{max} are maximum taxable earnings, which we obtain from the dollar values using

$$y_{max} = \left(e^{\$}_{max} / e^{\$}_{med}\right) y_{med}.$$

To balance the social security system, we need to find τ_{ss} such that $G_{ss} = \overline{T}_{ss}$. Equating $G_{ss} = \overline{T}_{ss}$ implies that

$$\tau_{ss} = \frac{\omega_{T_R-1}\overline{p}_R\overline{\epsilon}\Pi_R}{\sum_{t=1}^{T_R-1}\sum_p\sum_\epsilon \mu_t(t)\mu_p(p)\mu_\epsilon(\epsilon)\min\{y_{max}, \omega_t p\epsilon\}}. \tag{13}$$

B.2 Government budget

In this section, we derive an expression for the total amount of income taxes raised by the government. Before proceeding, we state the following useful definitions: We denote by $\overline{p}$ the average persistent labor shock,

$$\overline{p} = \sum_p \mu_p(p)p \tag{14}$$

and by Π_W the size of the labor force,

$$\Pi_W = \sum_{t=1}^{T_R-1} \mu_t(t) = 1 - \Pi_R$$

Additionally, average labor productivity can be defined as

$$\overline{y} = \Pi_W^{-1}\left[\sum_t^{T_R-1}\sum_p\sum_\epsilon \mu_t(t)\mu_p(p)\mu_\epsilon(\epsilon)\omega_t p\epsilon\right] = \Pi_W^{-1}\left[\overline{p}\cdot\overline{\epsilon}\sum_h^{T_R-1}\mu_t(t)\omega_t\right].$$

Now, consider the aggregate tax revenues raised from working individuals, which are given by

$$\begin{aligned} T_W = &\sum_{t=1}^{T_R-1}\sum_p\sum_\epsilon \mu_t(t)\mu_p(p)\mu_\epsilon(\epsilon)\Big[\big(\omega_t p\epsilon - T_{ss}\left(\omega_h p\epsilon\right)\big)w \\ &\qquad - \lambda\Big(\big(\omega_t p\epsilon - T_{ss}\left(\omega_t p\epsilon\right)\big)w\Big)^{1-\tau}\Big] \\ = &\Big[w\Pi_W\overline{y} - \lambda w^{1-\tau}\overline{y}_{-ss,\tau}\Big] - \overline{T}_{ss} \end{aligned}$$

where we define

$$\overline{y}_{-ss,\tau} = \sum_{t=1}^{T_R-1}\sum_p\sum_\epsilon \mu_t(t)\mu_p(p)\mu_\epsilon(\epsilon)\Big(\omega_t p\epsilon - T_{ss}\left(\omega_t p\epsilon\right)\Big)^{1-\tau}$$

to simplify the notation.

Income taxes raised from retired individuals amount to

$$
\begin{aligned}
T_R &= \sum_{t=T_R}^{N_t} \sum_p \mu_t(t)\mu_p(p) \left[y_R(p)w - \lambda \Big(y_R(p)w \Big)^{1-\tau} \right] \\
&= \sum_{t=T_R}^{N_t} \sum_p \mu_t(t)\mu_p(p) \left[\omega_{T_R-1} p_R(p)\overline{\epsilon} w - \lambda \Big(\omega_{T_R-1} p_R(p)\overline{\epsilon} w \Big)^{1-\tau} \right] \\
&= \Pi_R \left[w\omega_{T_R-1}\overline{p}_R\overline{\epsilon} - \lambda w^{1-\tau}\omega_{T_R-1}^{1-\tau}\overline{p}_{R,\tau}\overline{\epsilon}^{1-\tau} \right] \\
&= \overline{T}_{ss} - \lambda\Pi_R w^{1-\tau}\omega_{T_R-1}^{1-\tau}\overline{p}_{R,\tau}\overline{\epsilon}^{1-\tau}
\end{aligned}
$$

with

$$
\overline{p}_{R,\tau} = \sum_p \mu_p(p) p_R(p)^{1-\tau}
$$

Thus, the total revenue from income taxes is

$$
\begin{aligned}
T_{inc} &= T_W + T_R \\
&= w\Pi_W\overline{y} - \lambda w^{1-\tau} \left[\overline{y}_{-ss,\tau} + \Pi_R\omega_{T_R-1}^{1-\tau}\overline{p}_{R,\tau}\overline{\epsilon}^{1-\tau} \right]
\end{aligned}
$$

B.3 Equilibrium definition

A recursive competitive equilibrium is given by a set of prices $\{R, w\}$, tax rates $\{\tau_{ss}, \lambda\}$, decision rules $C(a, p, h, t)$ (for consumption) and $B(a, p, h, t)$ (for savings), and a stationary distribution Γ such that:

1. The decision rules solve the agents' problem for all (a, p, h, t).
2. Factor prices are given by:

$$
r = F_1(K, L) \qquad \text{and} \qquad w = F_2(K, L)
$$

3. τ_{ss} and λ are set so that both the social security system and the general government budget balance.
4. Capital and labor markets clear:

$$
K' = \int B(a, p, h, t)\, \mathrm{d}\Gamma \qquad \text{and} \qquad L = \sum_{t=1}^{T_R-1} \sum_p \sum_\epsilon \mu_t(t)\mu_p(p)\mu_\epsilon(\epsilon)\omega_t p\epsilon
$$

where $\mu_t(t)$, $\mu_p(p)$ and $\mu_\epsilon(\epsilon)$ are the ergodic distributions over age, the persistent and the transitory labor shocks, respectively.

