LONG-TERM-CARE UTILITY AND LATE-IN-LIFE SAVING

John Ameriks
Joseph S. Briggs
Andrew Caplin
Matthew D. Shapiro
Christopher Tonetti

ABSTRACT

Older wealthholders spend down assets much more slowly than predicted by classic life-cycle models. This paper introduces health-dependent utility into a model in which preferences for bequests, expenditures when in need of long-term care (LTC), and ordinary consumption combine with health and longevity uncertainty to explain saving behavior. To sharply identify motives, it develops strategic survey questions (SSQs) that elicit stated preferences. The model is estimated using these SSQs and wealth data from the Vanguard Research Initiative. A robust finding is that the desire to self-insure against long-term-care risk explains a substantial fraction of the wealthholding of older Americans.

John Ameriks
The Vanguard Group, Inc.
PO Box 2600, MS V36
Valley Forge PA 19482
john_ameriks@vanguard.com

Joseph S. Briggs
Federal Reserve Board of Governors
20th Street and Constitution Avenue N.W.
Washington, D.C. 20551
joseph.s.briggs@frb.gov

Andrew Caplin
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
andrew.caplin@nyu.edu

Matthew D. Shapiro
Department of Economics
University of Michigan
611 Tappan St
Ann Arbor, MI 48109-1220
and NBER
shapiro@umich.edu

Christopher Tonetti
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
tonetti@stanford.edu

1   Introduction

The elementary life-cycle model predicts a strong pattern of dissaving in retirement. Yet this strong dissaving is not observed empirically.[1] Establishing what is wrong with the simple model is vital for the optimal design of Social Security, Medicare, Medicaid, retirement savings plans, and private insurance products. Given the aging of the U.S. population, identifying the determinants of late-in-life saving behavior is an increasingly important endeavor.

At present there is no consensus on why there is so little spend down of assets. Explanations typically involve either bequest motives, precautionary motives associated with high late-in-life health and long-term-care (LTC) expenses, or both, but the quantitative contribution of each motive is in debate. Long-term-care needs are a quantitatively plausible driver of saving because the chance of needing care is non-trivial and the costs of care are large. Around one in three older Americans will spend time in a nursing home (Brown and Finkelstein (2008)) and, according to the Genworth (2016) survey available at `http://www.longtermcare.gov`, one year in a private room in a nursing home averages $92,000 and ranges from $55,000 to $250,000. This range in costs partially reflects large variation in quality of care and comfort, which suggests spending when in need of long-term care involves a choice with an intensive margin reflecting the utility of spending when in need of care.

Both bequest and health-related motives can by themselves generate the high levels of observed savings late in life and estimates of the importance of these motives range widely. Kopecky and Koreshkova (2014) finds LTC expenses to be a significant driver of savings and De Nardi, French, and Jones (2010) finds medical expenses to be important in replicating the slow spend-down of wealth. Others, including Hubbard, Skinner, and Zeldes (1994) and Palumbo (1999), estimate the contribution of such expenses to late-in-life savings to be low. Bequests as a saving motive have been studied extensively, with Kotlikoff and Summers (1981) and Hurd (1989) providing early analysis of the effect of a bequest motive on wealth decumulation. Most recent empirical work models the end-of-life bequest motive with the non-homothetic utility functional form proposed in De Nardi (2004). While such studies broadly agree that the bequest motive is present and active primarily for richer individuals (and even found in Lupton and Kopczuk (2007) to be present for individuals without children), its quantitative importance is debated. Lockwood (2016) estimates a near linear bequest utility function which can by itself largely explain the high savings rates of the elderly, but others such as De Nardi, French, and Jones (2010) and Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011) estimate the motive to be weaker, diminishing more rapidly in the bequest level.

We provide new estimates of the relative importance of bequest and precautionary motives. We find the precautionary motive associated with LTC to be, by various measures, more important than the bequest motive as a driver of late-in-life saving behavior. By contrast, our estimated bequest utility parameters suggest that the corresponding motive contributes more modestly to late-in-life saving.

Our results derive from four interrelated innovations. The first concerns the modeling strategy. We build a heterogeneous agent incomplete markets model of individuals, who save precautionarily when faced with health risks, the potential need for long-term care, and an uncertain life span. People value consuming, leaving a bequest, and receiving long-term care if they need it. From at least as early as Arrow (1974),

---

[1]Soto, Penner, and Smith (2009) find that the wealthiest 20 percent of the HRS report rising net worth until age 85, and Poterba, Venti, and Wise (2013) and Love, Palumbo, and Smith (2009) similarly show that household wealth is relatively stable at later ages absent death or divorce.

economists have postulated that utility may be state dependent and that health may be an important state that determines utility. A critical element of our modeling strategy is to allow for an intensive margin of LTC expenditure that is valued using a non-homothetic LTC-state dependent utility function. Specifically, we model LTC utility symmetrically with the bequest utility function proposed in De Nardi (2004). Existing models are asymmetric in this regard, allowing bequests to have a flexible state dependent utility, yet treating long-term care as either a fixed expense or as a portion of standard single period consumption, sometimes with a distinct marginal utility multiplier. Allowing this additional flexibility reflects the distinctive nature of the spending options and desires when in need of long-term care. We also model the option for individuals to utilize the publicly provided insurance against LTC and health risks. We incorporate these social insurance programs as means-tested consumption floors, with a separate provision for LTC and non-LTC health states. While clearly social insurance programs provide consumption for the U.S. population with no wealth, the perceived value of these social insurance programs affects savings across the wealth distribution.[2]

Our second innovation is one of measurement. We develop a series of strategic survey questions (SSQs) to help identify preference parameters (see Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011) and Barsky, Juster, Kimball, and Shapiro (1997)). Our use of this variant of the stated preference method is related to work by van der Klaauw and Wolpin (2008) and van der Klaauw (2012) by the use of non-behavioral data to estimate structural model parameters. In contrast, those papers use subjective expectations data, while we implement SSQs that elicit stated strategies in structured hypothetical scenarios. In addition to novel SSQs, we develop innovative wealth measures that are of particularly high quality, as can be confirmed through linkage to administrative records.

Our third innovation is our estimation approach. We estimate a structural life cycle model in the spirit of De Nardi, French, and Jones (2010), French and Jones (2011), Lockwood (2016), and Gourinchas and Parker (2002). Most importantly, we use not only standard behavioral data but also non-standard SSQ data to jointly estimate risk aversion, LTC utility parameters, and bequest utility parameters.[3] While our favored specification leverages both types of data and estimates the model jointly combining wealth and SSQ moments, we also provide separate estimates using moments of the wealth distribution alone and SSQs alone.[4] Additionally, in order to perform this estimation, we develop a method to efficiently compute optimal

---

[2]For evidence on the effect of such social insurance programs see, e.g., Hubbard, Skinner, and Zeldes (1995), Scholz, Seshadri, and Khitatrakun (2006), Brown and Finkelstein (2008), Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011) and Braun, Kopecky, and Koreshkova (2016).

[3]In previous literature there have been two primary empirical strategies used to identify health-state dependent utility. The first is to use panel data to analyze health profiles over time and the corresponding levels of consumption (Lillard and Weiss (1997)) or utility proxies (Finkelstein, Luttmer, and Notowidigdo (2013)). The primary alternative has been to use a compensating differentials approach (Viscusi and Evans (1990); Evans and Viscusi (1991)), asking survey respondents how much they would need to be paid to compensate for hypothetical health risks, often in the context of physically dangerous jobs. Finkelstein, Luttmer, and Notowidigdo (2009) provides an overview of the empirical strategies used to identify preferences in poor health states. See Hong, Pijoan-Mas, and Rios-Rull (2015) for an alternative method using Euler equations to estimate the effect of health on the marginal utility of consumption.

[4]Our estimation procedure is related to other strategies in which stated choices are used to estimate models in the same manner as are data on observed choices. For a recent example that highlights the similarities and differences between the classic stated-preference and our strategic survey methodologies, see Blass, Lach, and Manski (2010). The closest paper to ours in this dimension is Brown, Goda, and McGarry (2016), who use a related survey methodology to study the degree to which there exists health-state dependent utility. As in our paper, they do find evidence of state dependence. They do not estimate a state-dependent utility function.

policies that builds on the endogenous grid method of Fella (2014). The individual's value function is non-concave in wealth, which generates discontinuous optimal saving policies due to the interaction between free public care, public-care aversion, health-state utility, and minimum private LTC expenditure levels. Efficient computation of the optimal policies using this modified endogenous grid algorithm and parallelization greatly facilitates estimation of the model.

Our final innovation relates to the sample. We derive our results in the context of a new sample, the Vanguard Research Initiative (VRI), that explicitly targets the half of older Americans with non-trivial financial assets. While not randomly selected from the U.S. population, we document in detail in Ameriks, Caplin, Lee, Shapiro, and Tonetti (2014) that this sample has much in common with the appropriately conditioned Health and Retirement Study (HRS). In fact, in Section 9.2 we show that our estimated model predicts well saving patterns out-of-sample in the HRS population. Use of the VRI enables the use of SSQs while simultaneously providing new data on a previously under-sampled relevant population for the question at hand.[5]

Ultimately, we examine the implications of the estimated preferences for savings and expenditure profiles. These estimates suggest that spending when in need of long-term care is highly valued on the margin, and show the relatively greater importance of LTC-related than bequest-based saving motives.[6] It is striking that this broad conclusion holds not only for the estimates based on both wealth and SSQ data, but also for either type of data taken in isolation. With the flexible health-state utility functional form, even estimates targeting exclusively the traditionally-used wealth moments suggest the importance of LTC risk, but imprecisely. Combining SSQ and wealth moments confirms that much of late-in-life saving is driven by LTC-related desires by providing much sharper identification of the motives.

The remainder of the paper is organized as follows. Section 2 develops the model; Section 3 describes the financial and demographic data in the VRI; Section 4 details the strategic survey questions; Section 5 describes the estimation methodology that allows us to estimate the structural life cycle model without (and with) data on observed behavior; Section 6 presents our baseline parameter estimates obtained by matching both wealth and SSQ moments; Section 7 examines the resulting behavioral implications of the estimated preferences; Section 8 compares our baseline estimates to those obtained by exclusively targeting either wealth moments or SSQ moments to disentangle the relative contribution of the SSQs and also compares the baseline estimates to those found in the literature; Section 9 performs sensitivity analysis by re-estimating the model under alternative model and parameter assumptions; Section 10 concludes.

---

[5]Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2016a) uses the VRI to study the demand for insurance against risk of needing help with activities of daily living. It examines both model-implied and stated demands for idealized LTC insurance and compares them to demand for LTC insurance that the market provides. The findings point to substantial unmet demand for good LTC insurance and therefore points to market imperfections in the insurance market. This paper has no insurance market, so the hedge against risk is asset accumulation. Additionally, this paper presents method-of-moments estimates targeting wealth and SSQ data for a representative agent's preference parameters while Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2016a) targets SSQs exclusively to estimate different parameters for each respondent.

[6]In contrast to our findings, Koijen, Van Nieuwerburgh, and Yogo (2016), who allow for similar motives, but use different estimation methods and data on the observed demand for insurance products, find a strong bequest motive and a lower marginal utility when in poor health. Similarly, while not featuring health-specific utility, Lockwood (2016) matches moments of the cross-sectional wealth distribution and LTC insurance demand by cohort and finds strong bequest motives. These findings of low utility when in need of LTC could arise because the poor quality LTC insurance products available in the market are worse than the idealized state-contingent assets in the model.

2 The Model

2.1 Individual States

Individuals are heterogeneous over wealth ($a \in [0, \infty)$), income age-profile ($y \in \{y_1, y_2, \ldots, y_5\}$), age ($t \in \{55, 56, \ldots, 108\}$), gender ($g \in \{m, f\}$), health status ($s \in \{0, 1, 2, 3\}$), and health cost ($h \sim H_g(t, s)$). Time is discrete and the life-cycle horizon is finite. Consumers start at age t_0 and live to be at most T-1 years old, where in our parameterization t_0 is age 55 and T is age 108. Each period, consumers choose consumption (c), savings (a'), and whether to use government care ($G \in \{0, 1\}$). The model groups consumers into five income groups with deterministic age-income profiles.[7] Each person has a perfectly foreseen deterministic income sequence and receives a risk free rate of return of $(1 + r)$ on savings. The only uncertainty an individual has is over health/death.

2.2 Health and Death

There are four health states: $s = 0$ represents good health, $s = 1$ represents poor health, $s = 2$ represents the need for long-term care (LTC), and $s = 3$ represents death. People are defined to need LTC (s=2) if they need help with the activities of daily living (ADLs), such as bathing, eating, dressing, walking across a room, or getting in or out of bed. Thus, state 2 is interchangeably referred to as the LTC or ADL state. The health state evolves according to a Markov process, where the probability matrix, $\pi_g(s'|t, s)$ is gender, age, and health state dependent. Health status affects the distribution of out-of-pocket health expenditure shocks and the distribution of health status in the following year (including mortality risk). Out-of-pocket health expenditures (h) are lognormally distributed as a function of age, health status, and gender. Section 5.1.2 describes the mapping of health variables between model and data.

Health-State Dependent and Bequest Utility. In addition to affecting health costs and survival probabilities, health status affects preferences. There is a health-dependent utility function, such that spending when a consumer needs LTC is valued differently than spending when a consumer does not need LTC. When in good or poor health ($s \in \{0, 1\}$), consumers value consumption according to standard CRRA preferences with parameter $\gamma > 0$:

$$U_{s \in \{0,1\}}(c) = \frac{c^{1-\gamma}}{1 - \gamma}. \tag{1}$$

Utility when in need of help with LTC ($s = 2$) associated with chosen expenditure level c is:

$$U_{s=2}(c) = (\theta_{ADL})^{-\gamma} \frac{(c + \kappa_{ADL})^{1-\gamma}}{1 - \gamma}. \tag{2}$$

---

[7]The model abstracts from labor supply decisions, including retirement. These labor market decisions are taken into account through the exogenous income profiles.

Upon death ($s = 3$), the agent receives no income and pays all mandatory health costs. Any remaining wealth is left as a bequest, b, which the consumer values with a warm glow utility function:

$$v(b) = (\theta_{beq})^{-\gamma} \frac{(b + \kappa_{beq})^{1-\gamma}}{1 - \gamma}. \tag{3}$$

When an individual is healthy or sick, utility is given by a power utility function of consumption. Bequests are valued using the standard warm glow utility function developed in De Nardi (2004). When an individual needs long-term care, utility is given by a similar formula, which treats LTC and bequests symmetrically in theory, allowing differences in preferences to be determined empirically through estimated parameter differences.[8] Two key parameters are θ and κ; θ affects the marginal utility of an additional dollar spent and κ controls the degree to which an expenditure is valued as a luxury good or a necessity, in the sense that it provides a utility floor or need. Since it is raised to the power $-\gamma$, increases in θ decrease the marginal utility of a unit of expenditure. κ allows the model to represent non-homothetic preferences; an increase in κ indicates that the expenditure is valued as more of a luxury good; negative κ can be interpreted as the expenditure being a necessity.

2.3 Government

A person always has the option to use a means-tested government provided care program. The cost of using government care is that a consumer's wealth is set to zero, while the benefit is that the government provides predetermined levels of expenditure, which depend on the health status of the individual. If a person chooses to use government care when not in need of LTC (i.e., when $s = 0, 1$), then the government provides a consumption floor, $c = \omega_G$, that is designed to represent welfare.

If an individual needs LTC ($s = 2$), then he or she must either purchase private long-term care or use government care. Capturing the fact that LTC provision is essential for those in need and private long-term care is expensive, there is a minimum level of expenditure needed to obtain private LTC parameterized by χ, i.e., $c \geq \chi$ if $s = 2$ for those not using government care. In the model, government-provided care is loosely based on the institutions of Medicaid. If a person needs LTC and uses government care, the government provides $c = \psi_G$. The value ψ_G parameterizes the individual's value of public care, since that parameter determines the utility of an individual who needs LTC and chooses to use government care.

---

[8]We follow most papers in this literature, e.g., De Nardi, French, and Jones (2010, 2016) and Lockwood (2016), by using the same exponent in the healthy and bequest utility functions. We explore the sensitivity of our results to allowing for a bequest-specific exponent, γ_{beq}, in Section 9. De Nardi, French, and Jones (2016) also model a health-state dependent utility function to study late-in-life saving patterns. Their focus is on medical spending more generally among a less wealthy population; they model an additively separable homothetic medical expenditure utility function with a marginal utility multiplier that varies as a function of nursing home need and other medical states (e.g., broken bones). By contrast, our health-dependent utility function is over total consumption in the LTC state and represents non-homothetic preferences via κ_{ADL}. Thus, our parameterization provides the LTC utility function the same flexibility as the bequest utility function in terms of generating varied spending and saving behavior across wealth level. Indeed, as documented in Sections 6 and 7, we estimate κ_{ADL} to be negative and large, which significantly affects behavior.

2.4   The Individual Problem

The individual takes r as given and chooses a', c, and G to maximize utility. This problem, written recursively, is,

$$V(a, y, t, s, h, g) = \max_{a', c, G} \mathbb{I}_{s \neq 3} (1 - G) \left\{ U_s(c) + \beta E[V(a', y, t + 1, s', h')] \right\}$$

$$+ \mathbb{I}_{s \neq 3} G \left\{ U_s(\omega_G, \psi_G) + \beta E[V(0, y, t + 1, s', h')] \right\} + \mathbb{I}_{s=3} \{ v(b) \}$$

s.t.

$$a' = (1 - G)[(1 + r)a + y(t) - c - h] \geq 0$$

$$c \geq \chi \quad \text{if} \quad (G = 0 \ \wedge \ s = 2)$$

$$c = \psi_G \quad \text{if} \quad (G = 1 \ \wedge \ s = 2)$$

$$c = \omega_G \quad \text{if} \quad (G = 1 \ \wedge \ (s = 0 \ \vee \ s = 1))$$

$$b = \max\{(1 + r)a - h' \ , \ 0\}$$

$$U_s(c) = \mathbb{I}_{s \in \{0,1\}} \frac{c^{1-\gamma}}{1 - \gamma} + \mathbb{I}_{s=2} \ (\theta_{ADL})^{-\gamma} \frac{(c + \kappa_{ADL})^{1-\gamma}}{1 - \gamma}$$

$$v(b) = (\theta_{beq})^{-\gamma} \frac{(b + \kappa_{beq})^{1-\gamma}}{1 - \gamma}.$$

The value function has three components, corresponding to the utility plus expected continuation value of a living individual who chooses not use government care, that of one who chooses to use government care, and the warm glow bequest utility of the newly deceased individual.[9] $G = 1$ if the consumer chooses to use government care and $G = 0$ if the consumer chooses not to use government care. A person using government care has expenditure levels set to predetermined public care levels and zero next-period wealth. The budget constraint shows that wealth next period is equal to zero if government care is used, and is otherwise equal to the return on savings plus income minus chosen expenditures minus the health cost shock. The individual cannot borrow, cannot leave a negative bequest, and private expenditure when in need of LTC must be at least χ.

2.5   Describing Optimal Behavior

In this section, we explore key properties of optimal individual behavior to illustrate how each force in the model contributes to consumption and savings patterns over the life cycle and across the income and wealth distributions. The individual's saving behavior is largely determined by the confounding influence of the precautionary saving motive and bequest motive in the presence of government policies. Long-term care needs occur with non-trivial probability and paying for such care privately is very costly. The fact that the government offers a means-tested public care option induces interesting behavior. Because the individual has the option to choose government care, the value function is non-concave and the optimal saving policy is discontinuous. The model does not permit analytic solutions and must be solved numerically, with details of our solution algorithm presented in "Vanguard Research Initiative Technical Report: Long-term Care

---

[9]Technically, there is a fifth health state that is reached (with certainty) only in the period after death and is the absorbing state, so that the consumer only receives the value of a bequest in the first period of death.
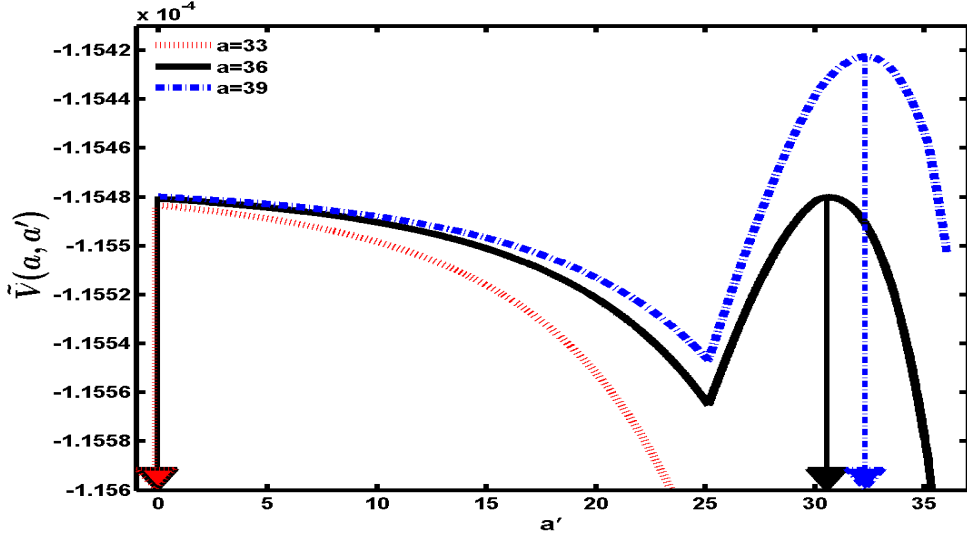
Figure 1: Saving Policy Discontinuity (Assets a, \$1000s)

Model."[10]

Discontinuous saving policies. The option to use means-tested government care induces discontinuous saving policies, in a manner similar to that studied in Hubbard, Skinner, and Zeldes (1995). Roughly speaking, very high wealth individuals have enough savings to ensure they will obtain a high level of personal consumption and leave a large bequest, regardless of whether or not they need to pay for private long-term care. For low wealth individuals, even if they saved almost all of their money and consumed very small amounts each year, they would not be able to save enough to make it optimal for them to purchase private long-term care if they eventually needed it. Thus, it is the middle-wealth people—like those in the VRI—whose actions are most likely to be affected by precautionary saving motives. If these middle-wealth individuals are frugal and save, they will have enough wealth to purchase private LTC if they need it late in life. If they do not save, but rather consume at a high rate over their life cycle, they will have higher utility when alive and healthy, but will forgo a bequest and rely on public provision of LTC if they need it later in life. There exists some threshold wealth level, conditional on all other idiosyncratic state variables, such that it is optimal for all agents with more wealth to follow the frugal path and for all agents below to follow the spendthrift path, with a discrete difference in their saving policy for a tiny difference in their wealth state. To illustrate optimal consumer behavior we present model simulations at certain parameter values. Parameters will be estimated and discussed further in Section 6.2.