5. The distribution Γ is stationary, i.e., for all relevant Borel sets $\mathcal{B}$

$$
\Gamma(\mathcal{B}, p, h, t) = \sum_{\overline{p}} \sum_{\overline{h}} \sum_{\overline{t}} \pi(p|\overline{p})\pi(h|\overline{h})\pi(t|\overline{t}) \int_{a:B(a,\overline{p},\overline{h},\overline{t})\in\mathcal{B}} \Gamma(da, \overline{p}, \overline{h}, \overline{t})
$$

B.4 Ergodic distribution over age and health

Figure 18 shows the relative cohort sizes and distribution of health states in the model, where for the purposes of illustration we rescale the cohort size of newborns to unity.

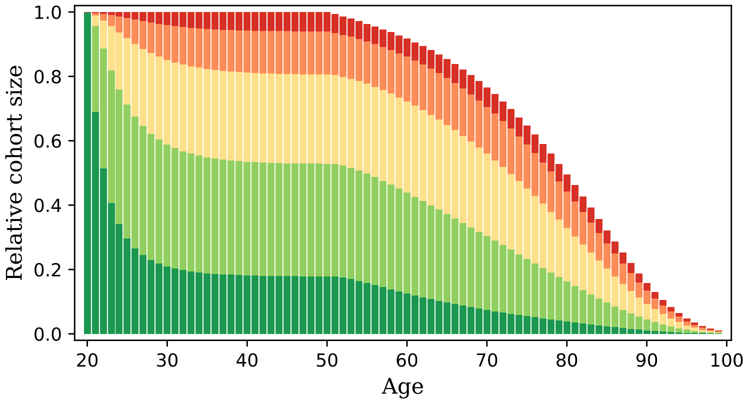

Figure 18: Cohort size and health state distribution. Colors denote health states, with dark green being *excellent* and red being *poor* health.

C Appendix: Additional model results

C.1 Policy functions

Figure 19 shows the savings policy functions for selected ages and productivity states for the model with objective survival heterogeneity (OSH). Panel (a) depicts how much an agent saves as a function of current cash-at-hand, while panel (b) shows how much the agent saves as a fraction of current cash-at-hand. To interpret these figures, it is helpful to know that the 10th percentile of the cash-at-hand distribution among 70-year-olds is approximately 0.4, while the median is around 2.

Figure 20 shows the same information from the model with subjective survival heterogeneity (SSH).

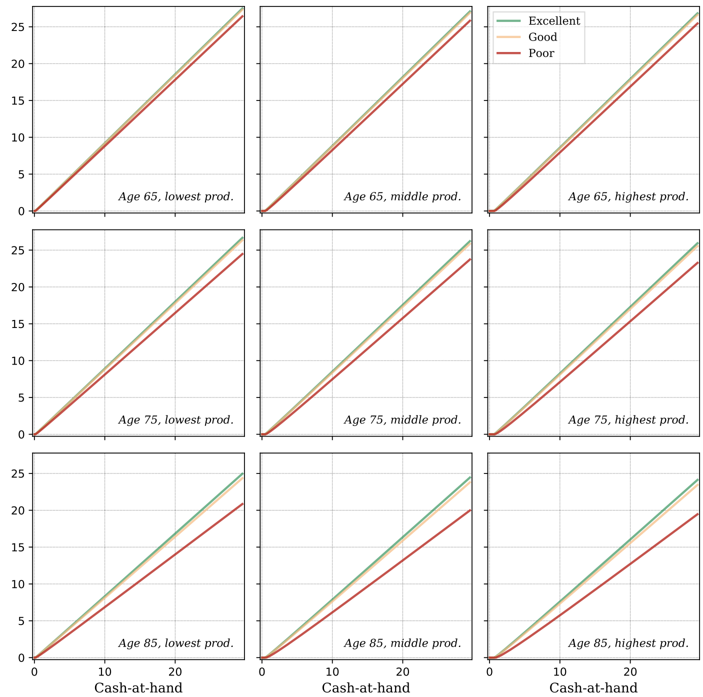

(a) Savings policy functions

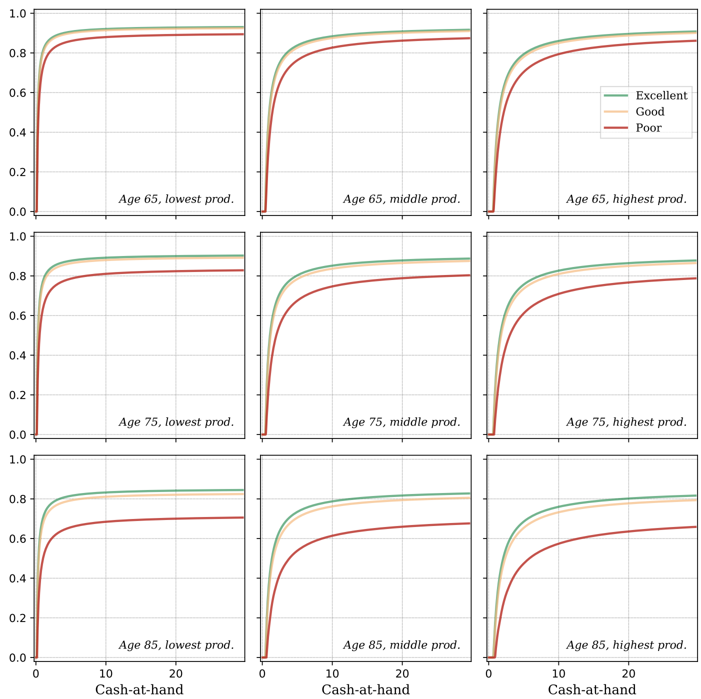

(b) Total savings rate (defined as fraction of total cash-at-hand saved to next period)