The discontinuity of the saving policy is demonstrated in Figure 1 by plotting the objective function that corresponds to the non-optimized value function across saving policies for different wealth states. Plotted on the horizontal axis is $t + 1$ wealth, and on the vertical axis is the associated value of that saving policy for a given level of period t wealth. The three lines depict the graph for an individual with identical states

---

[10]The non-concavity of the value function and the discontinuity in the optimal savings policy introduce computational complications. We use a modified endogenous grid method, building on insights from Fella (2014). The model solves approximately ten times faster when using the modified endogenous grid algorithm compared to value function iteration, which is essential since estimation of the model requires computational efficiency.

aside from the three different wealth levels: $33K in the red dashed line, $36K in the solid black line, and $39K in the dotted blue line.

In this example, the optimal value of saving for the low wealth individual is zero. This person is forgoing all precautionary saving and consuming as much as possible today. The higher wealth individual has enough money that it is worth it to save to help smooth consumption across states and time. This can be seen in the top line, in which the maximum of the value function is achieved with savings around $32K, attaining a significantly higher value than that of saving zero. As presented in the middle line, there exists a current wealth level for which the global maximum value jumps from the lower to the higher savings local maxima. An individual with $33K in this example is near indifferent between saving around $31K and saving zero. It is around this level of current wealth where there is a discrete jump in the optimal savings policy (the value function is kinked, but remains continuous).

How preferences and states determine saving behavior. Saving decisions are ultimately determined by the preferences of individuals and by their environment. As was highlighted by Dynan, Skinner, and Zeldes (2002), a dollar saved today is fungible in its future use. Saving early in the person's life could be to insure against future uncertain events like LTC as well as to ensure suitable savings remain at end of life to leave a desired bequest. If the bequest motive is weak, over-saving for an uncertain late-in-life event that never occurs is costly, as the individual would much rather have had a smooth higher consumption path over his life. However, with a strong bequest motive, "extra" savings at the end of life are highly valued, which reduces the cost of ex-post over-saving.
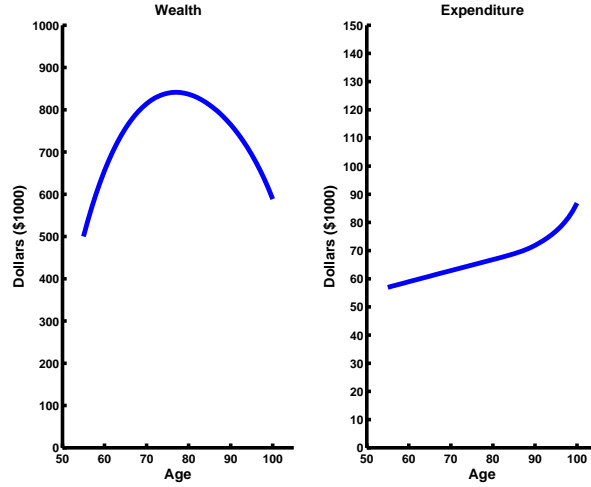
Figure 2: Wealth and Expenditure Profiles for Healthy Male

To demonstrate how savings are influenced by bequest and LTC-induced saving motives, we plot various age-profiles for wealth and expenditure for a simulated individual, in response to different sequences of health shocks, for different initial states, and for different preference parameters. Unless otherwise stated, the figures plot the wealth and expenditure profiles of a male who starts healthy at age 55 and has the median income profile, median wealth, and preference parameters from our preferred baseline estimation. Unshaded areas indicate behavior when healthy ($s = 0$) while gray-shaded regions indicate behavior when in need of long-term care ($s = 2$). It is important to note that these patterns will not be representative of

8

wealth and expenditure profiles of the population, as these are individuals and shocks selected to illustrate the workings of the model and are not necessarily typical or representative of the VRI sample or the U.S. population.

As a baseline, Figure 2 shows the wealth and consumption paths of a man who receives a shock sequence such that he remains in good health until death at $T = 108$. Wealth accumulates until age 75 and then steadily decumulates with age. Early on the individual saves, driven by a combination of LTC and bequest motives. As the individual ages, the probability of needing LTC for any given year tends to increase, but eventually the chance that LTC will be needed for any given large number of years decreases. The increase in consumption with age occurs because the individual was saving precautionarily for LTC and as he continues to receive such a good run of positive health shocks, he starts to consume the ex-post extra savings slowly.

(a) Median Wealth   (b) Low Wealth
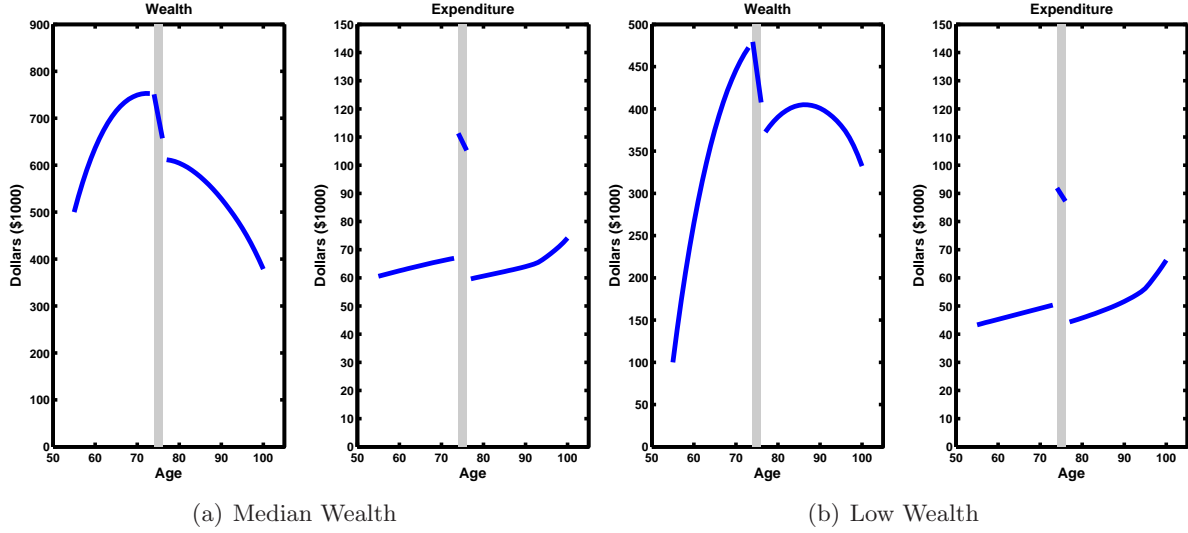
Figure 3: Wealth and Expenditure Profiles for Median Income Male

Figure 3(a) demonstrates the rapid dissaving and high expenditure associated with the need for LTC. This person received health shocks such that he was healthy his entire simulated life, except for one period in which he needed LTC for ages 74-76, highlighted by the gray shaded region. At the onset of needing LTC, expenditures jump from around $60,000 per year to around $110,000 per year, resulting in a large decrease in wealth. Expenditure remains high and roughly constant during the three year LTC period, as savings decline rapidly. After three years of LTC, the individual steadily dissaves and consumes, as no other adverse health shocks occur until death. Saving and expenditure behavior depend on an individual's level of wealth. Figure 3(b) plots the behavior of an individual that is similar in all ways except for having lower wealth at age 55. The low wealth individual saves more aggressively early on in order to build a buffer stock of wealth in case LTC is needed and in order to be able to leave a bequest. Similar patterns of rapid dissaving and high levels of expenditure are associated with the LTC event. However, the low wealth individual actually increases wealth after exiting long-term care to return to a desired buffer-stock level of wealth.

Figure 4(a) documents the behavior of a lower wealth individual who also has the lower first-quintile income profile. Furthermore, compared to the previous figure, in this simulation his need for LTC lasts for nine years instead of three. At first he purchases private LTC, but the high level of expenditure associated with his need for LTC depletes his wealth to near zero, at which point he chooses to use publicly provided

(a) Low Wealth   (b) Very Low Wealth

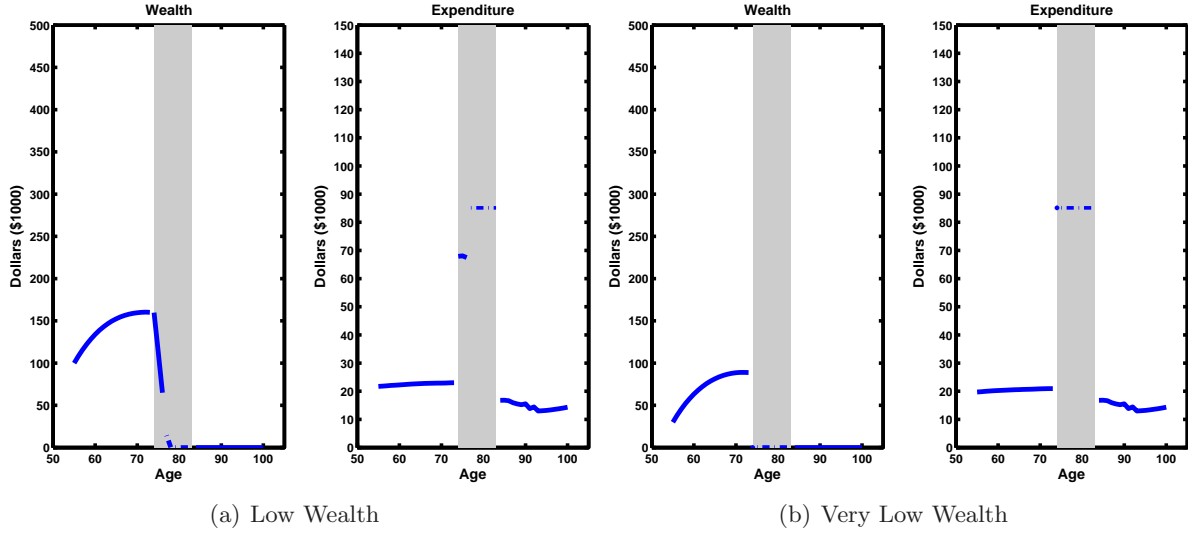Figure 4: Wealth and Expenditure Profiles with Publicly Provided LTC

LTC for the rest of his LTC episode, and then live hand to mouth afterwards. Note that public-care expenditure (dashed line) is included in the total expenditure reported.[11] Figure 4(b) shows what happens if the individual started at age 55 with $30,000 in savings instead of $100,000. He consumes very little and saves up until he needs LTC. His wealth is so low that he immediately uses public care as soon as he needs LTC. When he no longer needs public care, he simply consumes his roughly $20,000 a year income. As is apparent, the need for extended LTC rapidly depletes savings and can lead to extended periods of low consumption for the remainder of life.

Quantitatively, the levels of expenditure are quite reasonable across the wealth and income distribution. An individual who has $750,000 in wealth at age 74 and earns around $50,000 a year spends around $110,000 a year during a three-year LTC stay, while an individual who has $150,000 in wealth at age 74 and earns $20,000 a year spends around $70,000 a year for the same three-year LTC stay.

These saving and expenditure patterns are strongly influenced by people's preferences. To demonstrate the importance of the health-state utility function, Figure 5(a) recreates the simulation presented in Figure 3(a), except for an individual with preferences such that spending when in need of long-term care is valued just as spending when healthy ($\theta_{ADL} = 1$, $\kappa_{ADL} = 0$). The original behavior induced by baseline preferences is drawn with dashed lines and that associated with alternative preferences is drawn with solid lines. This analysis shows that much of the increase in wealth during the individual's 50's and 60's was driven by the precautionary saving motive associated with LTC. Furthermore, expenditure levels when in need of long-term care are much closer to expenditure when healthy, with a slight uptick due to the increased mortality risk associated with the worse health state. This major change in expenditure patterns foreshadows that our estimated health-state utility function induces higher marginal utility of expenditure when in need of LTC, not less.

Both the health-state utility function and the bequest function affect saving and spending behavior. Figure 5(b) plots the life-cycle behavior of the same median wealth and income individual, but with param-

---

[11]For a discussion of the level of public-care expenditure (ψ_G), see Section 6.2.

(a) No ADL-State Utility   (b) Strong Bequest Utility
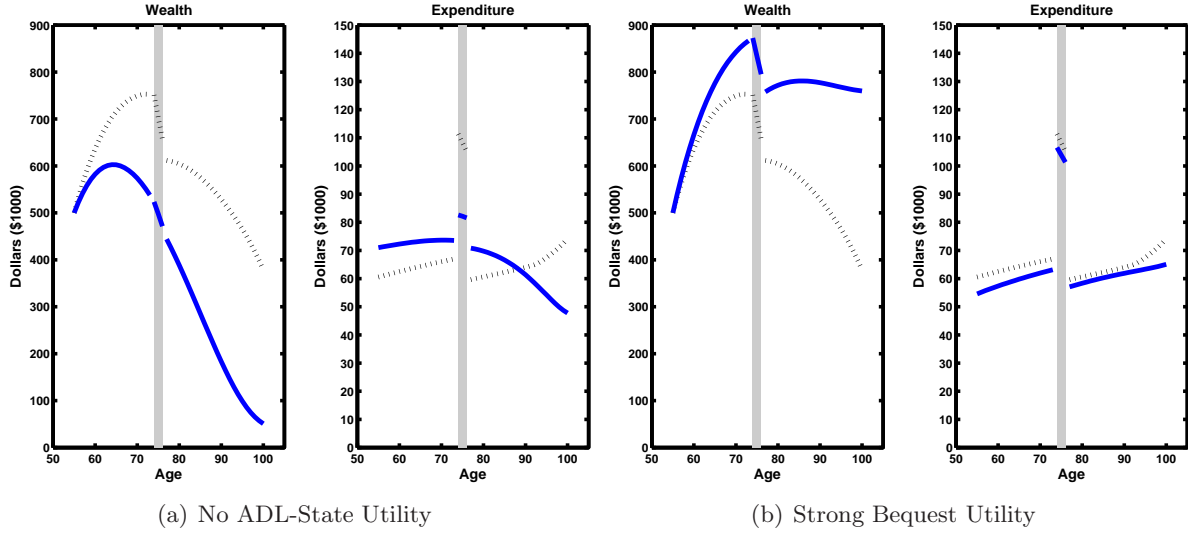
Figure 5: Wealth and Expenditure Profiles with Baseline (black dotted line) and
Alternative (blue solid line) Preferences

eters such that bequests are more strongly valued. As can be seen, the stronger bequest motive increases savings early on. Furthermore, the stronger bequest motive has a significant effect on late-in-life wealth levels, leading an individual to reach age 100 with near double the wealth of the baseline individual. This person needed to save so much early on because he had a strong desire to spend when in need of LTC and to leave a bequest. Later in life, expenditure patterns look similar, because consumption similar to that in the baseline case can be sustained without depleting wealth, due to the higher level of financial income generated by a larger stock of wealth.

(a) Needs Care Ages 85-89 then Dies   (b) Healthy Until Age 100
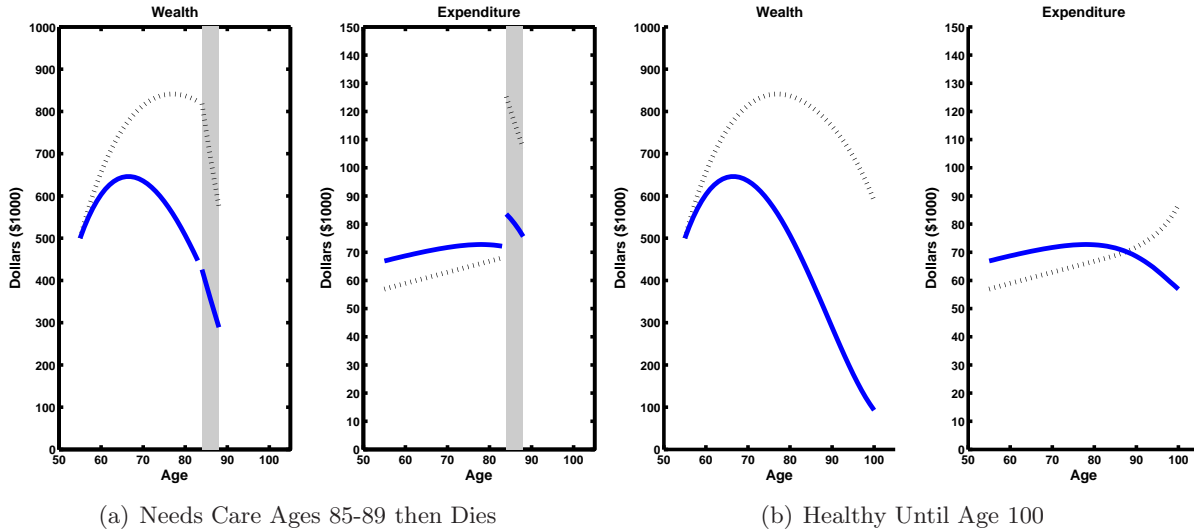
Figure 6: Wealth and Expenditure Profiles with (black dotted line) and without (blue solid line)
Health-State Utility for Different Health Paths

Finally, to give a sense of the effect of health-state utility on saving and bequest behavior in the data, we plot in Figure 6(a) a more typical health pattern in which a median wealth and income man is healthy from

11

age 55 to 84, needs LTC for age 85–89, and then dies. The figure plots wealth over the life cycle with the baseline parameters and also with no ADL-state utility. Without ADL-state utility, this person accumulates wealth roughly from ages 55 to 70, and then starts steadily dissaving. When in need of LTC, his rate of dissaving is only marginally faster, driven by changes in life expectancy. In contrast, with the baseline ADL-state utility function, this person accumulates wealth more rapidly from age 55-70 and continues to accumulate wealth until age 80. Furthermore, when the negative health shock realizes, dissaving is rapid, driven by the higher utility of expenditure in this health state. This desire to spend when in need of help with ADLs is to a large degree why he was saving in the first place. In this scenario, the baseline parameterization results in almost double the bequest compared to the no ADL-state utilty case, even with the rapid health-related spend down at the end of life due to the increased precautionary saving earlier in life. Many people die without ever needing long-term care. Figure 6(b) plots savings over the life cycle for a man who is healthy until death at age 100, with and without the health-state utility function. This figure demonstrates that the health-related precautionary saving motive generates a large incidental bequest even though there is the typical spend down late in life due to mortality risk absent LTC needs.

With an understanding of the key features of optimal saving behavior in the model and how they relate to important state and parameter values, we turn to a description of the data, with which these parameter values will be estimated.

3   Financial and Demographic Data

In order to examine late-in-life wealth patterns, it is essential to have data on a population with large enough financial resources to face non-trivial spending, saving, and giving decisions. This paper draws on the newly developed Vanguard Research Initiative (VRI) that combines survey and administrative account data.[12] In this section we briefly describe the VRI, highlighting the advantages of the sample population for addressing the question at hand.

The VRI consists of approximately 9,000 individuals drawn from Vanguard account holders who are at least 55 years old. Additionally, we require Vanguard assets of at least $10,000 (to assure non-trivial engagement with Vanguard) and Internet registration with Vanguard (to allow for surveys administered over the Internet). As a point of comparison, the VRI is cross-sectionally about the same size as the Health and Retirement Study (HRS) and around 4 times larger than the Survey of Consumer Finances (SCF) in the relevant age group. Surveys are administered over the Internet and ask respondents about their and their spouse's or partner's wealth, income, and decision-making motives.

A sample drawn from Vanguard account holders is, of course, not random or representative of the U.S. population. For example, by construction, the sample is drawn from individuals who have positive financial wealth. Hence, we exclude the large fraction of households who approach or reach retirement age with little or no financial assets. Use of this new dataset is a significant contribution of this paper. It provides a large sample of older Americans with sufficient financial assets to face meaningful trade-offs between consumption

---

[12]The VRI data contain non-public information that cannot be freely disclosed, so the dataset cannot be publicly distributed. Access to the VRI is at Vanguard's discretion, though Vanguard and the research team will work to make it available for replication and research within appropriate limits. The computer code written to solve and estimate the model is available publicly.

across time, between spending when well and when in need of assistance, between long-term care in private or publicly-funded facilities, and between leaving bequests versus spending while alive.

Since we do not explicitly model the family, in this paper we restrict our data to only include single respondents, who were oversampled to ensure a large single subsample. For the remainder of this paper we focus on a sample of $I = 1,241$ singles with no missing survey responses to mandatory questions.

Table 1: Income and Wealth Distribution Across Surveys: VRI-Eligible Single Households

		Financial Wealth					
	I	**Mean**	**10p**	**25p**	**50p**	**75p**	**90p**
VRI	1,241	808,007	101,000	262,113	527,600	993,800	1,602,000
HRS	1,201	376,432	24,000	68,000	178,000	445,000	920,000
SCF	265	487,234	18,500	58,500	159,000	410,700	1,019,000

		Income					
	I	**Mean**	**10p**	**25p**	**50p**	**75p**	**90p**
VRI	1,241	69,452	17,500	34,223	56,000	86,550	121,473
HRS	1,201	65,402	10,860	18,817	36,000	65,000	105,012
SCF	265	80,963	25,363	35,509	51,741	85,221	121,744

Financial wealth is the sum of IRA, employer sponsored retirement, checking, saving, money market, mutual fund, certificate of deposit, brokerage, and educational related accounts plus the current cash value (if any) of life insurance and annuities. Income is defined as the sum of labor income, publicly and privately provided pensions, and disability income. The sample comprises single households meeting VRI sample screens: age 55 years and older; assets of at least $10,000, and Internet access. Income is total household income (excluding distributions from defined-contribution pension plans). See Ameriks, Caplin, Lee, Shapiro, and Tonetti (2014), Appendices B and C, for a discussion of the definitions of variables in the HRS and SCF and for a detailed comparison of the VRI, HRS, and SCF.

For points of comparison, we construct "VRI-eligible" subsets of the Health and Retirement Study (HRS) and the Survey of Consumer Finance (SCF) by imposing sample screens to parallel the VRI: age 55 years and older, financial assets of at least $10,000, and access to the Internet. After imposing these screens, the characteristics of the VRI sample are similar in many dimensions to these subsets of the 2012 HRS and 2013 SCF, representing individuals in roughly the upper half of the wealth distribution. Financial wealth is defined as the sum of IRA, employer sponsored retirement, checking, saving, money market, mutual fund, certificate of deposit, brokerage, and educational related accounts plus the current cash value (if any) of life insurance and annuities. Income is defined as the sum of labor income, publicly and privately provided pensions, and disability income. Table 1 compares wealth and income of the VRI and VRI-eligible subsets of the HRS and SCF restricted to the single households considered in this paper. Our sample is well positioned to complement existing samples with a highly relevant population. In Table 1, we see VRI respondents have more wealth than the VRI-eligible HRS and SCF respondents, but the differences are much less stark than compared to the overall population, which has close to zero wealth at the median. Furthermore, although the income is somewhat higher in the VRI than in the VRI-eligible HRS, the VRI and the VRI-eligible SCF

have very similar levels of income.

For more details we refer the reader to Ameriks, Caplin, Lee, Shapiro, and Tonetti (2014), which provides an exhaustive analysis of the VRI, both on the survey methodology and on the resulting collected data. For the purposes of this paper, it is most important to note that the VRI contains high quality measures of individuals' wealth and income and, crucially, responses to SSQs that were specifically designed to identify parameters of the model just developed. The VRI also has measures of self-reported health status and need for help with activities of daily living, elicited using the same questions as in the HRS that we use to estimate health-state transition matrices. To estimate both the health-state transition probabilities and the out-of-pocket health expenditure shock distributions, we use HRS panel data, as detailed in Section 5.

4   Strategic Survey Questions

SSQs are stated-preference questions, with scenarios that are quite closely linked to important life decision faced by our respondents. Behavior in the model is driven by the preferences of individuals and the economic environment in which they make choices. Since a main goal of this paper is to identify the relative contributions of different saving motives associated with different preferences, it would be ideal if survey respondents could accurately and directly report their preference parameters. Of course, we can not ask survey respondents to report their coefficient of relative risk aversion, much less θ_{ADL}. Thus, if we want to develop direct measures of preferences, we need to develop survey instruments that allow respondents to provide us with information that identifies preference parameters in a language in which they are comfortable, but also in a format that allows a precise mapping to structural parameters of interest.

Along these lines, revealed preference methodology uses observed choices to perform inference about preferences. If a utility function is assumed to represent preferences, often these observed behaviors can be used to estimate preference parameters. In a similar vein, we develop strategic survey questions that use choices made in hypothetical scenarios to estimate preference parameters. In constructing our survey, we create a highly structured hypothetical environment with a very restricted choice set that allows us to make fewer assumptions on the unspecified economic environment. Though necessarily incomplete per se, our scenarios are significantly more detailed than in typical hypothetical questions. Our questions are designed to provide the survey respondent precise details on all relevant individual states of the world from the perspective of the structural model. SSQs ask the respondent to comprehend and imagine complex scenarios. As with any hypothetical question, there are legitimate concerns about whether survey respondents can understand the scenario and whether they can respond from the perspective of that scenario. To make these tasks as easy as possible for the survey participant, we paid close attention to the presentation of the material and developed the survey with input from survey design experts and cognitive psychologist in a process of exploratory, pilot, and then production surveys; we also perform tests of respondent comprehension and of the coherence and consistency of responses to address such concerns. See Section 4.3 for analysis of SSQ responses and Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2016a) for further evidence on the credibility and coherence of SSQ responses.

The SSQs are one ingredient in our overall estimation strategy to identify preference parameters. We treat each individual as an optimizer in a specific problem and characterize the individual by financial and demographic variables and preference parameters. The SSQs ask people to make choices that would be

very revealing about their preference parameters if only we were to observe them making such choices. A key feature of the SSQ approach is that it provides data about the priority of LTC risks and desires for all individuals in the sample, not just the ones who end up needing LTC. We add as moments in the GMM estimation the mean of each SSQ so that we are estimating preferences of an optimizer who would answer the SSQs in a manner consistent with the central tendencies of how the survey respondents answered the SSQs. Because these are deep structural preference parameters, they also affect all other behavior of these modeled individuals, including their saving behavior. In addition to the SSQs, we include wealth moments conditional on age to make sure our optimizing individual has preferences that are also consistent with the saving behavior of our survey respondents. There is a weighting matrix that determines the weight on the different moments, but all moments are informative of all preference parameters. In the context of this study, it turns out that two very different types of individuals could equally well have generated the wealth data: one with strong bequest motives and weaker LTC-related spending desires or vice versa. The SSQs allow us to sharply identify behavior that is only weakly identified in the wealth data alone: an individual with a very strong active bequest motive and weak LTC utility would not answer the SSQs as the survey respondents did, while an individual with strong LTC utility and a somewhat weaker active bequest motive would answer the SSQs like our survey respondents.

4.1 Detailing SSQ 2: Spending in Good Health vs. when Needing LTC

Ultimately the model is estimated on responses from four types of SSQs and wealth data. In this section, we will first illustrate the key features of SSQs by detailing one particular SSQ related to LTC (SSQ 2) and then we will present the other three SSQs. In Section 5, we detail more precisely how we construct moments and use model simulations to estimate parameters.

In SSQ 2, we are interested in understanding how individuals trade off having wealth in states of the world when they do not need LTC and when they do need LTC. At the core of the question, we are asking individuals to solve a simple portfolio allocation problem. The researchers specify that the respondent has some wealth (W), faces some chance they will need LTC ($1 - \pi$) and some chance they will not need LTC (π), and that they must allocate their resources by purchasing state-contingent assets (z_1, z_2) given a relative price of z_2 (p_2) to finance expenditure in the two possible states of the world. In the survey, we set $p_2 = \frac{1}{1-\pi}$. The mathematical representation of the survey question that we use for estimation is:

$$\max_{z_1, z_2} \quad \pi \frac{z_1^{1-\gamma}}{1-\gamma} + (1-\pi) \frac{(\theta_{ADL})^{-\gamma}(z_2 + \kappa_{ADL})^{1-\gamma}}{1-\gamma} \tag{4}$$
$$\text{s.t.} \quad z_1 + p_2 \, z_2 \leq W$$
$$z_1, z_2 \geq 0; \ z_2 \geq -\kappa_{ADL}.$$

Identification. The first order condition of the optimization problem gives the optimal allocation as a function of preference parameters. By inverting this function, we map the allocations chosen by survey respondents to preference parameters. For example, the optimal decision rule of the above problem is given

15

by:

$$z_2 = \begin{cases} 0 & \text{if } \pi\left(\frac{W}{p_1}\right)^{-\gamma}\left(\frac{p_2}{p_1}\right) - (1-\pi)(\theta_{ADL}\kappa_{ADL})^{-\gamma} > 0 \\ \dfrac{W\left(\frac{\pi p_2}{p_1(1-\pi)}\right)^{-\frac{1}{\gamma}} - p_1\theta_{ADL}\kappa_{ADL}}{p_1\theta_{ADL} + p_2\left(\frac{\pi p_2}{p_1(1-\pi)}\right)^{-\frac{1}{\gamma}}} & \text{otherwise.} \end{cases}$$

The scenario exogenously specifies π and W. By obtaining responses for two different combinations of W and π for which allocations to both states are positive, we are able to identify θ_{ADL} and κ_{ADL}, given γ. Combining these with the results of SSQ 1 that identify γ we are thus able to identify all three parameters from responses to these questions.

The key survey design challenge is that most individuals can not understand the allocation problem in the mathematical language of optimal control. We present below the SSQ that is designed to help survey respondents provide (z_1, z_2) such that they are making a choice that we are confident corresponds to that in the optimization problem, but in a format in which they are capable of doing so. Identification for all preference parameters proceeds similarly, as demonstrated below.

Survey Design. To ease respondent comprehension these questions are presented in five steps. In the first screen, we begin by telling the respondent explicitly what trade-off we are asking them to think about. This is done to prompt the respondent to weigh the relevant risks we are interested in, and to alleviate their concern over not understanding the point of the question and guessing about the motives of the survey designers. Second, the question presents the specific scenario and details the choices that the respondent must make. This screen is the complete scenario, and is made available to the respondents as they are giving their final answers if they would like to check any features of the scenario. Third, we present a set of rules that further defines the environment and clarifies the span of options and resources available to the survey respondent in each scenario. Fourth, we verify comprehension and reinforce key features of the question with a set of multiple choice questions about the scenario. Finally, we record an answer using a custom-designed interactive slider. See Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2016a) for a detailed discussion of the SSQ design process and further analysis of individual responses that demonstrate the credibility of SSQs.

The Scenario. The survey instrument first states the scenario precisely, but as simply as possible consistent with being precise. Specifically, the survey displays a screen with the following text.[13]

We are interested in how you trade off your desire for resources when you do and when you do not need help with activities of daily living (ADLs). This scenario is hypothetical and does not reflect a choice you are likely ever to face.

Suppose you are 80 years old, live alone, rent your home, and pay all your own bills. Suppose that there is a chance that you will need help with ADLs in the next year. If you need help with ADLs you will need long-term care.

---

[13]In previous sections of the survey the definition of "needing help with ADLs" is given and understanding is verified. Further, a reminder of this definition appears if respondents move their mouse over the word "ADLs" in the scenario.

- There is a **25%** chance that you **will** need help with ADLs for all of next year.

- There is a **75%** chance that you **will not** need any help at all with ADLs for all of next year.

You have **$100,000** to divide between two plans for the next year. This choice will affect your finances for next year alone. At the end of next year you will be offered the same choice with another $100,000 for the following year.

- Plan C is hypothetical ADL insurance that gives you money if you **do** need help with ADLs.

  - For every $1 you put in Plan C, you will get **$4** to spend if you need help with ADLs.
  - From that money, you will need to pay all your expenses including long-term care at home or in a nursing home and any other wants, needs, and discretionary purchases.

- Plan D gives you money only if you **do not** need help with ADLs.

  - For every $1 you put in Plan D, you will get **$1** to spend if you **do not** need help with ADLs.
  - From that money, you will need to pay for all of your wants, needs, and discretionary purchases.

Presenting the Rules of the Scenario. Immediately after the scenario is presented, the respondents are provided with a recap and elaboration of the specific rules that govern their choice. This recaps the previous screen but is presented in a bulleted, easy to read format. In addition, some features which were hinted at in the first screen, e.g., that there is no public care option and that determination of which plan pays out is made by an impartial third party, are stated explicitly. These rules are designed to ensure that the word problem corresponds as closely as possible to the intended optimal control problem.

- You can only spend money from Plan C or Plan D next year. You do not have any other money.
- If you want to be able to spend whether or not you need help with ADLs, you need to put money into both plans.
- If you need help with ADLs, all money in Plan D is lost.
- If you do not need help with ADLs, all money in Plan C is lost.
- Any money that is not spent at the end of next year cannot be saved for the future, be given away, or be left as a bequest.
- You must make your choice before you know whether you need help with ADLs. Once you make your choice, you cannot change how you split your money.
- Regardless of whether or not you need help with ADLs, your hospital, doctor bills, and medications are completely paid by insurance.
- Other than Plan C, you have no other resources available to help with your long-term care. **You** have to pay for any long-term care you may need from Plan C.

- There is **no public-care option or Medicaid** if you do not have enough money to pay for a nursing home or other long-term care.
- An impartial third party that you trust will verify whether or not you need help with ADLs immediately, impartially, and with complete accuracy.

Verification Questions. In order to reinforce details of the scenario and measure comprehension, we ask the respondents a sequence of questions about the specifics of the scenario, including payoffs in different states, potential uses of money, potential expenses, and rules regarding the payouts. When answering these questions the respondents do not have access to the screens describing the scenario, but have a chance to review the information before retrying any missed questions a second time. If the respondents fail to answer questions correctly a second time, they are presented the correct answers. As documented in Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2016a), the vast majority of individuals answered almost all verification questions correctly before recording their responses to the strategic survey questions.

Figure 7: SSQ 2 Response Screen

Recording the Response. Having reinforced and measured understanding, we are finally prepared to ask the question: how would respondents split their wealth between the two plans? After again presenting them with the original scenario, we present them a screen—a snapshot is shown in Figure 7—with a link in the top right corner to the full scenario. The responses are recorded through an interactive slider that we developed for the purpose of eliciting responses to SSQs. The slider allows the respondent to experiment with different answers and dynamically displays the trade-offs implicit in the SSQs—in this case the trade-off between spending when well and spending when needing LTC. The axis is not labeled with dollar amounts. Instead, the screen contains indications that moving the slider right places more money in Plan D and moving to the

18

	Question	Objective	Scenario Parameters	Preference Parameters
SSQ 1	Lottery over spending	$\lambda^* : \frac{1}{1-\gamma}(W)^{1-\gamma} = \frac{0.5}{1-\gamma}(2W)^{1-\gamma} + \frac{0.5}{1-\gamma}((1-\lambda^*)W)^{1-\gamma}$	(a) $W = \$100K$ (b) $W = \$50K$	γ
SSQ 2	Allocation between ordinary and ADL states	$\max_{z_1,z_2} \ \pi\frac{z_1^{1-\gamma}}{1-\gamma} + (1-\pi)\frac{(\theta_{ADL})^{-\gamma}(z_2+\kappa_{ADL})^{1-\gamma}}{1-\gamma}$	(a) $W = \$100K, \pi = 0.75$ (b) $W = \$100K, \pi = 0.50$ (c) $W = \$50K, \pi = 0.75$	$\gamma, \theta_{ADL}, \kappa_{ADL}$
SSQ 3	Allocation between ADL and bequest states	$\max_{z_1,z_2} \ (\theta_{ADL})^{-\gamma} \ \frac{(z_1+\kappa_{ADL})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma} \ \frac{(z_2+\kappa_{beq})^{1-\gamma}}{1-\gamma}$	(a) $W = \$100K$ (b) $W = \$150K$ (c) $W = \$200K$	$\gamma, \theta_{ADL}, \kappa_{ADL}$ $\theta_{beq}, \kappa_{beq}$
SSQ 4	Indifference between public and private LTC	$W^* : (\theta_{ADL})^{-\gamma} \ \frac{(\psi_G+\kappa_{ADL})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma} \ \frac{(W^*+\kappa_{beq})^{1-\gamma}}{1-\gamma} =$ $(\theta_{ADL})^{-\gamma} \ \frac{(z_1+\kappa_{ADL})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma} \ \frac{(W^*-z_1+\kappa_{beq})^{1-\gamma}}{1-\gamma}$	(a) Public Care Available	$\gamma, \theta_{ADL}, \kappa_{ADL}$ $\theta_{beq}, \kappa_{beq}, \psi_G$

Table 2: **Link between parameters and SSQs:** The first column briefly summarizes the tradeoffs, while the second lists the underlying optimization problem. The third column lists how question parameters were changed for different variations of each SSQ, where W is wealth and $1 - \pi$ is the probability of needing LTC. The z_1 in SSQ 4 is the optimal z_1 function calculated in SSQ 3. The fourth column lists the parameters that determine optimal responses in the model.

left places more money in Plan C. These amounts are displayed dynamically at the ends of the slider. See `http://ebp-projects.isr.umich.edu/VRI/survey_2.html` for an interactive demonstration of the SSQ survey instrument including the slider.

This mode of presentation likely contributes to the high-quality responses we were able to elicit to the SSQs. Use of the slider is effective because it summarizes the key tradeoff in each of the SSQs and allows the respondent to see how the tradeoff operates. The implementation of the slider also addresses several issues with surveys in general. First, there is no value on the slider when it is first presented. The respondent must click to establish an initial value. Second, the survey asks the respondent to move the slider from this initial point (and indeed requires that they do so before recording a response). Together, these two features of the presentation serve to reduce the effect of anchoring. Indeed, we observe little effect of a respondent's first click on their final answer.

After recording a response to this initial question, we ask two variations of this SSQ with different wealth levels and probabilities of needing LTC. This provides further information about how they value having wealth in different states and provides us with a consistency check of individual choices.

4.2 Overview of the Other SSQs

In addition to the SSQ presented above that examines the trade-off between wealth when in need and not in need of assistance with ADLs, the survey presents three other SSQs that examine trade-offs that are relevant to understanding the late-in-life saving motives. Like the previously presented SSQ, these questions present situations which individuals may be unlikely ever to face. Yet, they capture the same type and scale of risks and tradeoffs that individuals do need to confront in making lifetime saving decisions. The choices

that individuals make when confronted with these hypothetical situations provide information regarding the relative values of having wealth in different states of the world. These three additional SSQs are outlined below and the full text of all four SSQs, as well as histograms of the survey responses, are available in "VRI Technical Report: Strategic Survey Questions."

SSQ 1: Risk Aversion.   The first type of SSQ posed is a modified version of the Barsky, Juster, Kimball, and Shapiro (1997) question which examines an individual's willingness to trade a certain lifetime income for a lottery over lifetime income that has a higher expected payoff. Their original question measured tolerance for risk, and has been used frequently to identify the coefficient of relative risk aversion parameter in a power utility function. In the VRI formulation, we refine this question by specifying a more precise environment in which age, health expense, labor income, unexpected expenses, and outside sources of wealth are all controlled for. More generally, the framing of SSQ1 relative to BJKS represents an evolution of thought favoring being more explicit about the precise environment of SSQ scenarios at the expense of having longer and more complex survey questions. We also make the decision a (repeated) static choice, by allowing the individuals to only bet over a single year's spending at one time. This significant departure from the standard BJKS formulation is necessary to avoid confusion with late-in-life health-state utility and bequest preferences.

Specifically, in the VRI question we present individuals an option of choosing between two plans that affect their consumption for the upcoming one year. The first plan guarantees $100,000 for certain, and the second plan will with 50 percent probability double income to $200,000 and with 50 percent probability reduce income by some fraction. The individuals are then asked a series of questions that categorize them into ranges of proportional losses that they would be willing to risk, and then prompted to provide a point estimate of the largest possible loss for which they would choose the risky lottery over the certain income option. SSQ 1 follows BJKS by asking about preference over discrete gambles. At the end of the sequence, SSQ 1 uses plain language to ask respondents to provide the ξ that satisfies

$$\frac{W^{1-\gamma}}{1-\gamma} = 0.5\frac{(2W)^{1-\gamma}}{1-\gamma} + 0.5\frac{((1-\xi)W)^{1-\gamma}}{1-\gamma}, \tag{5}$$

with the choice of ξ point-identifying γ. A variant of this question is then presented, comparing lotteries to $50,000 in certain income.

SSQ 3: Bequest vs. LTC Utility.   In the third type of SSQ that we posed (the second being described in the previous section), we ask individuals to make an irreversible portfolio decision that allocates money between bequests and expenditure while alive when the individuals need help with ADLs. This question, which is similar to one posed in Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011), removes the possibility of an incidental bequest and thus allows us to focus on an intentional bequest motive. Because bequests observed in standard data sources also include unused precautionary savings, it is difficult to identify how strong the bequest motive is. By removing the option of saving money usable for both precautionary and bequest purposes, we are able to separately identify the relative strength of the two motives. To formulate this question, we present individuals with $100,000 and tell them that they have exactly one year left to live. Furthermore, they will need help with ADLs for the entire year. They then must allocate money

between two plans, the first that is available for them to spend during the coming year but can not be left as a bequest, and the second that is only accessible as a bequest upon their death. This response is then recorded and the individuals are asked how their portfolio allocation would change if they had \$150,000 or \$200,000 of wealth. SSQ 3 maps to the following optimization problem:

$$
\max_{z_1, z_2} \quad \frac{(\theta_{ADL})^{-\gamma}(z_1 + \kappa_{ADL})^{1-\gamma}}{1-\gamma} + \frac{(\theta_{beq})^{-\gamma}(z_2 + \kappa_{beq})^{1-\gamma}}{1-\gamma} \tag{6}
$$

$$
\text{s.t.} \quad z_1 + z_2 \leq W
$$

$$
z_1, z_2 \geq 0; \; z_1 \geq -\kappa_{ADL}; \; z_2 \geq -\kappa_{beq}.
$$

SSQ 4: The Value of Public Care. In the final SSQ, we focus on an individual's willingness to utilize public LTC. The environment is similar to that of SSQ 3 in that the respondents are told they only have one year to live, told they will need help with ADLs for the entire year, and the only two spending channels accessible to them are spending on themselves during the year and leaving money as a bequest. In this scenario there is a publicly funded care option that is available to them. Using the public care option will allow them to leave all of their wealth as a bequest, but they will receive the level of care that a typical public care facility would provide. We then ask for the level of wealth at which they would be indifferent between taking public care and paying for their own. Intuitively, for extremely low levels of wealth the respondents are likely to utilize public care, as they are unable to adequately fund their own care and a bequest. For wealth levels sufficiently high, they are likely to fund their own care as the value of public care becomes small compared to the value of private care and the total expenditures on LTC become small relative to their desired bequest level. This suggests there will be an interior response that provides a measure of the equivalent dollar amount an individual assigns to receiving public care. Ultimately, to identify ψ_G SSQ 4 asks respondents to provide the W that satisfies:

$$
\frac{(\theta_{ADL})^{-\gamma}(\psi_G + \kappa_{ADL})^{1-\gamma}}{1-\gamma} + \frac{(\theta_{beq})^{-\gamma}(W + \kappa_{beq})^{1-\gamma}}{1-\gamma} = \frac{(\theta_{ADL})^{-\gamma}(z_1 + \kappa_{ADL})^{1-\gamma}}{1-\gamma} + \frac{(\theta_{beq})^{-\gamma}(W - z_1 + \kappa_{beq})^{1-\gamma}}{1-\gamma},
$$

where z_1 is the optimal policy when there is no public care, as calculated in SSQ 3.

Identification. Because different utility functions are active in different states, different preference parameters control the marginal utility trade-offs that determine the decisions in each state of the world (see Table 2 to see which parameters influence optimal decisions in each SSQ). For instance, SSQ 1 asks individuals to make a risky bet regarding consumption when healthy and explicitly rules out the potential that this decision could influence consumption in other health states. Since relative risk aversion parameter γ is the only parameter which determines marginal utilities in the active states, this question identifies risk aversion. SSQ 2 examines the trade-off between having wealth when healthy and when in need of LTC with ADLs. This trade-off is optimally determined (abstracting from corner solutions for the moment) by equating marginal utility in the healthy state as determined by γ with marginal utility when in need of help, as determined by γ, θ_{ADL}, and κ_{ADL}. Utilizing the observed trade-offs at different wealth levels and state probabilities, we thus are able to identify the θ_{ADL} and κ_{ADL} necessary to align the model with SSQ responses. SSQ 3 examines a similar trade-off between wealth when in need of help with ADLs and wealth for a bequest, while SSQ 4 examines how the existence of a government LTC consumption floor effects the trade-off. In both

of these questions, the model implied optimal strategy is dictated by the marginal value of wealth in the ADL state (again, determined by γ, θ_{ADL}, and κ_{ADL}) and the marginal value of wealth allocated towards a final bequest (determined by γ, θ_{beq}, and κ_{beq}), while in the fourth SSQ the respondent must also take into account how the existence of a public care option affects this trade-off by determining how much he values public care (ψ_G).