Figure 19: Policy functions for the model with *objective* survival heterogeneity (OSH)

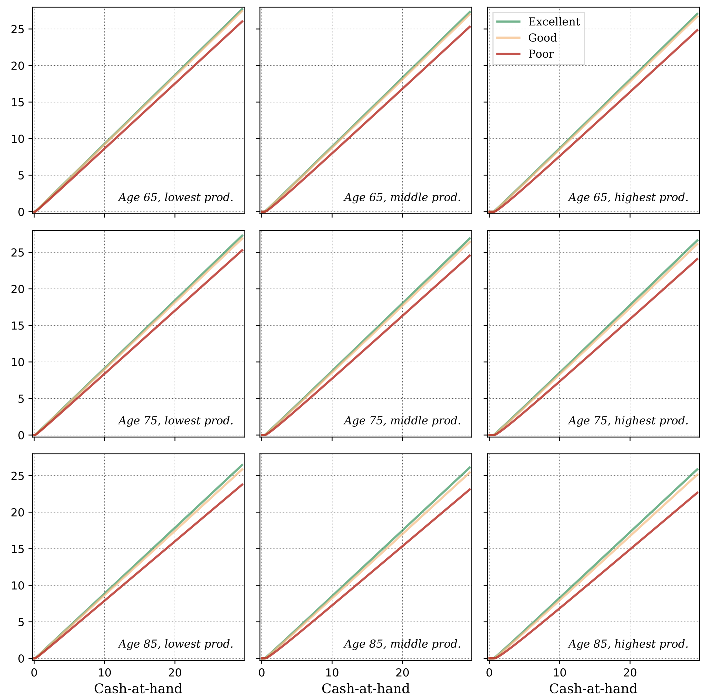

(a) Savings policy functions

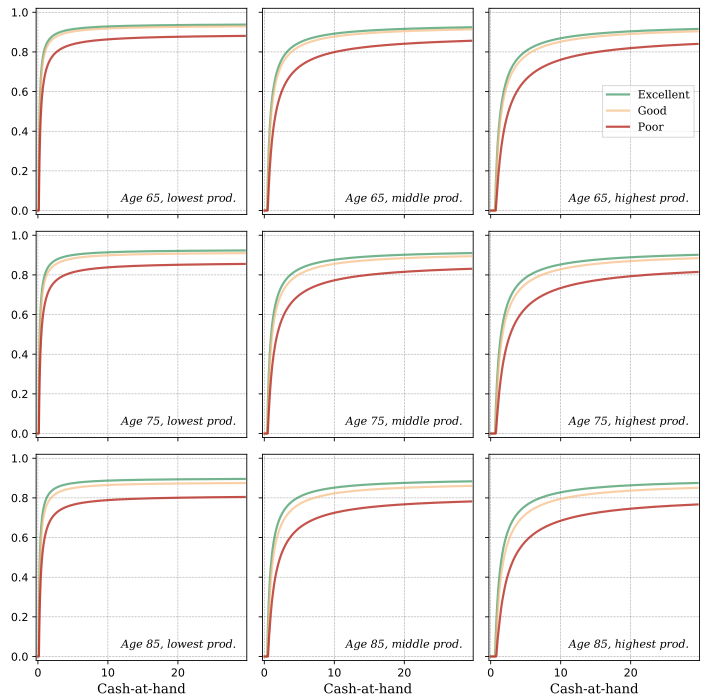

(b) Total savings rate (defined as fraction of total cash-at-hand saved to next period)

Figure 20: Policy functions for the model with *subjective* survival heterogeneity (SSH).

C.2 Wealth inequality within age groups

As noted in the main text, the wealth Gini is virtually unchanged between the three models. However, there are some (small) changes in inequality within cohorts. Figure 21 shows Gini by age group.

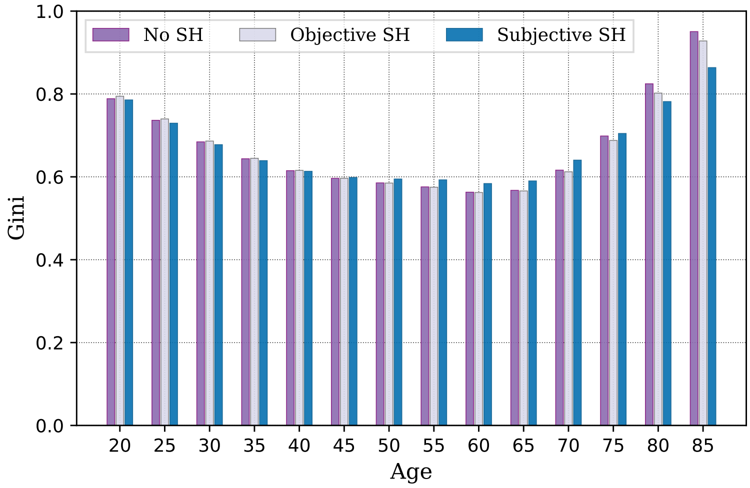

Figure 21: Wealth inequality by 5-year age groups.

C.3 Welfare cost of belief bias

Figure 22 shows the consumption equivalent variation for low, median and high labor productivity, and for poor (P10 of the unconditional CAH distribution), median, and rich (P90) agents, by health state and age.

The largest welfare costs in terms of CEV can be found for agents aged 90 who are very rich. The reason is that with subjective beliefs, these old agents are severely overestimating their survival probabilities. This means that they save "too much" and consequently consume too little today compared to what an agent with objective beliefs would have done. However, relative to the overall population, the fraction of old agents with high levels of wealth is negligible.

The other noteworthy group are individual who are about 50 years old and in bad health. These agents are overly pessimistic and underestimate their chances of survival. Thus, compared to an agent with objective beliefs, they are undersaving and consuming too much today at the expense of too little consumption in the future.

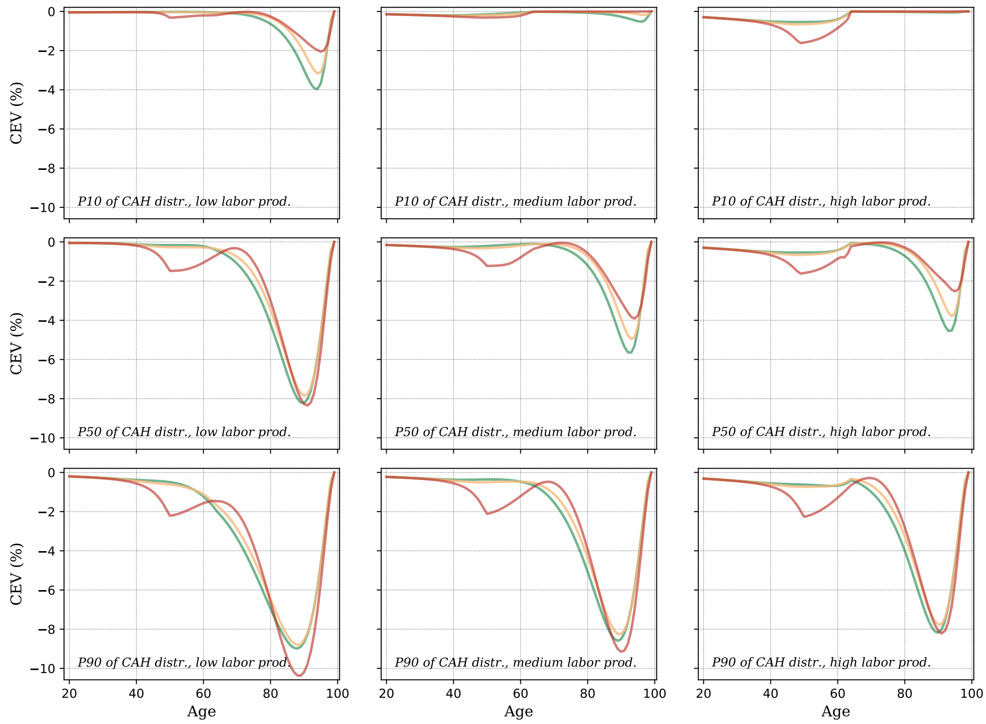

Figure 22: CEV for selected cash-at-hand levels and productivity states. The worst, middle, and best health state are shown, with color indicating the state: green is excellent while red is poor health..

C.4 Alternative calibration with Epstein-Zin-Weil preferences

In the model described in the main text, individuals hold few assets in old age, contrary to what we observe in the data. This is to be expected in a model with Social Security that guarantees non-zero consumption late in life and does not feature a bequest motive.

To alleviate this problem, we explored an alternative setup featuring Epstein-Zin-Weil preferences (Epstein 1988; Epstein and Zin 1989; Weil 1990) with a risk aversion of 4 and an elasticity of intertemporal substitution (EIS) of 1.5, following Kaplan and Violante (2014). Moreover, we introduce additional uncertainty in retirement to induce individuals to increase precautionary savings. Agents now face the same transitory income risk as during working age (but no persistent income shocks), in line with evidence in Blundell et al. (2020) who show that income is uncertain even after retirement. Furthermore, agents (incorrectly) perceive that there is a 5% chance that the Social Security system will partially default and their retirement benefits will permanently be reduced by 50%.

This alternative setup increases savings in old age only slightly, as shown in Figure 23

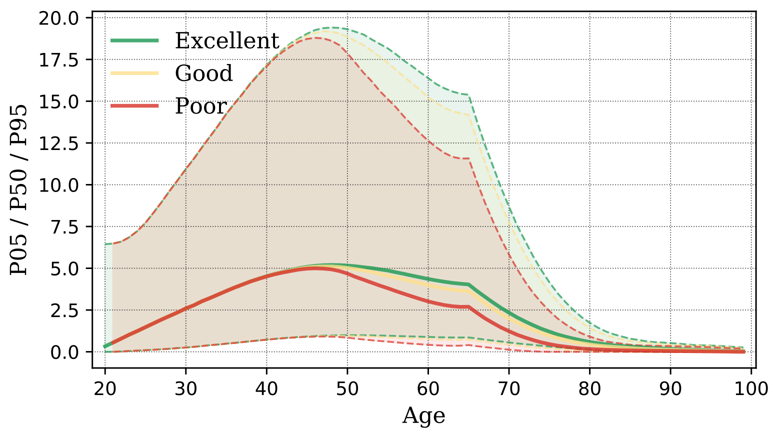

(a) Objective survival heterogeneity (OSH)