4.3 Descriptive Analysis of SSQ Responses

In this section, we seek to describe how the SSQ responses will ultimately inform the formal estimation of preference parameters. We do so by analyzing the histograms of responses to the three variants of SSQ 2 and three variants of SSQ 3, which are the questions most unique to our paper and most informative of the LTC and bequest utility parameters. Appendix B presents the histogram of responses for all SSQ questions.
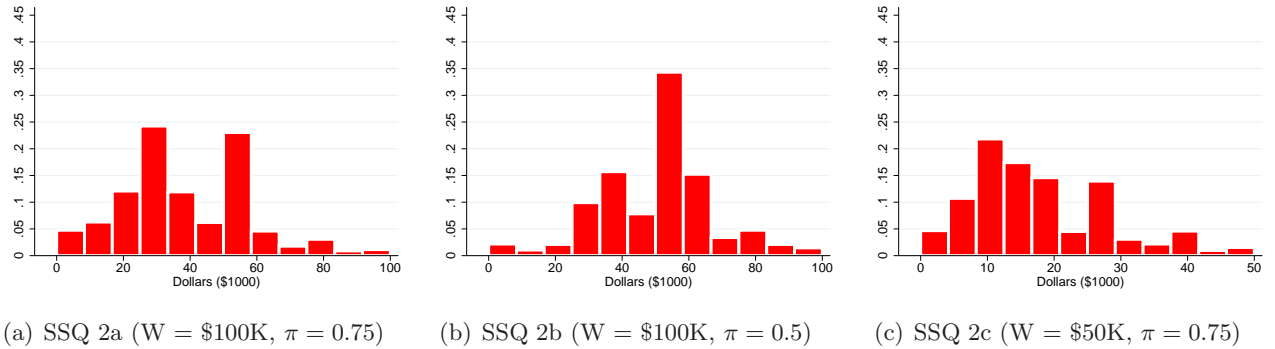

(a) SSQ 2a (W = $100K, $\pi = 0.75$)   (b) SSQ 2b (W = $100K, $\pi = 0.5$)   (c) SSQ 2c (W = $50K, $\pi = 0.75$)

Figure 8: Allocations to the LTC State in SSQ 2 at Different Wealth Levels and Probabilities of Needing LTC $(1 - \pi)$. Since the relative price of the LTC state contingent asset in the survey was set to $\frac{1}{1-\pi}$, respondents would have $\frac{1}{1-\pi}$ times their chosen allocation available for spending in the LTC state.

In Figures 8(a), 8(b), and 8(c) we observe how individuals tradeoff consuming when healthy and consuming when in need of LTC, across wealth levels and different probabilities of needing care. As documented in Figure 8(a), when the probability of needing LTC is 0.25 and people had $100K to allocate between the LTC state and the healthy state, most people allocated between $30-$50K to the LTC state. Since the relative price of the LTC state contingent asset in the survey was set to $\frac{1}{1-\pi}$, respondents would have $\frac{1}{1-\pi}$ times their chosen allocation available for spending in the LTC state. Thus, these choices reflect the respondents desire to have around $120-$200K available when they need LTC and $50-$70K available when they do not need LTC. Figure 8(b) shows the response distribution when the probability of needing LTC is increased to 0.5. In this case, more people allocate around $40-$70K to the LTC state, resulting in around $80-$140K when in need of care and $40-$60K when healthy. Thus, when the probability of needing care is higher, people chose to sacrifice consumption when healthy to keep significant resources available when in need of care. For most respondents, both consumption when healthy and when in need of care are lower, reflecting a tough tradeoff on the margin. That is, respondents do not fix a set amount of wealth in either the healthy or LTC state with the other state receiving the residual. Figure 8(c) plots the responses when people only have $50K in wealth for the year. Most people choose to allocate between $10K-$30K to the LTC state, delivering $40-$120K when in need of LTC and $20-$40K when healthy. Again, even when resources are

22

tight, people are choosing to divert significant funds away from the healthy state towards the LTC state, but are keeping substantial amounts for the healthy state as well. For all SSQ 2 variants, most respondents allocate resources so that they have more available in the LTC state than in the healthy state, often much more. These choices imply that LTC is a high utility state, generating low θ_{ADL} and negative κ_{ADL} in the estimation.

(a) SSQ 3a (W = $100K) (b) SSQ 3b (W = $150K) (c) SSQ 3c (W = $200K)
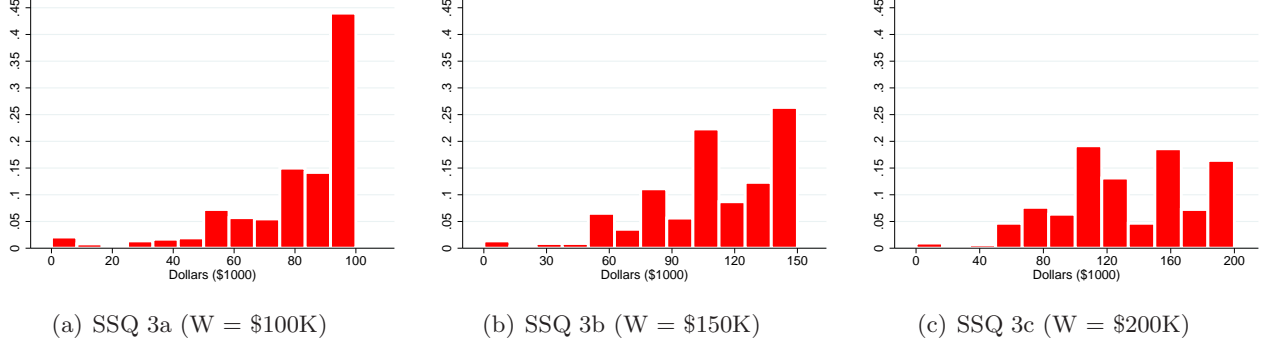
Figure 9: Allocations to the LTC state in SSQ 3 at Different Wealth Levels

In Figures 9(a), 9(b), and 9(c) we observe how individuals trade off leaving money as a bequest and having wealth when in the ADL state, across wealth levels of $100K, $150K, and $200K, respectively, the three variants of the bequest SSQ 3. The histograms show the amount that the individual would allocate towards the ADL state. Here, we clearly see that individuals react to the wealth level. Many respondents allocate almost all of their portfolio to the ADL state when wealth is $100K. When wealth is $150K, it is evident that the wealth constraint in scenario 3a severely restricts spending in the ADL state, as evidenced by the large mass of individuals responding with allocations to the ADL state above $100K. Similar patterns repeat when wealth is $200K, with significant response mass above $150K. Furthermore, even at $200K, many individuals gave zero bequest. This finding suggests that many people view LTC as the primary reason to save late in life at lower wealth levels, with bequest motives becoming more important at higher wealth levels. These responses are consistent with the view that bequests are considered a luxury good while LTC is a necessary good, which should be reflected in estimates of κ_{ADL} and κ_{beq}. It is harder to read directly how these responses translate into θ_{ADL} and θ_{beq}, especially as parameters are jointly estimated using all SSQs. These figures indicate, however, that estimators that target SSQ moments are likely to result in parameter estimates that indicate a strong desire to spend when in need of LTC compared to leaving a bequest. Having provided some descriptive evidence on how the SSQ responses will inform parameter estimates, we turn to the formal estimation strategy.

5   Estimation Methodology

We develop a two stage Method of Simulated Moments (MSM) estimator that is similar to those used in De Nardi, French, and Jones (2010), French and Jones (2011), Lockwood (2016), Gourinchas and Parker (2002), and Laibson, Repetto, and Tobacman (2007) to estimate parameter set Γ. $\Gamma := [\Xi, \Theta]$ is divided into two subsets, with the first subset, Ξ, consisting of parameters externally estimated without the use of the structural model (e.g., income, health transitions, and health costs) and the second parameter subset,

$\Theta := \{\gamma,\ \theta_{ADL},\ \kappa_{ADL},\ \theta_{beq},\ \kappa_{beq},\ \psi_G\}$, consisting of preference parameters that are estimated using moments generated by simulating the structural model.

First, we set some parameters to values informed by the literature and explore robustness to these externally chosen values. We restrict $\beta = 0.97$ and $\omega_G = \$30K$, as the empirical strategy using SSQs was not designed to estimate these parameters. The VRI sample is wealthy enough such that almost no one chooses to go on government welfare when healthy, and thus parameter estimates are identical for a wide range of ω_G.[14]

5.1 Estimated Inputs for the Model

5.1.1 Income

Income profiles are estimated from VRI Survey 1. In Survey 1, respondents report their income flows as the sum of labor income, pension and disability payments, and social security payments. For each age, we assign respondents to an income quintile based on their current rank amongst individuals of the same age. Using this cross-section of income, we use a quintile regression to estimate the age profile of earnings as a polynomial of age and gender. Since life-cycle parameters are estimated off the cross-section, we cannot control for cohort effects. The model groups consumers into five income groups with deterministic age-income profiles determined by the estimated coefficients. This allows us to capture income changes during retirement, but abstract from income fluctuations as a source of uncertainty. The estimated age profiles of income for each quintile are presented in Appendix A.1.

5.1.2 Health Transitions

Health transitions are estimated using HRS waves 2 through 10. We first apply the same sample selection criteria of the VRI to the HRS data to obtain a population with similar observable characteristics. This accounts for the fact that the VRI is wealthier and more educated than the U.S. population and health transition probabilities vary substantially in these dimensions. Then we construct the defined health states $(s = 0, 1, 2)$ from two sets of questions. The first utilizes self-reported subjective health status questions to classify individuals into good or bad health ($s = 0$ or $s = 1$).[15] The second set of questions is used to determine whether an individual is in the LTC/ADL state ($s = 2$). This set of questions presents five activities of daily living and asks whether respondents receive help with any of the five activities. If the respondent answers yes to any of these questions, then we define that respondent to be in the ADL health

---

[14]β is not identified from the SSQs. If β is estimated when targeting wealth moments, it is usually estimated to be close to 1 in order to help the model generate the large degree of wealth present in the upper deciles of the empirical wealth distribution. The effect of this on the other estimated parameters is not large. Robustness to alternative values of β are presented in Section 9. We pre-set $\omega_G = \$30K$ for all exercises because that is the point estimate—although with very large standard errors—when including this parameter to be estimated in our preferred baseline.

[15]Individuals are defined as in good health if they report health being good, very good, or excellent, and are defined to be in bad health if they report health being poor or fair. Self-reported assessments of health status have strong predictive power for future health realizations.

state.[16] Transitions are then estimated using a maximum likelihood estimator, with more information on the estimation methodology and the resulting estimates provided in Appendix A.2.

Given the HRS data, there are two alternative LTC/ADL state definitions that are feasible: needing help with ADLs or being in a nursing home. We prefer to use the "receive help" variable to define the ADL health state for two reasons. First, being in a nursing home is too restrictive to represent the general health status of needing help with ADLs. There are many people who need help with ADLs who do not reside in a nursing home and even when in need of help with ADLs, people may spend on goods other than LTC. Second, it is difficult to interpret a report of needing help, but not receiving help, as really needing help. Thus, we choose the more general needing help with ADLs status, but the more stringent requirement that they actual report receiving help. For consistency, this is the same ADL definition presented in the VRI survey.

5.1.3 Health Expense

The health expense distributions are estimated using the 2010 HRS data. Because we do not allow for persistence in the idiosyncratic cost state in the model, a single year of cross-sectional data is sufficient. We assume the mandatory out-of-pocket health expenditures are distributed log normally and estimate the mean and standard deviation of the lognormal distribution conditional on age, gender, and health state. For the VRI sample studied here, out-of-pocket (post-insurance) health cost shocks are relatively small compared to total wealth and thus do not substantively drive saving in the model. For more information on the estimation of costs and the resulting estimates see Appendix A.2.

5.2 Second Stage Estimates

In the second stage, we apply the MSM estimation procedure. The moments we use to estimate the model are derived from two distinct survey measurements. The first set of moment conditions is derived from behavioral data. We target age-conditional wealth percentiles, which are frequently used to estimate similar life-cycle saving models. A second set of moments is derived from SSQ responses. We first describe the estimator, and then detail the construction of both sets of moments and present the comprehensive moment set we target in our baseline estimation. In Section 6 we present and analyze the resulting baseline parameters. Furthermore, by design of the SSQ questions, it is possible to estimate the model using only wealth or only SSQ moments, which we explore in Section 8.1.

We define $X = (X_i)_{i=1}^I$ as the collection of measurements for all individuals, including behavioral responses, SSQ responses, and state variables; $x_i \subseteq X_i$ is used to denote relevant subsets of individual i's data. The moments used for estimation are the difference between statistics generated by the structural model $(m(\hat{\Xi}, \Theta, X))$ and empirical data $(s(X))$, defined as $g(\hat{\Xi}, \Theta, X) = \mathbb{E}\left[m(\hat{\Xi}, \Theta, X) - s(X)\right]$. We estimate second stage parameters $\hat{\Theta}$ that minimize a GMM quadratic objective function with moments $g(\hat{\Xi}, \Theta, X)$. That

---

[16]Results on saving behavior are robust to using receiving help with at least 2 or at least 3 ADLs as the cutoff for being in the $(s = 2)$ state; the effect of a lower probability of needing help with ADLs and the increased mortality rate associated with the more severe $s = 2$ is offset by the increase in persistence of the $s = 2$ state. Results of these alternative threshold are presented in Section 9.

is, the second stage estimator with weighting matrix W is:

$$\hat{\Theta} = \arg \min_{\Theta} g(\hat{\Xi}, \Theta, X)' W g(\hat{\Xi}, \Theta, X). \tag{7}$$

To estimate the model with the optimal weighting matrix, we use the standard two-step feasible MSM approach. In the first step we minimize the objective function defined in equation 7 using the identity weighting matrix, and denote the minimizing parameter set $\hat{\Theta}_1$. Using this parameter set, we calculate the moment vector $g(\hat{\Xi}, \hat{\Theta}_1, X)$ and the implied first-step covariance matrix $\hat{\Omega}_1$. We use the inverse of the first-step covariance matrix as the second-step optimal weighting matrix $\hat{W} = \hat{\Omega}_1^{-1}$. We then minimize according to equation 7 using $\hat{W}$ to estimate the final parameter set $\hat{\Theta}$. The typical asymptotic properties of the estimator and a derivation of the standard errors are presented in Appendix D.

5.2.1   Wealth Moments

As is common with many life cycle studies of late-in-life savings (e.g., De Nardi, French, and Jones (2010), Gustman and Steinmeier (1986), Lockwood (2016)), the first moment set consists of asset percentiles, conditional on a set of state variables x. The simulated wealth percentile p conditional on state variables x is denoted as $a_x^p(\Xi, \Theta, X)$, while a_i denotes individual i's empirical wealth holdings. The wealth moments conditional on x are:

$$g_{p,x}(\hat{\Xi}, \Theta, X) = \mathbb{E}\left[\mathbb{I}_{\{a_i < a_x^p(\Xi, \Theta, X)\}} - p | x_i = x\right], \tag{8}$$

an expression which can easily be converted to an unconditional expectation through the Law of Iterated Expectations.

For our baseline, we define the moment conditions as the 25^{th}, 50^{th}, and 75^{th} percentiles conditional on age t (a_t^p). In practice, we aggregate the age profiles into disjoint three year intervals, so that $t \in \{55 - 57, 58 - 60, ..., 88 - 90\}$.[17] This aggregation is done to smooth noise in the empirical asset profile that we observe in the cross-sectional data. Given the three targeted wealth percentiles for each of the 12 age intervals, the wealth moment set has 36 moments.

To generate the simulated wealth moments, we first solve the model to compute optimal decision rules as a function of all relevant state variables given parameters. We simulate the model for a large number of individuals (N) and then compute moments using the simulated behavior. To set the initial conditions for the simulation, we sample (with replacement) N individuals from the VRI data X, in which an individual is characterized by his or her vector of idiosyncratic state variables. Then, for each individual, we draw relevant shocks from the $\hat{\Xi}$ parameterized stochastic processes and simulate the behavior implied by the computed optimal policies (given parameters Θ). We then aggregate these individual behaviors to construct the simulated population moments $m(\hat{\Xi}, \Theta, X)$.

---

[17]We stop matching wealth moments when the sample for the age group becomes too small, defined as less than 20 individuals. This first occurs in the 91-94 age bin. Results are not sensitive around a reasonable range of the chosen cutoff.

5.2.2 Strategic Survey Question Moments

The second moment set is constructed entirely from SSQ responses. The structural model has 6 preference parameters to be estimated: relative risk aversion parameter, γ; LTC state utility parameters, θ_{ADL} and κ_{ADL}; bequest utility parameters, θ_{beq} and κ_{beq}; and public-care aversion, ψ_G. Using responses to the four types of SSQs described in Section 4.2 and their iterations at different wealth levels and state probabilities, we are able to identify these six parameters. Since the survey poses two variants of SSQ1, three variants of SSQ 2, 3 variants of SSQ 3, and one variant of SSQ 4, and we are estimating 6 preference parameters, using SSQ data alone we have an over-identified system suitable for estimation of these six preference parameters.

As moments, we match the model implied allocations to the empirical mean of each of the 9 SSQ variants, indexed by $m \in \{1, 2, ..., 9\}$. Let z_m^i be individual $i's$ response to SSQ variant m. The theoretical counterparts, the optimal SSQ answers conditional on preference parameters, are denoted as $s_m(\Theta)$. We then write each moment as:

$$g_m(\hat{\Xi}, \Theta, X) = \mathbb{E}\left[s_m(\Theta) - z_m^i\right], \tag{9}$$

providing us with 9 SSQ moments.

5.2.3 Combining Wealth and SSQ Moments

In the previous two sections we described the specification of wealth and strategic survey moments. In this section we describe our baseline estimation procedure that uses both SSQ and wealth data by combining these moments into a single moment vector that will be used to estimate the model. Utilizing both sources of information disciplines the estimator to match wealth data and SSQ data, with each source of data likely containing unique information. To combine the two sources of data, we concatenate the wealth moments formed from matching the cross-sectional wealth distribution percentiles and the SSQ moments formed by matching the empirical SSQ responses, resulting in a total of 45 moment conditions.

Letting $g_V(\hat{\Xi}, \Theta, X)$ denote the set of moments constructed from wealth data and $g_S(\hat{\Xi}, \Theta, X)$ denote a set of moments constructed from SSQ data, the baseline joint estimation's moment set is

$$g_J(\hat{\Xi}, \Theta, X) = \begin{bmatrix} g_V(\hat{\Xi}, \Theta, X) \\ g_S(\hat{\Xi}, \Theta, X) \end{bmatrix}. \tag{10}$$

By design, SSQ moments provide strong identification of the preference parameters, and thus the optimal weighting matrix assigns a disproportionate amount of weight to the SSQ moments, resulting in tight match of SSQ responses but a slight underprediction of wealth accumulation. In order to better match these wealth moments, which are the common target of other studies and helps facilitate comparison to the literature, we make use of a parameter, λ, that allows us to control the relative importance we assign to each moment set, $g_V(\hat{\Xi}, \hat{\Theta}, X)$ and $g_S(\hat{\Xi}, \hat{\Theta}, X)$. By over-weighting the wealth moments relative to the statistically optimal weighting matrix, we arrive at our baseline estimator.[18]

---

[18]See Appendix C for results using the optimal weighting matrix. Baseline estimates are quantitatively similar to those using the optimal weighting matrix. This similarity holds for a wide range of λ, suggesting the implications of preferences for saving, expenditure, and bequests are robust to this choice.

6 Parameter Estimates

6.1 Model Fit

(a) 75p, 50p, and 25p Wealth Moments in Model (dashed) and Data (solid)

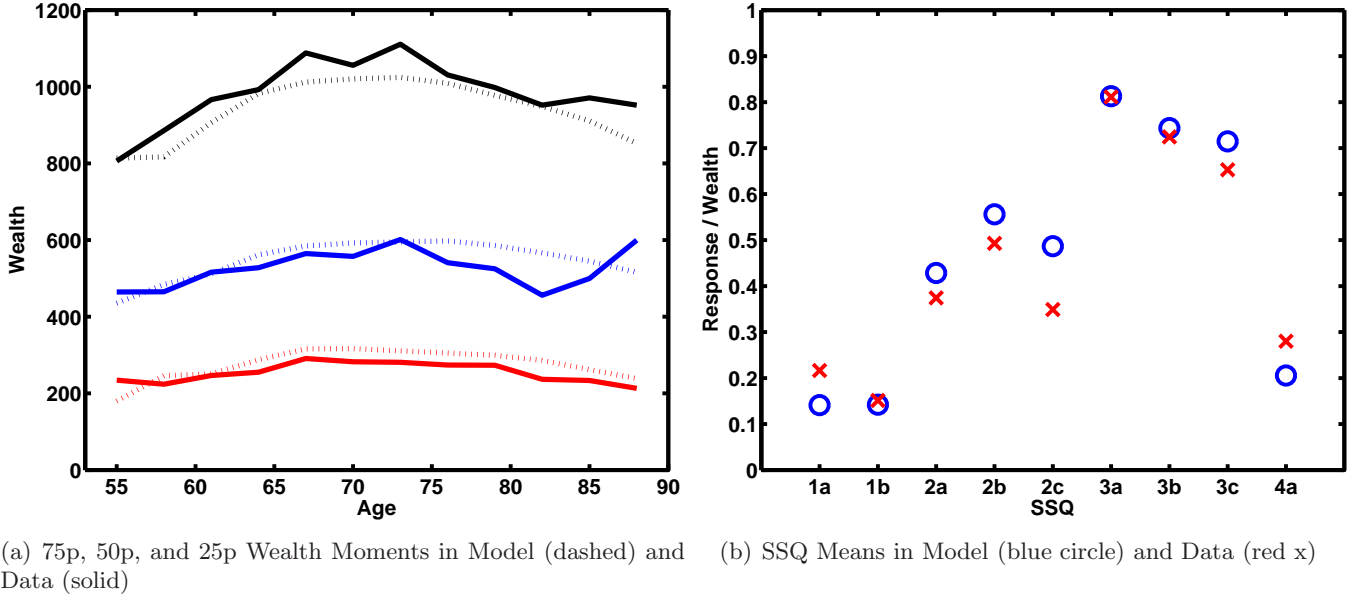(b) SSQ Means in Model (blue circle) and Data (red x)

Figure 10: Model Fit When Jointly Targeting Wealth and SSQ Moments

Before presenting the parameter estimates, we show the fit of the model to the data. Figures 10(a) and 10(b) document the model fit from our baseline estimation that jointly targets both wealth and SSQ moments. For exposition, we present the SSQ moments on a $[0, 1]$ scale, by normalizing the mean response by the maximum possible response.[19] Overall, the fit is good, as the percentiles of the wealth distribution are matched well, as are the SSQ moments. We will first analyze the baseline estimates and then analyze the distinct contributions of the SSQs and wealth moments in Section 8.

6.2 Baseline Estimated Parameter Values

Table 3 presents the parameter estimates for our baseline case that combines the wealth and SSQ moments. Of particular interest in the estimation are four groups of parameters. First, the estimate of risk aversion (γ) is of independent interest, but is also important because it is jointly estimated and strongly affects the estimated values of other preference parameters. Second, the estimates of health state utility (θ_{ADL} and κ_{ADL}) presented are the first we are aware of to estimate this functional form applied specifically to the LTC health state. Third are the bequest parameters θ_{beq} and κ_{beq}. Finally, ψ_G controls the degree to which there is public-care aversion.

The coefficient of relative risk aversion is estimated to be 5.27. This value is somewhat larger than that typically used in the literature and somewhat smaller than that typically estimated using similar survey techniques. For example, in an exercise using a similar model and very different data De Nardi, French, and

---

[19]Because there is no maximum response to SSQ 4, we set W_m to be the windsorized 95^{th} percentile of the raw response distribution.

Table 3: Parameter Estimates

	Joint Estimation: Baseline Model				
γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
5.27	0.67	-37.44	1.09	7.83	77.43
(0.05)	(0.37)	(0.22)	(0.25)	(0.48)	(9.49)

Standard errors are reported in parentheses.

Jones (2010) estimate $\gamma = 3.8$, while Kimball, Sahm, and Shapiro (2008) estimate a mean γ of 8.2 when using responses in the HRS to hypothetical lotteries over future income streams.

In examining the estimated preferences when targeting the joint moments, it is striking that the health-state utility function implies very different marginal utility of wealth in the state of the world when an individual needs LTC and when LTC is not needed. The estimated $\kappa_{ADL} = -37.44$ is negative and large, suggesting LTC is viewed as a necessary good with a spending floor of about \$37K.[20] $\theta_{ADL} = 0.67$ implies a high marginal utility of expenditure in the LTC state, especially when combined with κ_{ADL}. Note that if $\kappa_{ADL} = 0$, relative to ordinary consumption, $\theta_{ADL} > 1$ would imply that LTC expenditures provide less marginal utility for each dollar spent while $\theta_{ADL} < 1$ would imply a higher marginal utility for each dollar expenditure. Individuals with $\theta_{ADL} > 1$ and $\kappa_{ADL} < 0$ would view expenditure when in need of LTC as a strong necessity and so would optimally desire to consume the necessary amount, but not much more. Thus, the effect of precautionary savings on wealth accumulation would be similar to medical expenses such as those presented in De Nardi, French, and Jones (2010). For individuals with $\theta_{ADL} < 1$, marginal expenditures in the LTC state would have a high value, so there is a motive to have additional spending when in need of care even at expenditures above the necessary amount.

Compared to utility from expenditure in the LTC state, individuals receive less utility from leaving a bequest. First, $\kappa_{beq} = 7.83$, which suggests that bequests are viewed as a luxury good. Together with a positive κ_{beq}, the estimate of $\theta_{beq} = 1.09$ implies a low marginal utility of bequests, relative to LTC. This is very different from estimates in the literature like those from Lockwood (2016) and Koijen, Van Nieuwerburgh, and Yogo (2016), who find much stronger bequest motives, as discussed in Section 8.2.

Because expenditure when in need of LTC is so highly valued, $\psi_G = \$77,430$ is estimated to be large compared to previous estimates in the literature. To facilitate comparison to the literature, this estimate implies an equivalent utility level to that provided by a \$24,395 government provided public-care expenditure in a model with the same risk aversion but without state dependent preferences.[21]

7   Behavioral Implications of Estimated Preferences

To help further interpret the parameter values in economically meaningful ways and to investigate the contribution of the different saving motives in determining life-cycle wealth patterns, in this section we use

---

[20]In our baseline, we set the minimum private LTC expenditure $\chi = \$40,000$. The estimation results are independent of the choice of χ for any $\chi \leq -\kappa_{ADL}$ and χ is quantitatively irrelevant when greater than but close to $-\kappa_{ADL}$ as is the case here.

[21]To calculate this expenditure equivalent in a model without the health state utility function, we find the expenditure level $\bar{\psi}$ that would equate utility across the two specifications: $\frac{\bar{\psi}^{1-\gamma}}{1-\gamma} = (\theta_{ADL})^{-\gamma} \frac{(\psi_G + \kappa_{ADL})^{1-\gamma}}{1-\gamma}$.

the model to document behavior induced by the estimated preferences. We first do so in a simple static optimization problem to give a sense of the curvature of the estimated utility functions, before turning to an examination of expenditure in the estimated full model.

7.1 An Illustrative Synthetic Static Choice Problem

As a first pass on interpreting the estimated parameters, we plot optimal allocations in a simple synthetic static choice problem, in which an individual regards consumption when healthy, when in need of LTC, and bequests as three different "goods" that can be purchased simultaneously at a relative price of 1 and are valued according to the corresponding estimated utility functions. This problem is not one that maps to a situation ever faced by individuals; for example, there is no probability weighting to reflect the likelihood of being in each state. It is nonetheless a convenient illustration device to present the marginal utility of expenditures by health status implied by the estimated parameters before turning to the full model that incorporates dynamics and risks. Specifically, the static allocation problem for a person with wealth W is:

$$\max_{z_1, z_2, z_3} \frac{(z_1)^{1-\gamma}}{1-\gamma} + \frac{(\theta_{ADL})^{-\gamma}(z_2 + \kappa_{ADL})^{1-\gamma}}{1-\gamma} + \frac{(\theta_{beq})^{-\gamma}(z_3 + \kappa_{beq})^{1-\gamma}}{1-\gamma} \tag{11}$$
$$\text{s.t. } z_1 + z_2 + z_3 \leq W$$
$$z_1, \ z_2, \ z_3 \geq 0; \ z_2 \geq -\kappa_{ADL}; \ z_3 \geq -\kappa_{beq}.$$

Recall that this synthetic choice models is meant to summarize the parameters implications succinctly, not to provide quantitative predictions reflecting an actual choice problem. The main advantage of this simple choice problem is transparency, as the interaction between motives that is expressed in the full dynamic model makes it challenging to easily determine the relative strength of each motive.
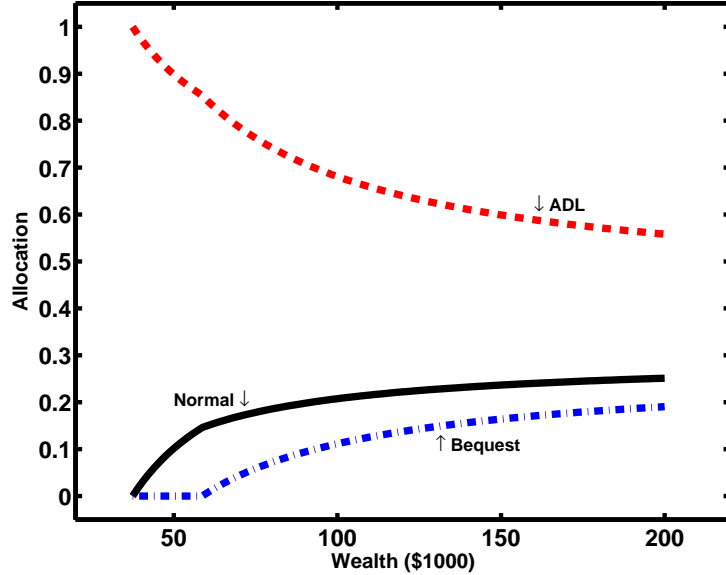

Figure 11: Expenditure in Synthetic Static Choice Problem

Figure 11 plots the resulting optimal allocations. The most striking feature of this figure is that ADL-

state expenditure is allocated the majority of wealth up to rather high wealth levels. Because ADL-state expenditure is a necessary good at least $-\kappa_{ADL}$ dollars are always allocated to this motive, giving rise to an initial ADL-allocation share nearly equal to 1 that falls as wealth is allocated to other expenditures to take advantage of diminishing marginal utility from any one type of expenditure. In addition, we see that because bequests are estimated to be a luxury good, there is initially no spending on bequests until wealth is sufficiently high that the bequest motive becomes active. The effect on expenditure of the non-homotheticity parameter κ diminishes at higher wealth levels and the differences in θ become more dominant. Even at high wealth levels, people spend the most on the ADL-state good, then the ordinary-state good, followed by a bequest. Although the outcome of a simple synthetic static optimization problem, these results highlight that the estimated utility functions reflect a strong desire to spend when in need of help with LTC and the relatively weak bequest motive that will drive life-cycle saving behavior in the full dynamic model with health and mortality risk.

7.2   Implied Expenditures using the Baseline Model

In this section, using the full estimated model, we analyze expenditure behavior over the life cycle and in the cross section.[22] The simulated expenditures show the implications of the parameter estimates on an important non-targeted dimension and provide credibility for the quantitative behavior of the model.

We begin by documenting the cross-section of expenditure across health status in Figure 12. The heterogeneity in expenditure comes from differences in income, wealth, age, gender, and the realization of health shocks. When in good health, the modal expenditure is around $40K, with a long upper tail. As alluded to when analyzing the estimated parameters, spending when in need of LTC is higher than when healthy. When in need of LTC, over 60 percent of people spend around $80K to $100K. Furthermore, there are around 15 percent of people spending over $150K when in need of LTC and 5 percent spending over $200K.

Although modeled expenditure reflects total expenditure when in the ADL state, to get a sense of comparison to the data, we can compare to nursing home costs as reported in Genworth (2016). There is a wide range nursing home costs; an annual stay in a private room in a nursing home averages $92K and ranges between $60K to $250K, reflecting varying costs and qualities. These numbers align well with expenditure in the model.

Figure 13 presents the distribution of expenditure by health status conditional on age. The distribution of expenditure for people in good health is near constant across age. The 90th percentile is around $100K per year, while the 10th percentile is around $30K. Expenditure when in need of LTC is also fairly constant across age, with substantially higher skewness at all ages and more variation in the upper percentiles. These expenditure patterns are consistent with the interpretation that spending when in need of help with ADLs is viewed as a necessary good with an additional intensive margin valuation.

Figure 14 puts this all together by plotting the share of total expenditure for each age by ADL-status. This calculation combines the fraction of people with each health status and the distribution of expenditures by health status. The most dramatic change over the life cycle is the decrease in the share of expenditure by those in good health and the increase in the share associated with those needing help with ADLs at older ages. At younger ages, very few people need help with ADLs, and it is not until around age 80 when

---

[22]Expenditure is defined as $c + h$, summing all voluntary expenditure and the mandatory out-of-pocket health cost shocks.
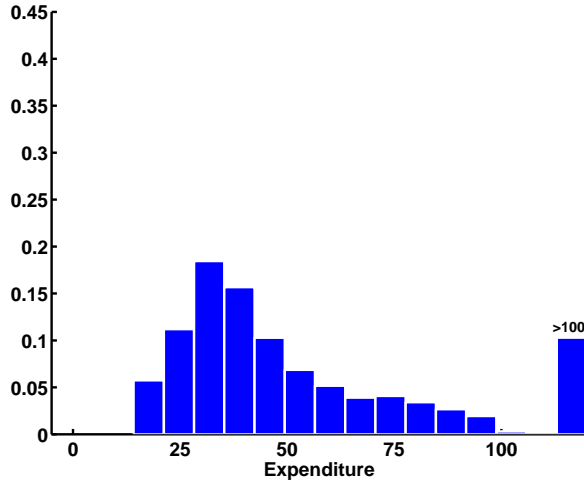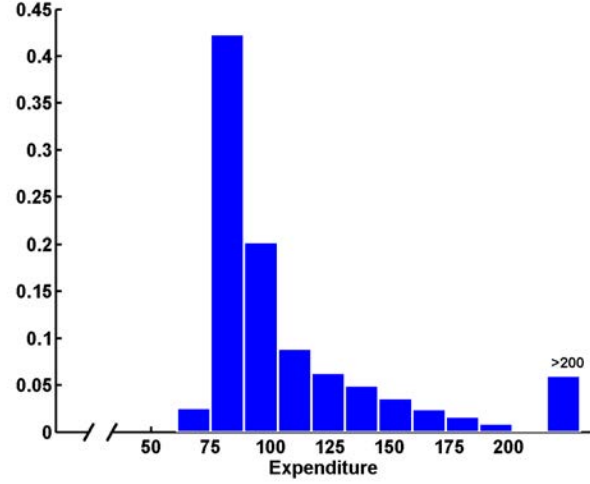
Figure 12: **Distribution of Expenditure**: This figure plots the cross-section of expenditure for those in good health and for those in need of help with ADLs.
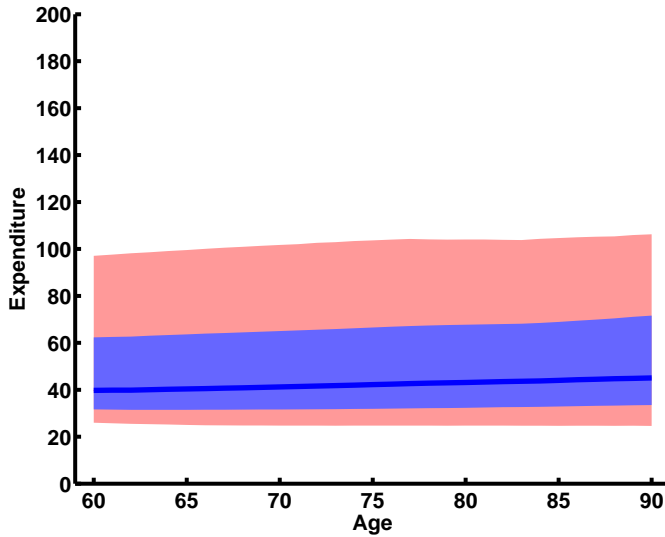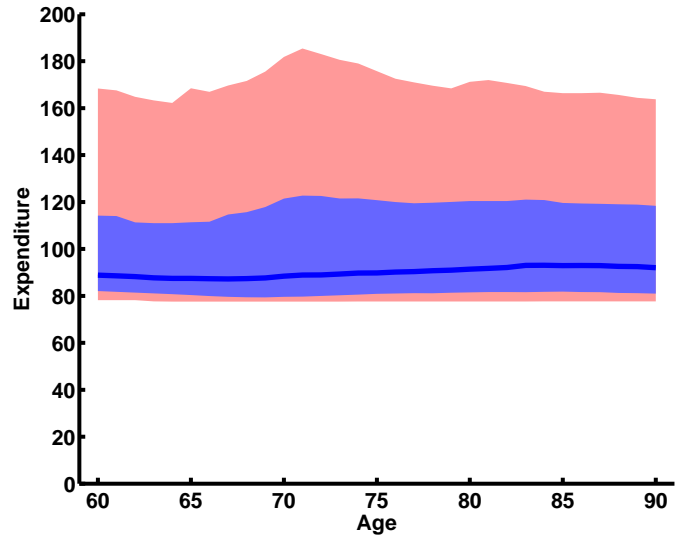
Figure 13: **Distribution of Expenditure by Age**: This figure plots expenditure c by age for those in good health and for those in need of help with ADLs. The blue line is median expenditure, the blue area marks the 25th-75th percentiles and the pink area contains the 10th-90th percentiles.

those people account for over 10 percent of total expenditure. After age 80, however, the expenditure share associated with those needing LTC increases rapidly. The probability of needing care increases at these ages and the expenditure by those needing care is significantly higher than those in good health. At age 90 those needing help with ADLs account for about one third of expenditures. This shows that risks and associated preferences combine to make LTC a significant driver of saving for the elderly. Taken together, these figures show that the estimated model predicts intuitively reasonable expenditure patterns in the cross-section and

over the life cycle, at the individual and population levels.

Figure 14: Share of Expenditure by Age and Health Status

Finally, to document the relative strength of the different saving motives, we present in Table 4 percentiles of the distribution of bequests under alternative parameter values. Bequests are a useful statistic to capture saving over the life cycle, since they are the result of a life time of asset accumulation. The differences presented in Table 4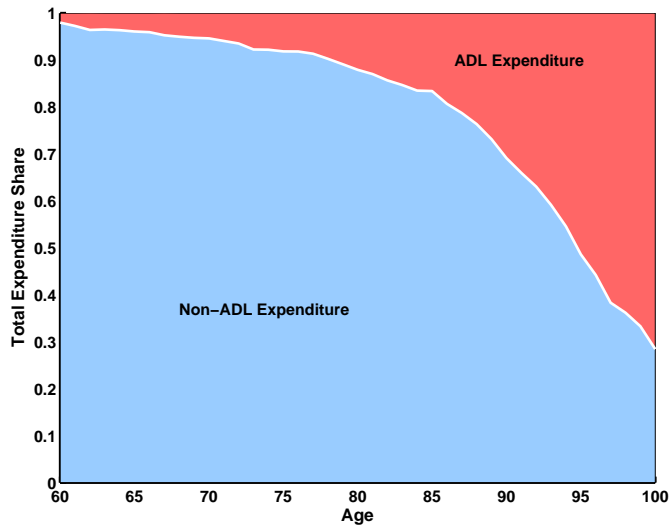 reflect the relative strength of the LTC and bequest utility functions estimated using the wealth and SSQ data using the full model, as opposed to the synthetic static choice problems featured earlier. The first row shows bequests under the baseline parameterization, the second row shows bequest when there is zero utility from leaving a bequest, the third row shows bequests when utility when in need of help with ADLs is equal to utility when healthy, and the fourth row combines these alternatives by showing bequests when there is no bequest utility and no health-state dependent utility.

Comparing these parameterizations provides insight on the contribution of the various saving motives to overall late-in-life wealth accumulation. In the baseline model, realized bequests reflect an active desire to leave a bequest and an incidental bequest driven by precautionary saving combined with mortality risk. These motives interact to reinforce each other: with strong bequest motives, the cost of ex-post oversaving for the ADL health state is lower and strong ADL-state utility induces high saving rates that may result in a large incidental bequest. Because of this interaction this exercise does not provide a strict decomposition. Nonetheless, the magnitude of the differences are informative of the strength of the motives. These interdependencies highlight the benefit of SSQs that can help to separately identify motives in light of these interactions.

When shutting off the bequest motive, the average bequest is two-thirds its baseline level, suggesting that the majority of saving over the life cycle is driven by precautionary motives related to health and death and that most bequests are incidental as opposed to being driven by an active warm-glow bequest margin. Furthermore, average bequests are larger in the model with no bequest motive than in the model with no health-state dependent utility. Even absent the bequest motive that reduces the cost of ex-post oversaving, precautionary saving related to health generates more than half of baseline bequests. Absent the bequest motive, the LTC motive alone makes bequests 43 percent larger than would longevity uncertainty alone. On

33

average, the model without either the bequest motive or the health-state utility function generates just under half of the baseline bequest, indicating that longevity risk is a significant contributor to realized bequests, but not the whole story quantitatively. There are also different responses to the alternative parameterizations across the bequest distribution. The median bequest is lower in the case in which there is no bequest utility compared to when there is no health-state dependent utility, while the 75th and 90th percentiles of the bequest distribution are larger. This pattern arises because the means-tested public care option dissuades lower-wealth individuals from saving for LTC reasons and because the health-state utility encourages rapid dissaving when in need of care towards the end of life that is subject to diminishing within-period returns. Overall, Table 4 shows that precautionary motives contribute substantially to wealth accumulation over the life cycle and that health-state dependent utility related to long-term care needs is a quantitatively significant determinant of saving patterns at older ages.

Table 4: Decomposing Saving Motives: Bequests Under Alternative Preferences

	mean	10p	25p	50p	75p	90p
Baseline	505	0	157	455	707	1,024
Health State Utility Only	336	0	0	158	530	891
Bequest Utility Only	302	0	59	200	408	677
No Bequest or Health State Utility	235	0	0	74	340	638

This table presents percentiles of the bequest distribution in thousands of dollars under different parameterizations of health-state dependent and bequest utility functions.

7.3   Fit of Non-Targeted Wealth Moments

To demonstrate the saving behavior across the wealth distribution implied by the baseline estimates, Figure 15 plots the model fit to non-targeted wealth moments. Even though the 10th percentile was not targeted in the estimation, the model fit to data is almost exact. Preference parameters estimated by targeting the higher wealth moments successfully replicate savings patterns of the less affluent. The model does a good job of "out of sample" prediction for low levels of wealth, which corresponds to a large fraction of the U.S. population who may be the target of policy changes and innovations to promote retirement saving. The use of SSQs permits predictions over households who do not—because of their level of financial resources or institutional and market constraints—face the same economic trade-offs as members of the VRI sample. We explore this further by analyzing model predictions for the HRS sample in Section 9.2. In contrast, the model predicts less wealth than appears in the data at the 90th percentile, typically underpredicting wealth by around 10 to 15 percent. Our interpretation is that the model is missing certain features that are particularly apt for the wealthiest members of the VRI sample, which itself oversamples high wealth individuals relative to the U.S. population.
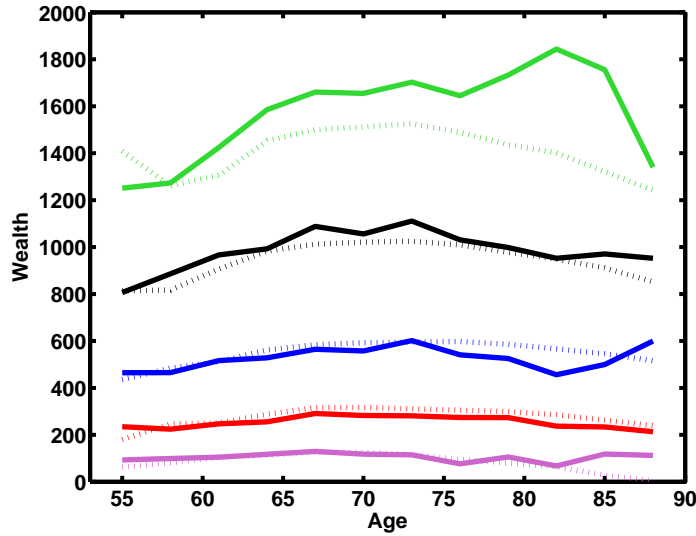
Figure 15: 90p, 75p, 50p, 25p, and 10p Wealth Moments in Model (dashed) and Data (solid)

8   Comparison to Alternative Parameter Estimates

8.1   Matching only SSQ or only Wealth Moments

In our baseline estimation, we target both wealth moments traditionally used in this literature and our newly-created SSQ moments. In order to disentangle the relative contribution of the SSQs in the estimation results, in this section we present and analyze the resulting estimated preferences when we target wealth moments or SSQ moments exclusively. We are able to use the SSQ moments exclusively because we have collected sufficient non-behavioral data to identify the preference parameters solely from the SSQ responses. Although our preferred estimates use both wealth and SSQ data, we do not need to augment the non-behavioral data with observed behaviors to gain identification. We are unaware of any existing study that has undertaken a similar estimation of a structural life cycle model without requiring behavioral data for parameter identification.