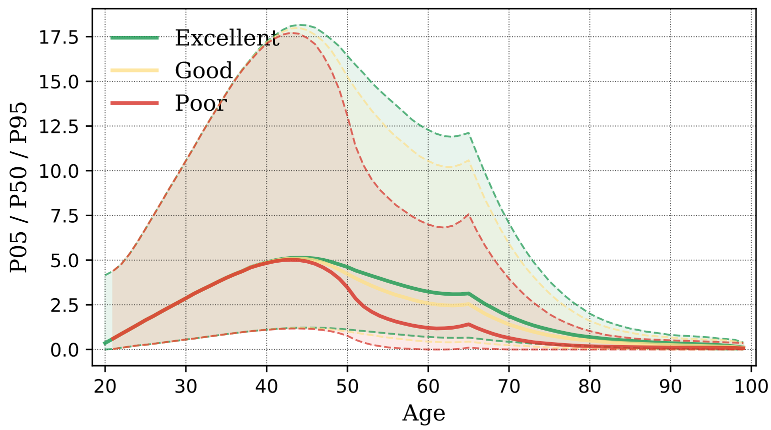

(b) Subjective survival heterogeneity (SSH)

Figure 23: Resulting life-cycle profiles for wealth, model with EZW preferences.

for both objective and subjective beliefs. In the model with subjective beliefs, the average 65- to 69-year-old in the best health state now holds 115% more wealth than the average agent in the worst health state in the same age group.

D Appendix: Model with a bequest motive

D.1 Building intuition in a two-period model

To build intuition, we start with a simple two-period model with a bequest motive where the only uncertainty is survival risk. The survival probability between the first and second periods is denoted by π. After the second period, the agent dies with certainty. We assume that the agent has some initial assets a_0, the discount factor is $\beta < 1$, and the gross interest rate is R. The agent solves

$$\max_{c_0, c_1} \left\{ u(c_0) + \beta \left(\pi \Big[u(c_1) + \beta V_b(a_2) \Big] + (1-\pi) V_b(a_1) \right) \right\} \tag{15}$$

subject to the constraints

$$\begin{aligned} & c_t > 0 \quad \forall t \\ & a_1 = a_0 - c_0, \quad a_2 = Ra_1 - c_1. \end{aligned}$$

The first-order conditions for this problem are given by

$$\begin{aligned} \frac{\partial u}{\partial c_0} &= \beta^2 \pi R \left. \frac{\partial V_b}{\partial a_2} \right|_{a_2^*} + \beta(1-\pi) \left. \frac{\partial V_b}{\partial a_1} \right|_{a_1^*} = 0 \\ \frac{\partial u}{\partial c_1} &= \beta \left. \frac{\partial V_b}{\partial a_2} \right|_{a_2^*} = 0 \end{aligned}$$

where a_1^* and a_2^* denote the optimal choice of savings in each period. The impact of an increase in survival probability on c_0^*, optimal first period consumption, is ambiguous and depends on how the bequest motive is parametrized.

The utility function u has the usual CRRA functional form,

$$u(c) = \frac{c^{1-\sigma} - 1}{1-\sigma}$$

and we use the usual warm-glow bequest motive as in De Nardi (2004),

$$V_b(a) = \theta_B \frac{(a+\kappa)^{1-\sigma} - 1}{1-\sigma} \tag{16}$$

where θ_B determines the strength of the bequest motive, and κ determines to what extent bequests are a luxury good. For simplicity, we assume $\kappa = 0$ and solve for the agent's problem as given in (15). Some algebra shows that whether an increase in the survival probability makes the agent consume more or less in the first period depends on

$$\theta_B \lesseqgtr \left(\frac{R^{\frac{1-\sigma}{\sigma}}}{1 - R^{\frac{1-\sigma}{\sigma}} \beta^{1/\sigma}} \right)^{\sigma} \equiv \widehat{\theta}_B \tag{17}$$

If $\theta_B < \widehat{\theta}_B$, an increase in the survival probability makes the agent consume less in the current period and save more, given that it is more likely to survive to the next period.

On the other hand, if $\theta_B > \widehat{\theta}_B$, an increase in the survival probability leads to *decreased* savings and more consumption in the first period. There are two mechanisms behind this. First, when the probability of surviving increases, the effective discounting of the next-period bequest utility increases, and hence the incentive to save decreases. We call this the expected-date-of-handover channel. Second, an increase in the survival probability leads to a higher expected interest rate income over the remaining life, and therefore the agent can afford more consumption also in the first period. We call this channel the income effect. If the weight on the bequest motive is high, these two effects dominate the effect of wanting to save for a longer expected life.

In the above example, we assumed $\kappa = 0$. It can be shown that if we assume $R = 1$, the sign of $\frac{\partial c_0^*}{\partial \pi}$ is independent of κ (although it still affects the level of savings). However, if we allow for a positive interest rate, the extent to which the bequest is a luxury good in combination with the level of initial assets a_0 matters. Moreover, we assumed no second-period income in our above example. If we allow for income in the second period, the level of income compared to the initial assets affects the incentives. For a given a_0, the higher the second-period income, the less bequest weight is needed to get $\frac{\partial c_0^*}{\partial \pi} > 0$ – since then the bequest becomes relatively more important as a reason to save.

D.2 Quantitative model with a bequest motive

D.2.1 Model setup

Compared to the model from the main text, we now include a warm-glow bequest motive and estate taxes (which enter the government budget constraint). The remaining building blocks remain unchanged.

Retired agents. A retired individual in state (a, p, h, t) now maximizes utility according to

$$V_R(a, p, h, t) = \max_{c,\, b'} \left\{ u(c) + \beta \left(\pi^s_{th} \mathbf{E} \left[V_R(\boldsymbol{x}') \,\middle|\, h, t \right] + (1 - \pi^s_{th}) V_b(a'_b) \right) \right\}$$

subject to the same constraints as in (6). Bequests are given by

$$a'_b = Rb' - T_b(Rb') \tag{18}$$

and are subject to estate taxes given by the function $T_b(\bullet)$ which we describe below.