Although it is difficult to compare across estimates that target different moments, some broad messages are clear. Every set of estimates suggests that the marginal utility of expenditures when in need of LTC is larger than that from bequests. κ_{ADL} is always estimated to be negative and large in magnitude. Often θ_{ADL} is smaller than θ_{beq}, and in the SSQ-moment case when they are almost equal, the large negative κ_{ADL} dominates the positive and small κ_{beq}. Thus, it is not just because we target SSQs that we find the importance of the precautionary saving motive induced by LTC risk. Hence, many features of the data strongly support allowing for a non-homothetic LTC health-state utility function, that spending when in need of LTC is viewed as a necessity and it is highly valued on the margin. Moreover, the SSQs responses imply that the bequest motive is not as strong as the wealth data alone suggests.

8.1.1   Model Fit

The first set of figures documents model fit when only wealth moments are targeted. Figure 16(a) presents the model-generated and empirical 25th percentile, 50th percentile, and 75th percentile of the wealth distribution across ages. The model-generated data matches the cross-sectional wealth distribution well for most ages.

(a) 75p, 50p, and 25p Wealth Moments in Model (dashed) and Data (solid)

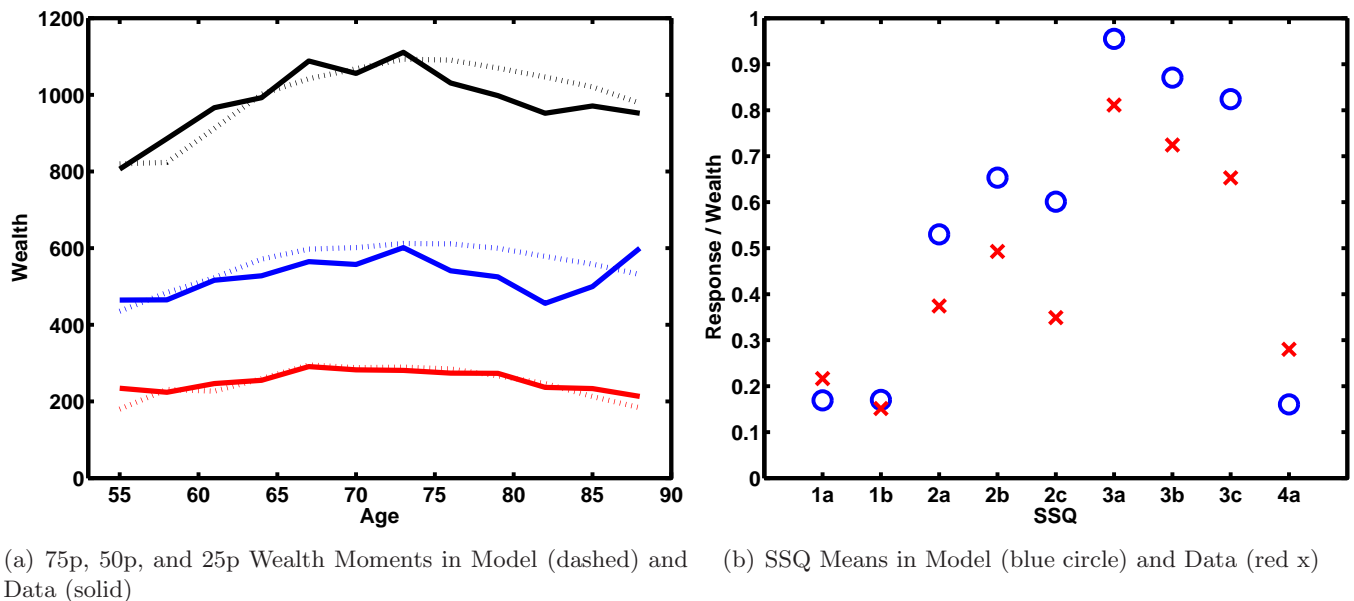(b) SSQ Means in Model (blue circle) and Data (red x)

Figure 16: Model Fit When Exclusively Targeting Wealth Moments

As displayed in Figure 16(b), the parameters that best match the wealth moments generate SSQ moments that are somewhat distant from those measured in the data. Roughly speaking, the resulting parameters suggest individuals that are more risk averse and have stronger desires to spend when in need of LTC and to leave bequests than the SSQ data imply.

Now consider targeting the SSQ moments only. As can be seen in Figures 17(a) and 17(b), the SSQ moments are hit almost exactly. The success in fitting these moments should not be surprising, given that SSQs were designed to ensure identification. When only targeting the SSQ moments, the fit to the wealth moments deteriorates, with the model predicting undersaving for the 75th percentile and over saving for the 50th and, especially, the 25th percentile.

An important consideration is that in our wealth estimation and in other studies that target wealth moments, a model's ability to match wealth moments does not necessarily mean that preferences and all associated savings motives are well identified. Including SSQ moments introduces more information on saving motives and disciplines the channels through which a model can match wealth moments. The tension in matching both datasets simultaneously highlights both the need for further model development (as we have started here by introducing the non-homothetic health-state utility function) and that the information contained in the SSQs provides a source of identification for richer models.

8.1.2 Estimated Parameter Values from Alternative Moments

Table 5 documents the parameter estimates that result from exclusively targeting wealth or SSQ moments. It is important to note that, due to the non-linearity of the model, it is not necessary that this estimation procedure yields estimates that are a convex combination of or are contained within the interval defined by estimates that result from exclusively targeting one type of data (i.e., using only wealth or only SSQ moments).

Compared to baseline, when matching the SSQ moments alone the estimation procedure delivers a
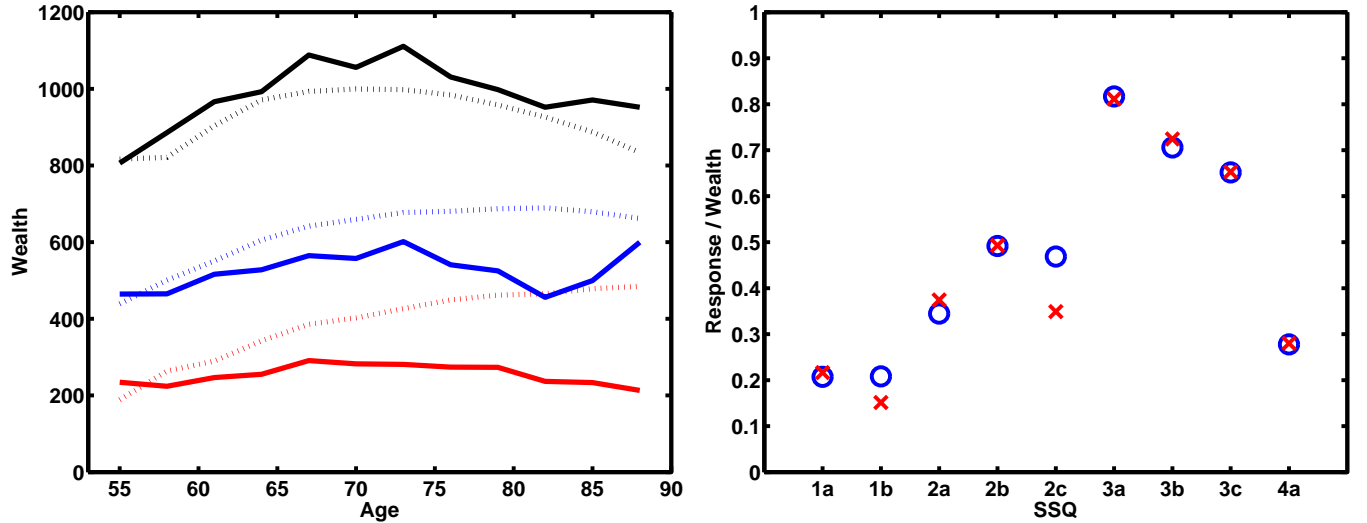
(a) 75p, 50p, and 25p Wealth Moments in Model (dashed) and Data (solid)

(b) SSQ Means in Model (blue circle) and Data (red x)

Figure 17: Model Fit When Exclusively Targeting SSQ Moments

Table 5: Estimated Parameters: Alternative Estimates

Joint Estimation: Baseline Model					
γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
5.27	0.67	-37.44	1.09	7.83	77.43
(0.05)	(0.37)	(0.22)	(0.25)	(0.48)	(9.49)

Wealth Moments Only					
γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
4.41	0.51	-53.12	1.17	14.06	79.71
(3.85)	(8.86)	(4.75)	(1.07)	(7.26)	(7.18)

Strategic Survey Questions Only					
γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
3.56	1.77	-63.97	1.61	0.34	91.682
(0.03)	(0.18)	(1.32)	(0.14)	(1.73)	(2.61)

This table presents parameter estimates for the estimation targeting jointly both sets of moments, targeting wealth moments only, and targeting SSQ moments only. Standard errors are reported in parentheses.

lower γ and higher θ_{ADL} and θ_{beq}. The estimates targeting wealth moments only result in a higher risk aversion coefficient and a higher marginal utility when in need of long-term care, generated by both lower θ_{ADL} and more negative κ_{ADL}. In the wealth estimation we find that most parameters are estimated with very large standard errors, which reflects the fundamental difficulty in identifying richer preference parameters from wealth data alone. This difficulty was discussed extensively in Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011), and commented on by De Nardi, French, and Jones (2010) in discussion of their Table 2. Additionally, because we only have a cross-section, we cannot control for cohort effects. The

use of SSQs allows us to design questions that overcome these difficulties and facilitate identification.

Comparing Parameters Estimated to Match Alternative Moments. It is difficult to compare the economic interpretation of estimated parameters in isolation because the value of γ affects the interpretation of the θ and κ parameters, just as the value of κ_{ADL} affects the interpretation of the value of θ_{ADL}. As a result, we present the following exercise, which illustrates the relative strength of the different expenditure motives induced by the different parameter values. To compare across sets of parameters we return to the simple synthetic static choice problem presented in equation 11. Figures 18(a), 18(b), and 18(c) plot the resulting optimal allocations for the parameter sets that result from targeting different moments.

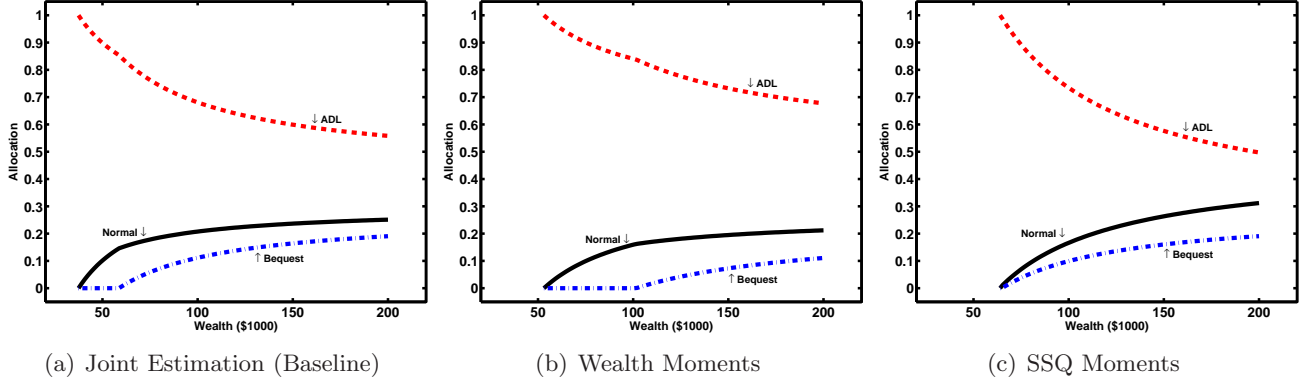(a) Joint Estimation (Baseline) (b) Wealth Moments (c) SSQ Moments

Figure 18: Expenditure Fractions Across Parameter Estimates

Across all estimates, LTC expenditure is allocated the majority of wealth. In comparing across the specifications, several patterns are observable that reflect differences in parameter estimates. We observe that the estimates that exclusively target wealth moments indicate that bequests are more of a luxury good, and thus receive a lower share of allocated wealth. In addition, these wealth estimates indicate a very high marginal value of wealth when in the LTC state, which is reflected in panel 18(b) by the slow decline in LTC expenditure as wealth increases. The SSQ estimates indicate high θ_{ADL} and θ_{beq}. The SSQ estimates imply a low marginal value of wealth in these states, and is reflected by the steeper allocation profile of normal consumption in panel 18(c). Finally, note that the lines in panel 18(a) generally fall between the corresponding lines in the other two panels. Although this does not hold at lower levels of wealth exactly, this broad pattern reflects the trade-off between matching the two sets of moments in the joint estimation. In sum, the SSQ-only, wealth-only, and baseline estimation all result in estimates that indicate a strong desire to spend when in need of help with activities of daily living.

8.2 Comparison of Preferences Across the Literature

Different preference parameters can induce radically different saving and spending behavior. To demonstrate the degree of difference and to illustrate the large impact of these differences on implied saving behavior, we contrast the implications of the estimated baseline parameters with those in the literature.

In Figures 19(a) and 19(b) we compare the relative strength of these motives induced by the parameters we have estimated to those estimated in De Nardi, French, and Jones (2010) (DFJ) and Lockwood (2016), as these are some of the most closely related papers to ours. We translate the parameters from other papers
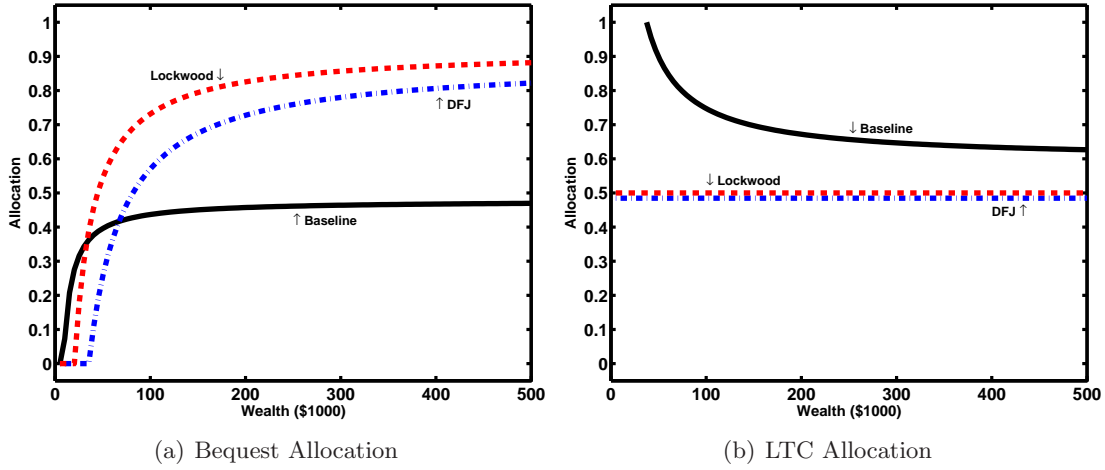
(a) Bequest Allocation	(b) LTC Allocation

Figure 19: Expenditure Levels Across Studies

to be compatible with our utility function specification, yielding DFJ parameters of $\gamma = 3.81$, $\theta_{ADL} = 0.79$, $\kappa_{ADL} = 0$, $\theta_{beq} = 2360$, and $\kappa_{beq} = 273$ and Lockwood parameters of $\gamma = 3$, $\theta_{ADL} = 1$, $\kappa_{ADL} = 0$, $\theta_{beq} = 1460$, and $\kappa_{beq} = 231$. In a similar analysis to that presented in Figure 18, in Figure 19(a) we present the optimal bequest allocation (z_2) from the following optimization problem when calibrated according to each paper's baseline estimate of risk aversion and bequest parameters:

$$\max_{z_1, z_2} \quad \frac{(z_1)^{1-\gamma}}{1-\gamma} + \frac{(\theta_{beq})^{\gamma}(z_2 + \kappa_{beq})^{1-\gamma}}{1-\gamma} \tag{12}$$
$$\text{s.t.} \quad z_1 + z_2 \leq W$$
$$z_1, \ z_2 \geq 0; \ z_2 \geq -\kappa_{beq}.$$

We observe that relative to these other studies, in this paper the bequest motive is estimated to be somewhat stronger at low levels of wealth, but much weaker at high wealth levels. Whereas these studies suggest a steadily increasing allocation to bequests even for wealth levels above \$400,000, our estimates suggest the allocation share is relatively stable for wealth levels above \$100,000. Thus, in our estimation bequests are considered less of a luxury good compared to both DFJ and Lockwood. Furthermore, for high levels of wealth, the bequest share asymptotes to 48% in our baseline, compared to 89% for DFJ and 92% for Lockwood. These differences reflect that we estimate a much lower bequest multiplier and hence a lower marginal value of bequests at high wealth levels than either other study. This further demonstrates that in our model most of bequeathed wealth is given incidentally rather than as a result of a bequest motive *per se*.

DFJ show that health risks play an important role in determining late-in-life savings patterns, while Lockwood suggests that highly valued bequests can explain much of late-in-life wealth holdings. This difference between the two studies in the strength of the bequest motive implied by the estimated parameters can be large at the lower wealth levels featured in their samples. At higher wealth levels these differences shrink, documenting the value of the higher-wealth VRI sample in identifying bequest motives. Indeed, DFJ are well aware of this, writing "Our sample of singles may not contain enough rich households to reveal strong bequest motives." We can extend DFJ's claims that "most people in [their] sample do not have

strong bequest motives" by documenting that a sample with higher wealth individuals yields a similar lack of individuals with strong bequest motives.

One reason we estimate weaker bequest motives than DFJ and Lockwood is that our modeling of non-homothetic utility in the ADL health state reduces the need for such a strong bequest motive to match wealth patterns. In Figure 19(b) we present the z_2 allocation for the problem presented in equation 12, with the bequest function replaced by the LTC state utility function. Lockwood does not have state dependent utility in the baseline model, and hence would suggest an equal allocation across states. DFJ do allow for a differential marginal utility multiplier when unhealthy, but no marginal utility additive shifter. (Additionally, DFJ allow for health-state dependent utility when an individual is sick, but do not model the LTC state separately.) The estimated state dependent utility parameter in DFJ is insignificant and assigns almost equal marginal utility to both states, resulting in an almost equal allocation. In this paper, we estimate LTC expenditure to be valued strongly as a necessity, resulting in a much higher allocation to the LTC state expenditure, not only at low levels of wealth, but also at higher wealth levels.

Analyzing the Behavioral Implications of Parameters from the Literature. We have documented substantial differences in the preferences estimated in the literature and that these differences induce very different savings patterns and imply different motives determining saving behavior. Using both new wealth data on high wealth individuals and new SSQ measurements, we have estimated preferences that imply that savings are driven to a significant degree by health related precautionary saving motives.
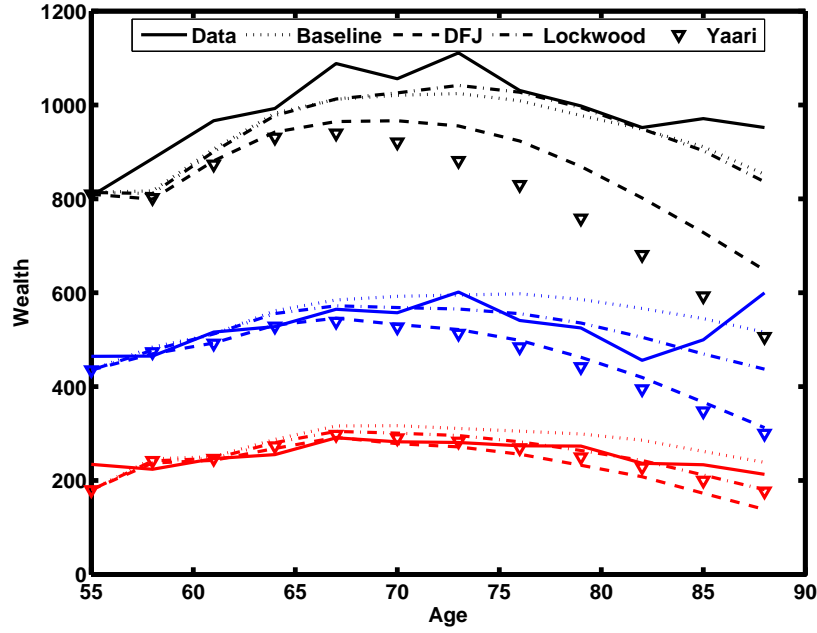
Figure 20: Model Fit to Wealth Moments for Alternative Parameters

Figure 20 compares moments of the population wealth profiles induced by these model parameters to those in the data. In addition to Lockwood and DFJ parameter values, we include a set of parameters designed to capture the classic setup similar to that in Yaari (1965) in which there is no health-state dependent utility function, zero bequest utility, no government provided care, and no indivisibility in the purchase of private LTC ($\gamma = 3$, $\theta_{ADL} = 1$, $\kappa_{ADL} = 0$, $\theta_{beq} = 0$, $\kappa_{beq} \approx \infty$, $\psi_G = 0$, $\omega_G = 0$, and $\chi = 0$).

Since our baseline parameters were in part chosen to match these wealth targets, it is no surprise that they do so well, as discussed previously. In terms of induced savings profiles, our estimated parameters look most similar to those of Lockwood, with DFJ and especially Yaari (1965) preferences leading to less saving over the life cycle than is observed in the data.[23] These differences show that the wealth data can be used to help distinguish across preference parameters, however, it is in the SSQ moments that the different implied motives really stand out.

(a) Lockwood

(b) DFJ

(c) Yaari
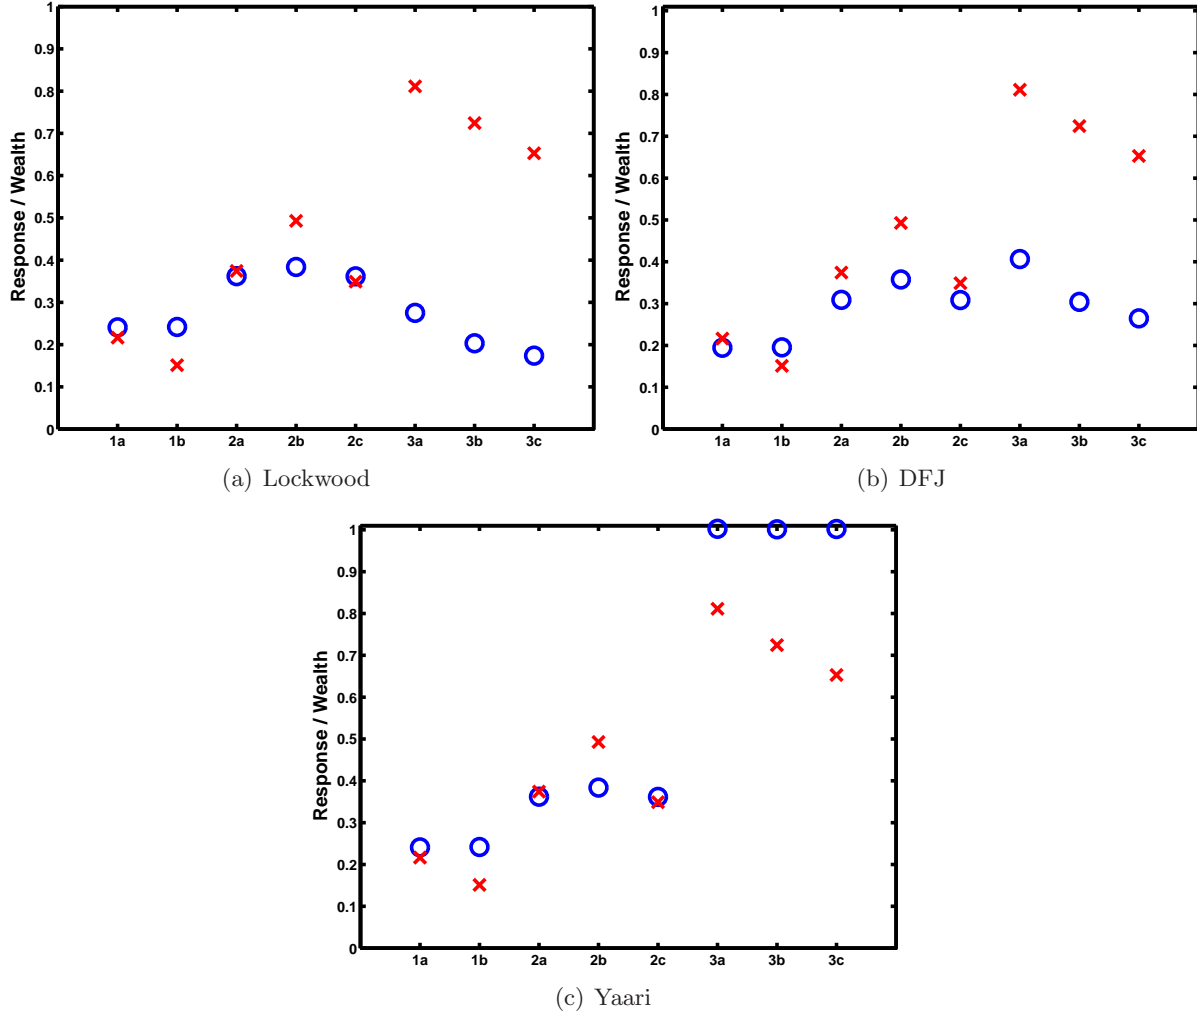
Figure 21: SSQ Means in Model (blue circle) and Data (red x) for Alternative Parameters

Figure 21 compares parameter-implied and empirical SSQ means across parameter sets. The results are striking, particularly for questions 3a, 3b, and 3c, which ask respondents to make a trade-off between spending on themselves when in need of LTC and leaving a bequest. The parameters estimated by Lockwood and DFJ generate data that match SSQs 1 and 2 well, but miss on SSQ 3, as they imply a much too high propensity to allocate towards bequests given their estimated θ_{beq} and κ_{beq}. In contrast, since the Yaari parameters place zero value on bequests, they overshoot and would predict much too small of an allocation towards bequests compared to the data. Together, Figures 20 and 21 show that models can be consistent

---

[23]Many papers in the literature, including DFJ, have estimates that are derived from less affluent samples. Thus, it is not surprising that the estimated parameters do not match moments for the high wealth individuals in the VRI.

with the preference information provided by SSQs without sacrificing fit with the empirical wealth-age distribution. Furthermore, SSQ methods can be used as a powerful parameter identification device.

Given that these parameters represent significantly different preferences and induce very different saving behavior in different contingencies, SSQs provide extra information that helps to identify these parameter values. Modeling an LTC-state dependent utility function allows us to match wealth moments across the wealth and age distribution, while also being consistent with survey evidence on stated preferences.

9 Additional Analyses

9.1 Sensitivity Analysis

In this section we provide additional analysis which demonstrates the robustness of the main conclusion that long-term-care related utility and risks are a significant driver of late-in-life saving behavior. To show this we perform sensitivity analysis by re-estimating the model under the following alternative assumptions: (i) we vary the CRRA parameter (γ); (ii) we vary the time discount factor (β); (iii) we allow for different curvature exponents in the utility functions for bequests and when alive; (iv) we allow for a more severe definition of needing help with ADLs; and (v) we add housing wealth to financial wealth and match the model to this measure of total wealth.

Table 6: Sensitivity Analysis

Alternatives	Parameter Estimates						Synthetic Expenditure Shares			
							W=100		W=200	
	γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G	ADL	Beq	ADL	Beq
Baseline	5.27	0.67	-37.44	1.09	7.83	77.43	0.68	0.11	0.56	0.19
Risk Aversion										
$\gamma = 2$	2.00	0.71	-34.49	0.45	22.13	58.48	0.61	0.20	0.46	0.34
$\gamma = 8$	8.00	0.86	-42.05	1.06	12.82	82.05	0.68	0.09	0.53	0.20
Discount Factor										
$\beta = .98$	4.91	0.69	-41.55	0.76	12.37	81.55	0.69	0.12	0.54	0.24
$\beta = .99$	4.65	0.88	-41.27	1.09	6.92	81.27	0.66	0.13	0.51	0.21
$\beta = .999$	4.90	0.85	-39.79	0.99	9.10	79.79	0.65	0.13	0.51	0.22
Beq. Risk Aversion										
$\gamma_{beq} = 2$	3.01	0.61	-36.98	1.23	18.00	69.73	0.68	0.14	0.48	0.34
$\gamma_{beq} = 8$	7.56	0.80	-40.59	1.05	12.48	96.74	0.69	0.08	0.55	0.17
ADL State Definition										
$\geq$ 2 ADLs	5.27	0.68	-37.38	1.09	7.79	82.39	0.68	0.11	0.56	0.19
$\geq$ 3 ADLs	5.48	0.69	-34.20	0.95	8.34	79.20	0.65	0.14	0.53	0.22
Housing Wealth										
15% Transaction Cost	5.47	0.66	-37.27	0.89	5.90	82.28	0.66	0.15	0.54	0.23
No Transaction Cost	6.02	0.68	-39.10	0.81	9.83	84.10	0.67	0.14	0.53	0.24

Table 6 reports the alternative parameter estimates associated with each sensitivity analysis exercise. In

addition, to provide a sense of the behavior that would be induced by the alternative parameters, the table reports expenditure shares associated with the synthetic static allocation problem presented in equation 11. The striking feature across all exercises is the large and negative κ_{ADL} coupled with $\theta_{ADL} < 1$, which drives a large allocation to ADL expenditure.

Different γ. Table 6 presents estimates of all other parameters when setting $\gamma = 2$ or $\gamma = 8$, in contrast to the baseline estimate of $\gamma = 5.27$. When risk aversion is higher, not much changes. With a smaller coefficient of relative risk aversion, there is slightly more need for a stronger bequest utility to match the same wealth data. This appears in the large decrease in θ_{beq}, since κ_{beq} actually increases. This change in preferences nearly doubles bequests in the synthetic allocation problem. Nonetheless, for this wide range of γ, LTC remains by far the dominant expenditure in the simple static allocation problem.

Different β. A higher β, reflecting more patient individuals, is associated with minor reductions in ADL expenditure and slight increases in bequests. The CRRA parameter estimate is lower than in the baseline, as increased patience helps match the wealth data with a somewhat weaker precautionary motive. Although there are some changes in bequest and ADL parameters, associated synthetic expenditure shares do not move much. In total, changes in the discount factor between the baseline $\beta = 0.97$ and $\beta = 0.999$ do not much affect the estimated parameters or the main conclusion that LTC related preferences and risks are a significant driver of savings over the life cycle.

Bequest-Utility-Specific Exponent. In this exercise, we allow for a bequest-specific exponent in the utility function. That is, different from γ, we allow for γ_{beq} such that: $\nu(b) = \frac{(\theta_{beq})^{-\gamma_{beq}}(b+\kappa_{beq})^{1-\gamma_{beq}}}{1-\gamma_{beq}}$.

Since the parameter is not precisely-estimated using wealth data alone and we do not have a strategic survey question to separately identify a bequest curvature parameter, we explore $\gamma_{beq} = 2$ and $\gamma_{beq} = 8$. We choose one value below and one above the ordinary CRRA parameter γ to explore whether the LTC-related saving motive remains strong with a more flexible parameterization. In particular, the lower γ_{beq} makes the bequest utility function less concave and provides more potential for a large bequest motive. *Ceteris paribus*, these changes represent significant differences in the marginal utility of a bequest relative to baseline. Nonetheless, when reestimating the model most parameters are little affected; the ordinary CRRA parameter γ changes the most from baseline. Aside from the estimate of γ, qualitatively and quantitatively the other parameter estimates remain similar to baseline. When $\gamma_{beq} = 2$, the bequest motive generates more saving and the model can match the same wealth moments with less precautionary savings, resulting in a lower estimate of $\gamma = 3.01$. At $\gamma_{beq} = 8$, the precautionary saving motive needs to be stronger to match the wealth moments, resulting in a larger than baseline $\gamma = 7.56$. Expenditure when in need of LTC is still viewed as a strong necessary good with substantial marginal valuation, indicated by a large and negative κ_{ADL} and $\theta_{ADL} < 1$.

At lower levels of expenditure, e.g., $W = \$100K$, the synthetic expenditure shares on ADL and bequests remain close to those resulting from the baseline parameters. At higher expenditure levels, e.g. $\$200K$, when $gamma_{beq} = 2$, the synthetic expenditure share on bequests is larger than in the baseline. Both when γ_{beq} equals 2 or 8, bequests are estimated as more of a luxury good (larger κ_{beq}). With a more linear bequest function, the higher $\gamma_{beq} = 8$ combines with the higher κ_{beq} and similar θ_{beq} to generate less bequests across

wealth levels. When the utility function is less concave, at higher wealth levels the lower $\gamma_{beq} = 2$ more than offsets the higher κ_{beq} to generate an increased desire to leave a bequest.

In sum, whether the bequest function is closer to linear or is more concave than in the baseline, whenever an individual needs help with LTC, long-term-care utility is always estimated to represent a strong necessary desire to spend with a high marginal valuation.

Different Health Risks. When mapping the model to the data, we defined $s = 2$ as meaning "needs and is receiving help with at least one activity of daily living." In this exercise, we change the definition of ADLs to only capture people who are in need of help with multiple ADLs. First we classify $s = 2$ if an individual needs and receives help with at least two ADLs, and then with at least three ADLs. This has three effects that partially offset each other. First, the ADL state is more severe, with higher persistence in the ADL state which increases the precautionary saving motive all other things equal. Second, the more severe ADL state is associated with lower survival rates, reducing expected longevity. Third, individuals are less likely to enter the ADL health state, reducing precautionary saving for any individual and reducing the fraction of people dissaving rapidly in the ADL state. Because of these offsetting effects the different estimated health state transition probabilities leave the estimated preference parameters largely unchanged.

Housing Wealth. In the baseline model, we map the wealth variable a to total financial wealth in the data. We do this because housing is a complicated asset that is in some ways like financial wealth and in some ways more like a durable consumption good. Housing is like financial wealth in that it can be sold and transformed into liquid wealth, albeit in a market subject to frictions. Housing is unlike financial wealth in that it is a physical good that delivers housing services. In this exercise, we map a to the sum of housing and financial wealth. In addition to using the full value of the house, to account for costs associated with the sale of a house, we also do this by subtracting 15 percent of the house price. Either way, this results in only modest changes in the estimated parameters and associated expenditure shares.

9.2 External Validity using HRS Data

Our final exercise is to use the estimated baseline parameters to compare expenditure patterns in the data and the model using the U.S. representative HRS population. In this section we provide a summary of this analysis, with details provided in Appendix E. First, we estimate health transitions, survival probabilities, health costs, and income profiles using the HRS sample, as opposed to the VRI-eligible HRS subsample. Then we sample single households from the HRS and record age, health status, and permanent income quintile in the year in which they enter the HRS. We generate a sequence of health transitions and cost shocks for each household and simulate their saving and expenditure patterns using the optimal policies associated with our estimated baseline model. Finally, we compare percentiles of the simulated wealth distribution with those measured by the HRS in the subsequent waves. When possible, we do this analysis separately for each cohort, given the different sampling procedures and aggregate environment associated with each cohort.

Overall, our model with our baseline parameter estimates performs well in the HRS sample and matches the 25th, 50th, and 75th percentiles of the wealth distribution for the cohorts that we consider with a few notable exceptions. For the War Babies and Early Baby Boomers cohorts, our model matches closely the

wealth distribution by age. The notable exception is that the model overpredicts wealth in the 2010 wave, possibly because we do not model the financial crisis. When using the Asset and Health Dynamics of the Oldest Old (AHEAD) sample, the model again matches the data well, although it slightly overpredicts the spend down of wealth in the median wealth percentile later in life.

10   Conclusion

Individuals face multiple risks and have multiple objectives late in life. The risks include longevity, health, and the need for care. The objectives include sustaining consumption in potentially different health states and providing a bequest. With the decline of defined benefit pensions, with increased longevity, with high and variable cost of care, and absent or highly-imperfect insurance markets for late-in-life risks, older households are increasingly responsible for managing financial assets late in life to address these multiple objectives across states of the world.

In this paper, we present novel preference elicitation and novel data to answer questions about late-in-life saving and spending. The modeling, measurement, and data were designed in tandem:

- We build an incomplete markets model of individuals, who save precautionarily when faced with health risks, the potential need for long-term care, and an uncertain life span and who value consuming, leaving a bequest, and receiving long-term care if they need it. Expenditures when in need of long-term care can be valued differently than ordinary consumption, depending on estimates of a health-state dependent utility function, and individuals choose on the margin the amount to spend when in need of LTC.

- We develop strategic survey questions (SSQs) that use a novel application of stated-preference methodology. These SSQs are designed to elicit responses that identify the preference parameters of the model.

- We develop a novel data infrastructure—the Vanguard Research Initiative—that combines information on income, assets, and demographics with the SSQs responses on a large sample of older Americans with sufficient financial assets to make operational the financial tradeoffs implied by the model.

We estimate the parameters of the model using moments from the wealth distribution alone, SSQs alone, and both wealth and SSQs. The point estimates based on the traditional approach of using the wealth distribution and the novel approach of using the SSQs are fairly similar. The estimates with the wealth data alone are, however, imprecise. Because multiple motives generate wealth accumulation and wealth is fungible across outcomes, wealth data alone do not strongly identify motives. In contrast, the SSQs provide relatively sharp identification of parameters. This precise identification of parameters allows us to make inferences that are difficult to support with behavioral data alone.

The findings in the paper point toward the risk of needing LTC when old as a substantial motive for accumulating assets. The marginal utility of expenditures when in need of LTC is larger than that from bequests. Due to the strength of the estimated health-state dependent utility function, the precautionary saving motives associated with LTC contribute significantly to late-in-life savings behavior, strongly affecting wealth accumulation—tending to increase savings both across the wealth distribution and over the life cycle— for a large fraction of the U.S. population. For a typical wealthholder in our sample, peak assets in retirement are about twice what they would be absent a difference between utility of spending in the healthy versus

45

need-for-care state. Health-state-dependent utility also helps explain the post-retirement pattern of assets. It is a motive for continued asset accumulation in retirement. LTC associated expenses also lead to a rapid drawdown of assets late in life and therefore generate lower bequests on average than the continued buildup of assets would otherwise suggest.

Finally, the model and SSQ-based estimates help distinguish the motives for leaving bequests—something difficult to do with the behavioral data alone because many bequests will be incidental, i.e., occur in the event where an earlier than expected death or less-costly than expected old age result in a larger than expected bequest. Bequests in the baseline model resulting from longevity risk, precautionary saving for LTC risk, and a warm-glow bequest motive are on average more than twice the incidental bequest owing to longevity uncertainty alone. Overall, we show that the desire to self-insure against long-term-care risk explains a substantial fraction of the wealthholding of older Americans.

References

AMERIKS, J., J. BRIGGS, A. CAPLIN, M. D. SHAPIRO, AND C. TONETTI (2016a): "Late-in-life Risks and the Under-Insurance Puzzle," *Vanguard Research Initiative Working Paper*.

———— (2016b): "Long-term Care Model," *Vanguard Research Initiative Technical Report*, http://ebp-projects.isr.umich.edu/VRI/papers/VRI-TechReport-LTC-Model.pdf.

———— (2016c): "Long-term Care Strategic Survey Questions," *Vanguard Research Initiative Technical Report*, http://ebp-projects.isr.umich.edu/VRI/papers/VRI-TechReport-LTC-SSQs.pdf.

AMERIKS, J., A. CAPLIN, S. LAUFER, AND S. VAN NIEUWERBURGH (2011): "The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Public Care Aversion from Bequest Motives," *Journal of Finance*, 66(2), 519–561.

AMERIKS, J., A. CAPLIN, M. LEE, M. D. SHAPIRO, AND C. TONETTI (2014): "The Wealth of Wealthholders," *Vanguard Research Initiative Working Paper*.

ARROW, K. (1974): "Optimal Insurance and Generalized Deductibles," *Scandanavian Actuarial Journal*, 3, 1–42.

BARSKY, R. B., F. T. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics*, 112, 537–579.

BLASS, A. A., S. LACH, AND C. F. MANSKI (2010): "Uing Elicited Choice Probabilities to Estimate Random Utility Models: Preferences for Electricity Reliability," *International Economic Review*, 51(2), 421–440.

BRAUN, R. A., K. A. KOPECKY, AND T. KORESHKOVA (2016): "Old, Sick, Alone and Poor: A Welfare Analysis of Old-Age Social Insurance Programs," *Review of Economic Studies*, forthcoming.

BROWN, J. R., AND A. FINKELSTEIN (2008): "The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market," *The American Economic Review*, 98(3), 1083–1102.

BROWN, J. R., G. S. GODA, AND K. MCGARRY (2016): "Heterogeneity in State-Dependent Utility: Evidence from Strategic Surveys," *Economic Inquiry*, 54(2), 847–861.

CHAMBERLAIN, G. (1992): "Comment: Sequential Moment Restrictions in Panel Data," *Journal of Business and Economic Statistics*, 10, 20–26.

DE NARDI, M. (2004): "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71(3), 743–768.

DE NARDI, M., E. FRENCH, AND J. JONES (2010): "Differential Mortality, Uncertain Medical Expenses, and the Saving of Elderly Singles," *Journal of Political Economy*, 118, 49–75.

———— (2016): "Medicaid Insurance in Old Age," *American Economic Review*, 106(11), 3480–3520.

DYNAN, K., J. SKINNER, AND S. ZELDES (2002): "The inportance of bequests and life-cycle savings in capital accumulation: A new answer," *American Economic Review*, 92, 274–278.

EVANS, W., AND W. K. VISCUSI (1991): "Estimation of State-Dependent Utility Functions Using Survey Data," *Review of Economics and Statistics*, 73(1), 94–104.

FELLA, G. (2014): "A Generalized Endogenous Grid Method for Non-Smooth and Non-Concave Problems," *The Review of Economic Dynamics*, 17(2), 329–344.

FINKELSTEIN, A., E. LUTTMER, AND M. NOTOWIDIGDO (2009): "Approaches to Estimating the Health State Dependence of the Utility Function," *American Economic Review: Papers and Proceedings*, 99(2), 116–121.

——— (2013): "What Good is Wealth without Health? The Effect of Health on the Marginal Utility of Consumption," *Journal of European Economic Association*, 11, 221–258.

FRENCH, E., AND J. JONES (2011): "The Effects of Health Insurance and Self Insurance on Retirement Behavior," *Econometrica*, 79, 693–732.

GENWORTH (2016): "Genworth 2016 Cost of Care Study," Discussion paper, Genworth.

GOURINCHAS, P.-O., AND J. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.

GUSTMAN, A., AND T. L. STEINMEIER (1986): "A Structural Retirement Model," *Econometrica*, 54(3), 555–584.

HONG, J. H., J. PIJOAN-MAS, AND J.-V. RIOS-RULL (2015): "Health Heterogeneity and Preferences," *working paper*.

HUBBARD, R. G., J. SKINNER, AND S. P. ZELDES (1994): "Expanding the Life-Cycle Model: Precautionary Saving and Public Policy," *The American Economic Review*, pp. 174–179.

HUBBARD, R. G., J. SKINNER, AND S. P. ZELDES (1995): "Precautionary Saving and Social Insurance," *The Journal of Political Economy*, 103(2), 360–3999.

HURD, M. D. (1989): "Mortality Risk and Bequests," *Econometrica*, 57(4), 779–813.

KIMBALL, M. S., C. R. SAHM, AND M. D. SHAPIRO (2008): "Imputing Risk Tolerance from Survey Responses," *Journal of the American Statistical Association*, 103(483), 1028–1038.

KOIJEN, R. S. J., S. VAN NIEUWERBURGH, AND M. YOGO (2016): "Health and Mortality Delta: Assessing the Welfare Cost of Household Insurance Choice," *Journal of Finance*, 71(2), 957–1010.

KOPECKY, W., AND T. KORESHKOVA (2014): "The Impact of Medical and Nursing Home Expenses and Social Insurance Policies on Savings and Inequality," *American Economic Journal: Macroeconomics*, 6(3), 29–72.

KOTLIKOFF, L. J., AND L. H. SUMMERS (1981): "The Role of Intergeneration Transfers in Aggregate Capital Accumulation," *Journal of Political Economy*, 89(4), 706–732.

LAIBSON, D., A. REPETTO, AND J. TOBACMAN (2007): "Estimating Discount Functions with Consumption Choices over the Lifecycle," *NBER Working Paper 13314*.

LILLARD, L., AND Y. WEISS (1997): "Uncertain Health and Survival: Effects on End of Life Consumption.," *Journal of Business and Economic Statistics*, 15(2), 254–68.

LOCKWOOD, L. (2016): "Incidental Bequests: Bequest Motives and the Choice to Self-Insure Late-Life Risks," *NBER Working Paper 20745*.

LOVE, D. A., M. G. PALUMBO, AND P. A. SMITH (2009): "The trajectory of wealth in retirement," *Journal of Public Economics*, 93(1), 191–208.