In the terminal period with $t = N_t$, the individual solves

$$
\begin{aligned}
V_R\left(a, p, h, t\right) = \max_{c,\, b'} & \left\{ u(c) + \beta V_b\left(a'_b\right) \right\} \\
\text{s.t.} \quad & a \geq c + b' \,, \quad c \geq 0 \,, \quad b' \geq 0 \\
& a'_b = Rb' - T_b\left(Rb'\right)
\end{aligned}
$$

Working-age agents. Individuals of age $t < T_R - 1$ who continue to work in the next period now solve

$$
V_W\left(\boldsymbol{x}\right) = \max_{c,\, b'} \left\{ u(c) + \beta \left(\pi^s_{th} \mathbf{E}\left[V_W\left(\boldsymbol{x}'\right) \,\middle|\, p, h, t \right] + (1 - \pi^s_{th}) V_b\left(a'_b\right) \right) \right\}
$$

subject to the same constraints as in (7) plus (18). In the final period of their working life, i.e., when $t = T_R - 1$, they solve

$$
V_W\left(a, p, h, t\right) = \max_{c,\, b'} \left\{ u(c) + \beta \left(\pi^s_{th} \mathbf{E}\left[V_R\left(a', p, h', t+1\right) \,\middle|\, h, t \right] + (1 - \pi^s_{th}) V_b\left(a'_b\right) \right) \right\}
$$

subject to (6) and (18), which is identical to a retired individual's problem.

Estate taxes. The estate tax schedule is defined as

$$
T_b(b) = \begin{cases} 0 & \text{if } b \leq \chi_b \\ \frac{\tau_b}{2}\left[\sin\left(\pi \left[\frac{b - \chi_b}{B} - 1 \right] \right) \frac{B}{\pi} + b - \chi_b \right] & \text{if } \chi_b < b \leq \chi_b + B \\ \tau_b(b - \chi_b - B) + \frac{\tau_b}{2} B & \text{else} \end{cases}
$$

This formulation is effectively a step function: estates valued at less than χ_b are exempt from taxes. For values b in an interval $\chi_b < b \leq \chi_b + B$, the marginal tax rate is increasing, and for $b > \chi_b + B$, the marginal tax rate is τ_b. The tax schedule is twice continuously differentiable, which is required by our solution algorithm.

D.2.2 Calibration

We follow De Nardi (2004) and model the bequest motive as

$$
V_b(a) = \theta_B \frac{(a + \kappa)^{1-\sigma} - 1}{1 - \sigma} \tag{19}
$$

where θ_B determines the strength of the bequest motive, and κ determines to what extent bequests are a luxury good.

We use the same calibration as in the main text, where applicable, and set $\sigma = 1$. The marginal tax rate on estates τ_b is set to 30%.[11] We determine the discount factor β, the newly introduced parameters governing the bequest motive and the estate tax exception threshold χ_b using the method of simulated moments, i.e., we minimize the weighted sum of squared distances between targeted and simulated moments. For this exercise, we use the model where all agents have the same survival expectations conditional on age (no survival heterogeneity)

Target moments. In the model with a bequest motive, we again target a capital-to-output ratio of 3, but we additionally try to match the life-cycle profile of assets. To this end, we use the median wealth levels at ages 55, 60, 65, 70, 75, 80 and 85 observed in the HRS. The capital-to-output ratio and the life-cycle profile are jointly matched by choosing the parameters for the discount factor (β), the bequest utility weight (θ_B), the bequest utility shifter (κ), and the estate tax exemption (χ_b).

Estimation results. The estimated parameters are listed in Table 5 and the asset holdings by age and their data counterparts are shown in Table 6. Our model is too simplistic to match the data moments exactly. For example, because we impose an exogenous retirement age of 65, the lifecycle profile of assets peaks exactly at this age, whereas this is not the case in the data.

To assess the magnitude of bequests in the model, it is informative to look at non-targeted moments. The bequest-to-wealth ratio in the model is 1.9%. It is difficult to precisely measure this figure in the data, but according to Gale and Scholz (1994), using SCF data, it should be closer to 0.9%. However, according to others, 2% is "a conservative estimate".[12] Another measure is the fraction of estates left by deceased individuals that are subject to estate taxes, which is 0.74% in the model. This figure has varied during the last 15 years: in 2004 it was 0.8% whereas in 2013 is was 0.2%.[13] Consequently, these model quantities are mostly in line with the data.

A last measure is estate tax revenue as a fraction of GDP. In the model, this figure is 0.02%. In the U.S., this figure varied during the last few years, from a high of 0.17% in 2007 to a low of 0.07% in 2011.[14] The model cannot replicate these magnitudes since it is unable to generate the right tail of the wealth distribution and hence also lacks the estates left behind by the richest.

[11] The top marginal tax rate today is 40% (see https://www.irs.gov/pub/irs-pdf/i706.pdf); however, not all taxable estates fall into the top category. We choose 30% as an approximation.

[12] See Larry Summer's opinion at https://www.reuters.com/article/column-summers/column-how-to-target-untaxed-wealth-lawrence-summers-idUSL1E8NG2MC20121216?irpc=932.

[13] Calculated as the ratio of the number of estates subject to estate taxes, as reported by the IRS, and the number of deaths taken from CDC records.

[14] Calculated as the ratio of net estate taxes paid, as reported by the IRS, and official GDP figures.

Parameter	Description	Value
Model with bequest		
β	Discount factor	0.961
θ_B	Bequest weight	10.33
κ	Bequest shifter	1.98
χ_b	Estate tax exemption	27.32

Table 5: Estimated parameters for model with a bequest motive.

Median asset holdings							
Age	**55**	**60**	**65**	**70**	**75**	**80**	**85**
Data	1.30	1.65	1.69	1.73	1.62	1.47	1.41
Model	1.26	1.61	2.13	1.78	1.61	1.43	1.35

Table 6: Life-cycle asset profile in the model with a bequest motive, relative to the median asset holdings at age 50.

D.2.3 Results

Figure 24 shows the changes in total savings rates across the three scenarios in the presence of a bequest motive. These are strikingly different from those in the main text: 50-year-old individuals in poor health save *more*, which is the complete opposite of the behavior when no bequest motive is present. This is a direct consequence of expected utility: since agents in poor health are more likely to die soon, the probability attached to the bequest utility increases and they consequently have an incentive to save more.

This mechanism is always operational in the presence of a warm-glow bequest motive as in (19). With a low weight θ_B on the bequest utility, the life-expectancy channel could still dominate and the net effect could still be that individuals with a shorter life expectancy save less, but the underlying mechanism – that a decrease in expected longevity makes the agent want to save more for bequest reasons – is still present.

Figure 25 shows the life-cycle profiles for the three scenarios. Due to the bequest motive, individuals in older ages do not decumulate their wealth, and the resulting average asset profile is more in line with the data. There is hardly any difference between individuals in poor vs. excellent health, neither in the model with objective survival heterogeneity, nor in the model with subjective heterogeneity. This is a direct result of what we saw in Figure 24: the total savings rates at older ages are quite close to those from the model with no survival heterogeneity.

As the life-cycle profiles for the SSH model show, the median wealth among individuals

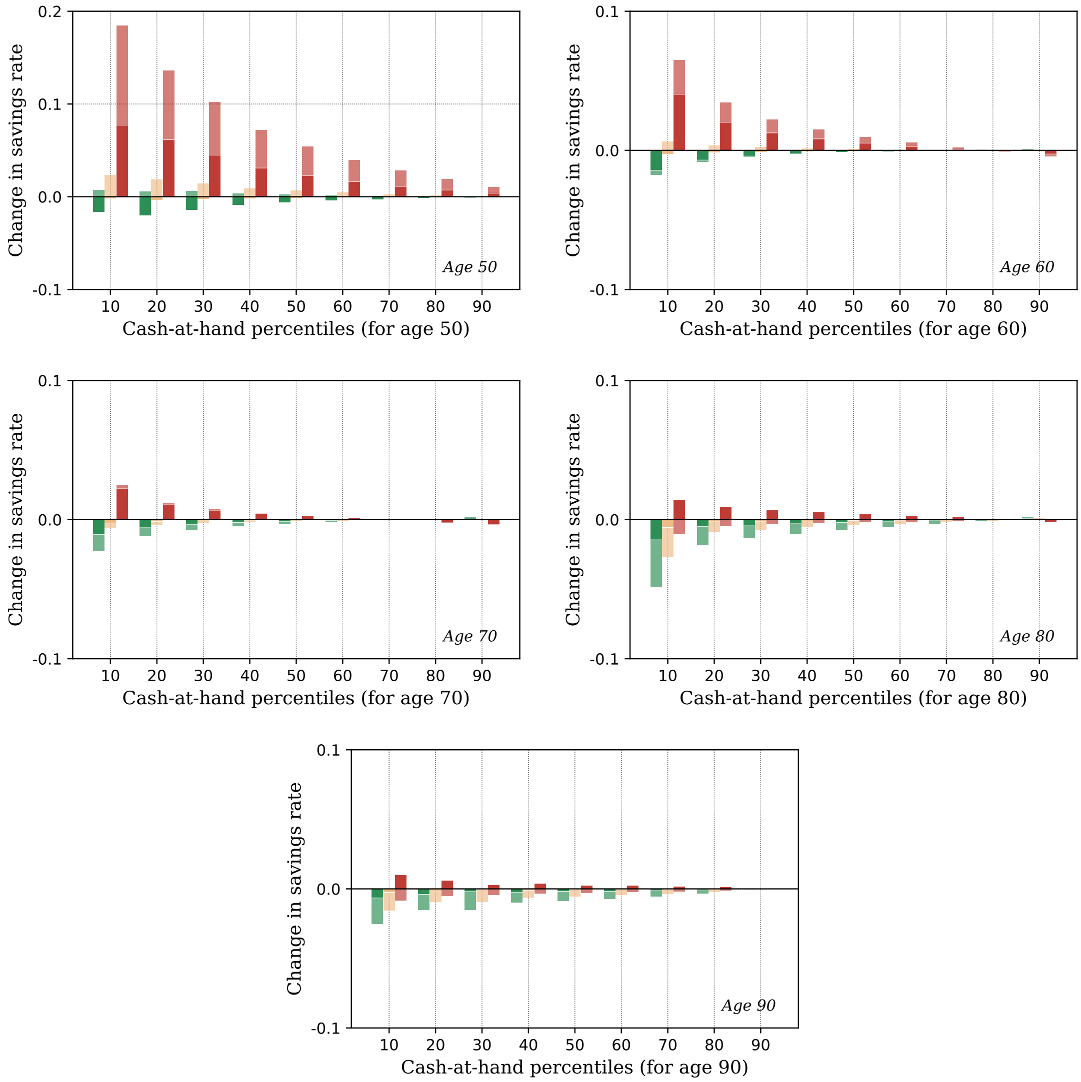

Figure 24: Difference in the absolute savings rate for the model with a *bequest motive*. The darker color shows the model with objective survival heterogeneity compared to the baseline model with no survival heterogeneity. The lighter color shows the additional effect of adding subjective survival beliefs. The x-axis depicts cash-at-hand percentiles by age (equilibrium values for the baseline model). The color indicates health state: green is excellent while red is poor health.

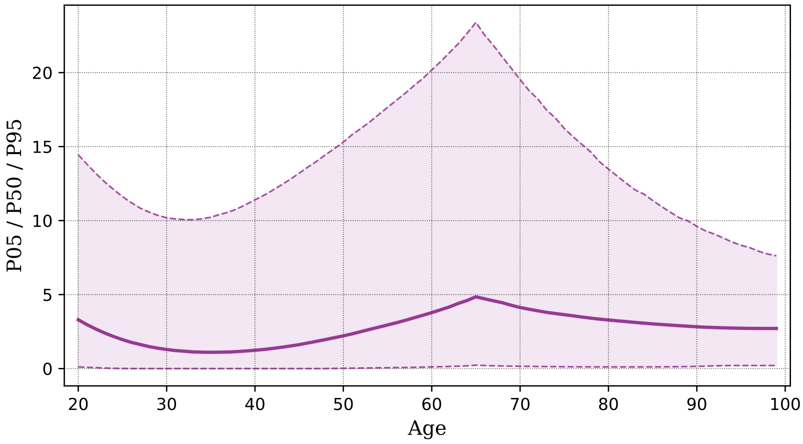

(a) No survival heterogeneity (NSH)

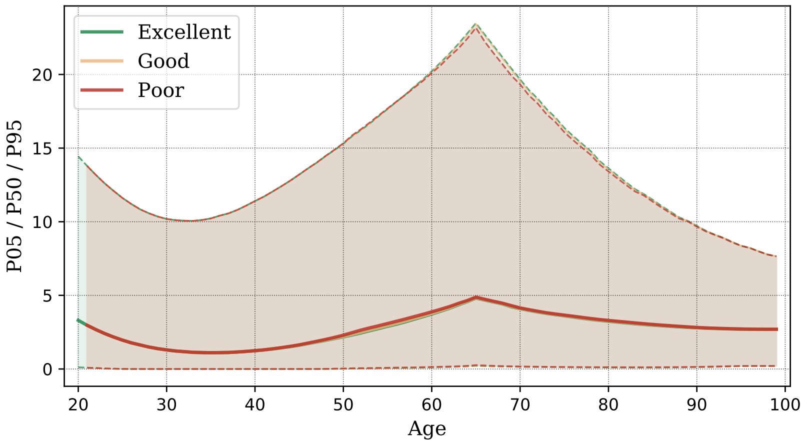

(b) Objective survival heterogeneity (OSH)

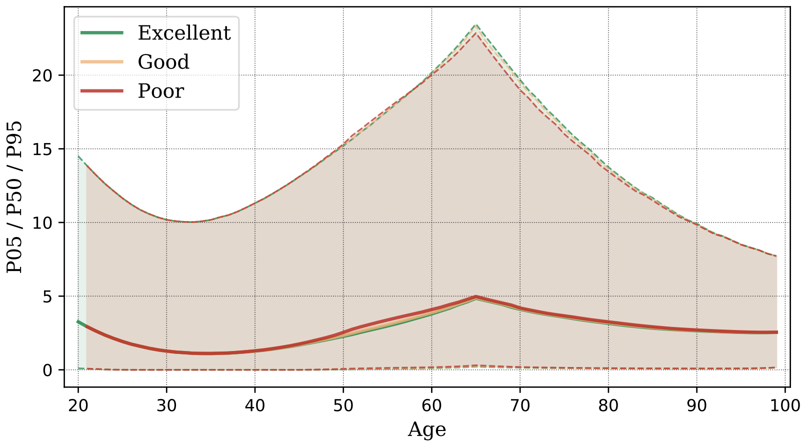

(c) Subjective survival heterogeneity (SSH)

Figure 25: Life-cycle profiles for wealth, model with a *bequest motive*.

in their 50s in poor health is slightly higher than among individuals in excellent health. This is the effect of the higher savings due to the bequest motive. However, wealth in the top P95 is slightly higher for 65-year-olds in excellent health than for 65-year-olds in poor health. The reason is that for very rich individuals, the savings motive to smooth consumption in the event of a longer life becomes relatively more important again, after having saved up so much that the marginal utility from the bequest is small.

These results are difficult to square with cross-sectional data (recall Figure 1) and give rise to the question of what an alternative formulation of a warm-glow bequest motive might look like. However, an answer to this question is beyond the scope of this paper. The inclusion of other economic effects related to the health could be important: for instance, if unhealthy individuals were subject to lower wages (as they are in the data), expensive medical shocks or had some inherent low-savings characteristics, they would have lower asset holdings on average. Since leaving a bequest is a luxury good, individuals who are poor because they lived through many periods of bad health would not experience a strong bequest motive and would consequently save less. This would be a remedy for some of the counter-factual cross-sectional implications: we would most certainly get a model where the unhealthy individuals have lower asset holdings, as in the data.

However, the behavioral mechanism described above would still be present: imagine an asset rich person (for whom the bequest motive is operational) who receives a bad health shock (e.g., a cancer diagnosis) which shortens his/her expected life span. The implication for this person's savings behavior would be that he/she saves more, despite the shorter expected life span, due to the bequest motive. Hence, even though more health-related shocks could be a remedy for the cross-sectional counterfactual health-wealth gradient, the channel of decreasing survival expectancy leading to increased savings due to the bequest channel would still exist.