LUPTON, J., AND W. KOPCZUK (2007): "To Leave or Not To Leave: The Distribution of Bequest Motive," *Review of Economic Studies*, 74(1), 207–235.

NEWEY, W. (1985): "Generalized Method of Moments Specification Testing," *Journal of Econometrics*, 29, 229–256.

NEWEY, W. K., AND D. L. MCFADDEN (1994): "Large Sample Estimation and Hypothesis Testing," *Handbook of Econometrics*, 4, 1027–1057.

PAKES, A., AND D. POLLARD (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57, 1027–1057.

PALUMBO, M. (1999): "Uncertain Medical Expenses and Precautionary Saving near the End of the Life Cycle," *Review of Economic Studies*, 66, 395–421.

POTERBA, J. M., S. F. VENTI, AND D. A. WISE (2013): "Health, education, and the post-retirement evolution of household assets," Discussion paper, National Bureau of Economic Research.

POWELL, J. (1994): "Estimation of Semi-Parametric Models," *Handbook of Econometrics*.

SCHOLZ, J. K., A. SESHADRI, AND S. KHITATRAKUN (2006): "Are Americans Saving "Optimally" for Retirement?," *The Journal of Political Economy*, 114(4), 771–807.

SOTO, M., R. PENNER, AND K. SMITH (2009): "How Seniors Change Their Asset Holdings During Retirement," *Available at SSRN 1553268*.

VAN DER KLAAUW, W. (2012): "On the Use of Expectations Data in Estimating Structural Dynamic Choice Models," *Journal of Labor Economics*, 30(3), 521–554.

VAN DER KLAAUW, W., AND K. I. WOLPIN (2008): "Social Security and the Retirement and Savings Behavior of Low-income Households," *Journal of Econometrics*, 145(1-2), 21–42.

VISCUSI, W. K., AND W. EVANS (1990): "Utility Functions that Depend on Health Status: Estimates and Economic Implications," *American Economic Review*, 80(3), 353–374.

YAARI, M. (1965): "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer," *The Review of Economic Studies*, 32(2), 137–150.

Appendix

A Estimated Inputs for Model

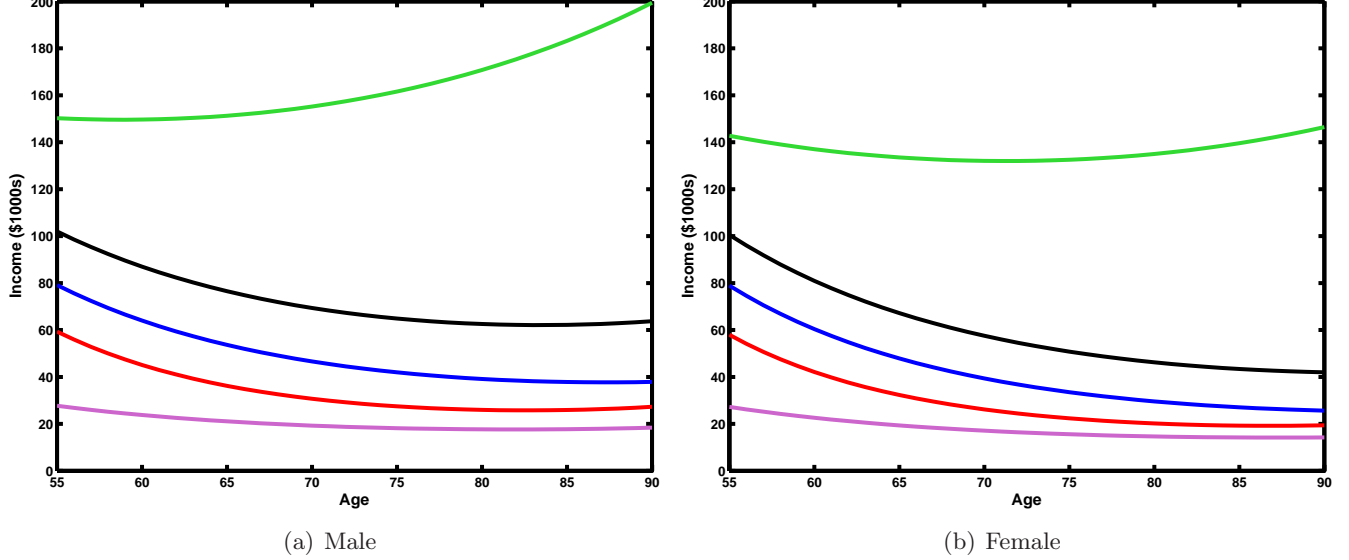### A.1 Income

(a) Male

(b) Female

Figure A.1: Income Profile Quintiles

We estimate a deterministic income process from the cross-sectional income distribution. Income is defined as the sum of labor income, publicly and privately provided pensions, and disability income, as measured in VRI Survey 1. The income processes are estimated to be a function of a constant, age, age squared, gender, and the interaction of gender and age. To ensure that income is positive in all periods, we estimate a quantile regression of log income on these variables. Because we allow for 5 income profiles, the quantile regression is estimated for the 10^{th}, 30^{th}, 50^{th}, 70^{th}, and 90^{th} percentiles of the income distribution. We calibrate our income processes to the resulting estimates and group individuals into income profile quintiles.

A.2 Health

Health-State Transition Matrix. Using appropriate health state definitions we estimate a sequence of health transition matrices conditional on a vector $x_{i,t}$ which includes individual i's age, t, and gender, g. The HRS only records 2 year health state transitions which we use to identify the one-year transition probabilities in a manner similar to De Nardi, French, and Jones (2010). To do this, we write the two year transition probabilities as:

$$Pr(s_{t+2} = j|s_t = i) = \sum_{k=0}^{3} Pr(s_{t+2} = j|s_{t+1} = k)Pr(s_{t+1} = k|s_t = i) = \sum_{k=0}^{3} \pi_{kj,t+1}\pi_{ik,t}$$

where,

$$\pi_{ik,t} = \frac{\phi_{ik,t}}{\sum_{m=0}^{3} \phi_{im,t}} \text{ and } \phi_{ik,t} = exp(x_{i,t}\beta_k).$$

We then estimate β_k using a maximum likelihood estimator, and use these estimates to construct the corresponding cells in the health transition matrices. Health transitions are estimated using HRS waves 2 through 10. The questions necessary to make the health state assignment are not available in the 1992 survey, so we exclude this wave from the health transition estimates.

Figures A.2 and A.3 display the estimated health state transition probabilities ($\pi_g(s'|t,s)$). Section 5.1.2 describes the estimation methodology. An additional consideration is how to define the "needs long term care" health status. There are 3 measures in the HRS that could potentially be used. The first is nursing home stay, the second is needs help with the activities of daily living, and the third is receives help with the activities of daily living.

Nursing home stay (more than 120 nights in a nursing home before the current interview or currently in a nursing home at time of interview) is what De Nardi, French, and Jones (2010) used. Given that we allow people in the model to choose their type of care, we want a less restrictive definition for $s = 2$. The ADL questions in the RAND version of the HRS list many activities of daily living and asks if the respondent has difficulty completing those tasks without help. In some sense, these questions provide the broadest possible definition of the ADL state, since many people could report having difficulty, but would still be able to live without receiving help. We choose to implement the intermediate measure: we categorize and individual as needing help with LTC if they have difficulty with at least one ADL and they also receive outside help completing the ADL task.

Health Cost. To estimate the mean of the health cost distribution, $\mu(t, g, s)$, we regress annualized log out-of-pocket medical expenditures (variable r10oopmd in the RAND HRS) on age, gender, health state, and interaction terms. Log medical expenses are modeled as a linear function of a quartic in age, prior health status, gender, and health interacted with age. This relationship is estimated using OLS. Using the residuals from this first regression, we regress the squared residuals on the same set of state variables as in the first regression to find the conditional variance of medical expenses, $\sigma^2(t, g, s)$. The resulting mean and variance estimates parameterize the lognormal distribution for medical expenses conditional on age, gender and health status. We then discretize the error term using the Tauchen method to generate the medical expense process.

Figure A.4 plots the median mandatory health costs spent over the life cycle by men and women of different health status. Men in poor health spend around $100 more per year out of pocket for health costs than healthy men. Later in life, men in need of LTC spend about $600 more than healthy men for non-LTC health costs. Since saving behavior is also affected by the risk of high costs, and not just the typical expenditures, in Table A.1 we present examples of the health cost distribution across gender, age, and health states. Overall, these out of pocket health costs are much smaller than LTC expenditures and thus contribute little to the overall precautionary savings motive.
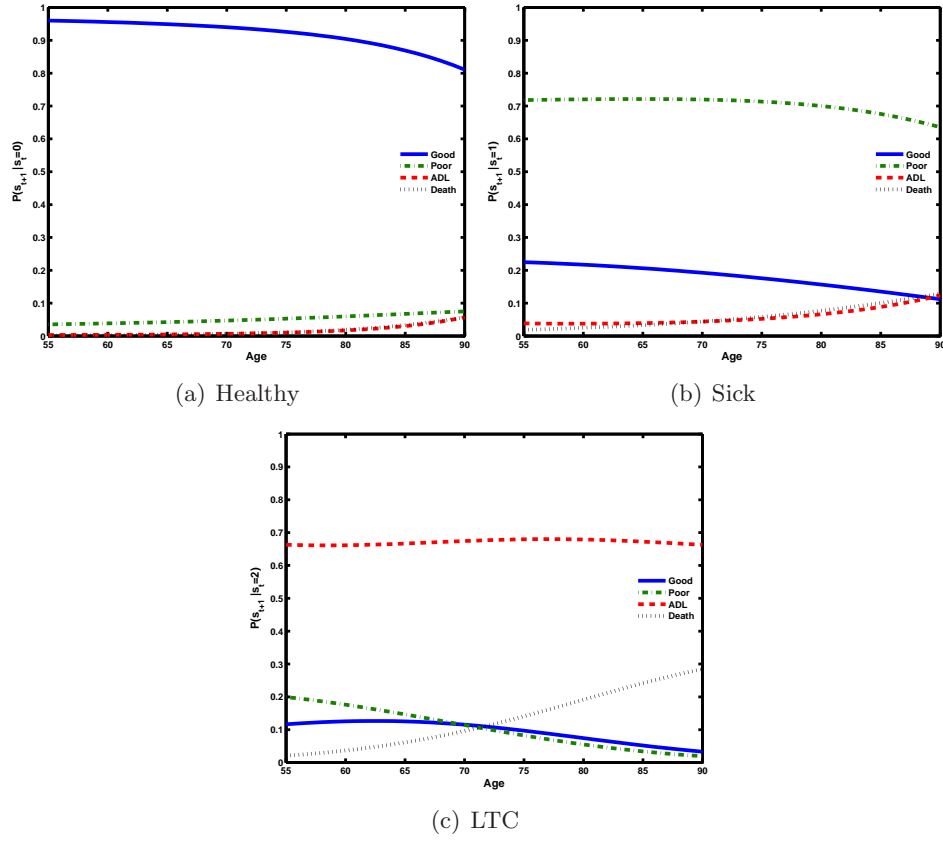
(a) Healthy

(b) Sick

(c) LTC

Figure A.2: Male Health State Transition Profile

Table A.1: Mean and Standard Deviation of the Health Expense Lognormal Distributions ($ thousands)

Female Health Expenses

Health State		55	60	65	70	75	80	85	90	95	100
							Age				
1	mean	2.4	2.0	1.8	1.7	1.9	2.4	3.1	3.9	4.2	3.1
	sd	6.5	4.7	3.7	3.6	4.2	6.1	9.6	14.2	16.2	10.9
2	mean	3.4	2.9	2.6	2.4	2.5	3.1	4.0	4.9	5.1	3.7
	sd	12.4	8.7	6.6	6.1	6.8	9.3	13.9	19.7	21.4	13.8
3	mean	3.5	3.1	2.8	2.9	3.3	4.2	5.7	7.5	8.2	6.3
	sd	16.6	12.1	9.5	8.9	10.4	14.4	22.1	32.1	36.0	23.8

Male Health Expenses

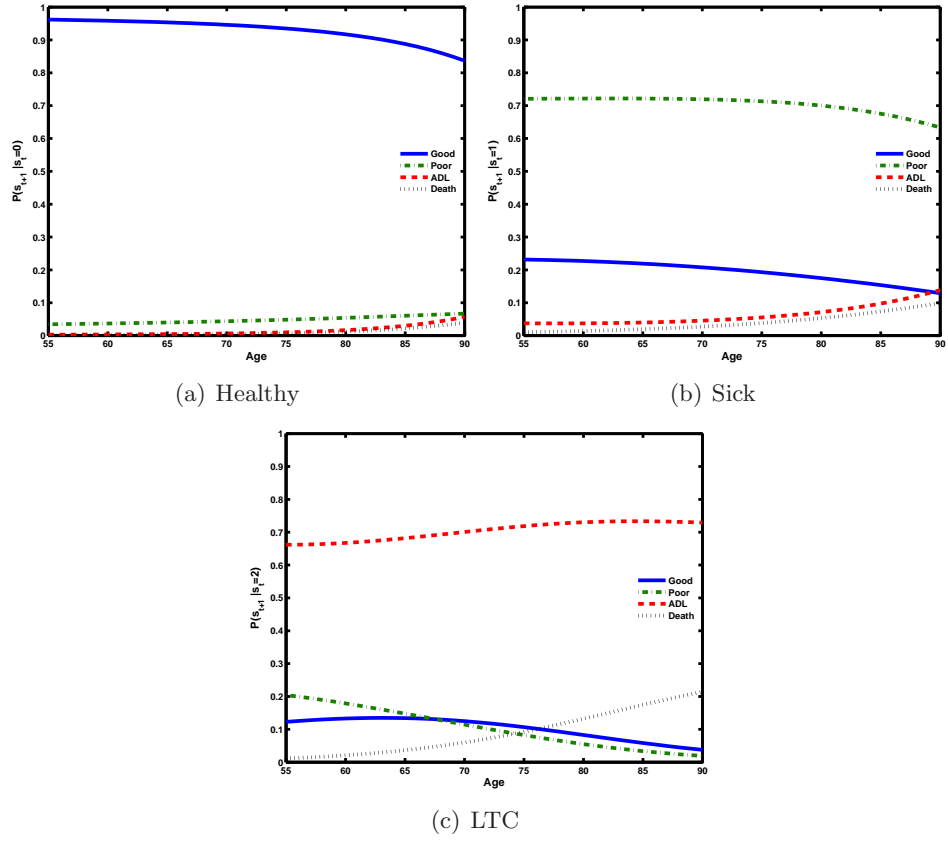| Health State		55	60	65	70	75	80	85	90	95	100
							Age				
1	mean	2.8	2.3	1.9	1.7	1.8	2.1	2.6	3.1	3.1	2.2
	sd	10.1	6.5	4.6	3.8	4.0	5.0	6.9	9.0	9.1	5.3
2	mean	4.1	3.2	2.6	2.4	2.4	2.7	3.3	3.9	3.8	2.6
	sd	19.1	11.9	8.0	6.5	6.4	7.7	10.1	12.6	12.1	6.8
3	mean	4.2	3.5	3.0	2.9	3.1	3.7	4.8	5.9	6.2	4.5
	sd	25.5	16.4	11.5	9.5	9.7	11.9	16.1	20.6	20.3	11.7

(a) Healthy

(b) Sick

(c) LTC

Figure A.3: Female Health State Transition Profile

(a) Male

(b) Female

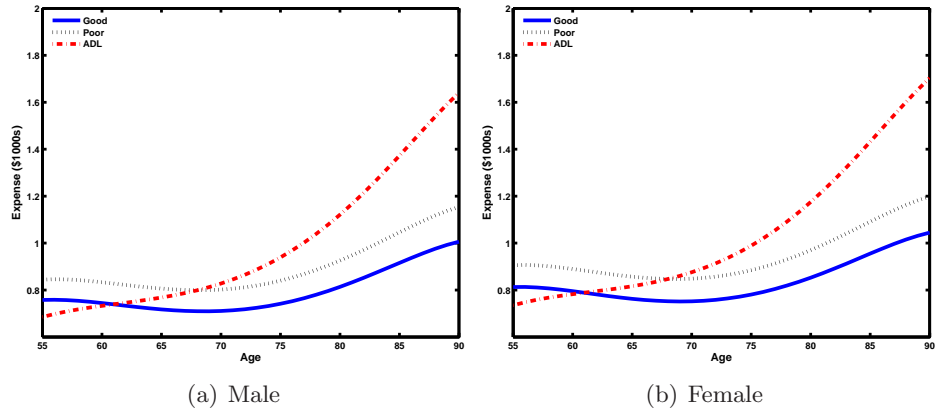Figure A.4: Median Health Cost Profile

B   SSQ Response Histograms

In Figure B.5 we present the responses to all of the SSQ variants. In SSQs 1a and 1b, a response indicates how much income a respondent would be willing to risk. In SSQs 2a, 2b, 2c, 3a, 3b, a response indicates the amount of wealth allocated to the ADL state. In SSQ 4a a response indicates the wealth level at which a respondent is indifferent between a world with and without public LTC funding.
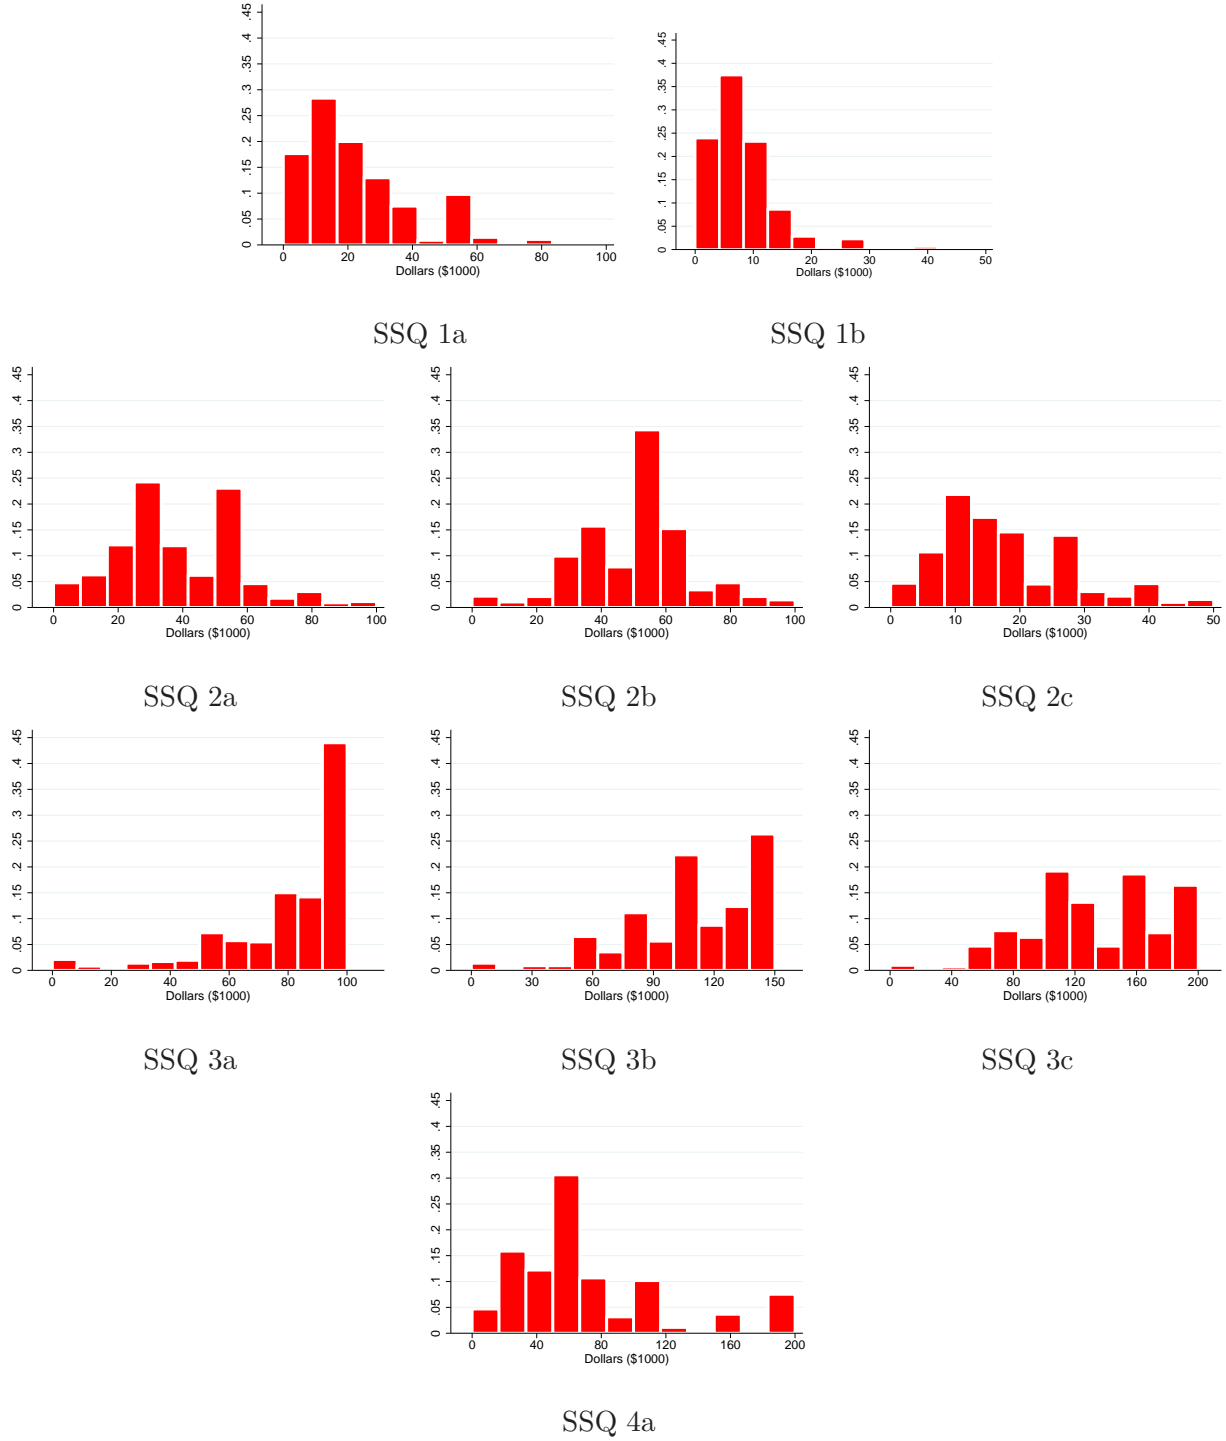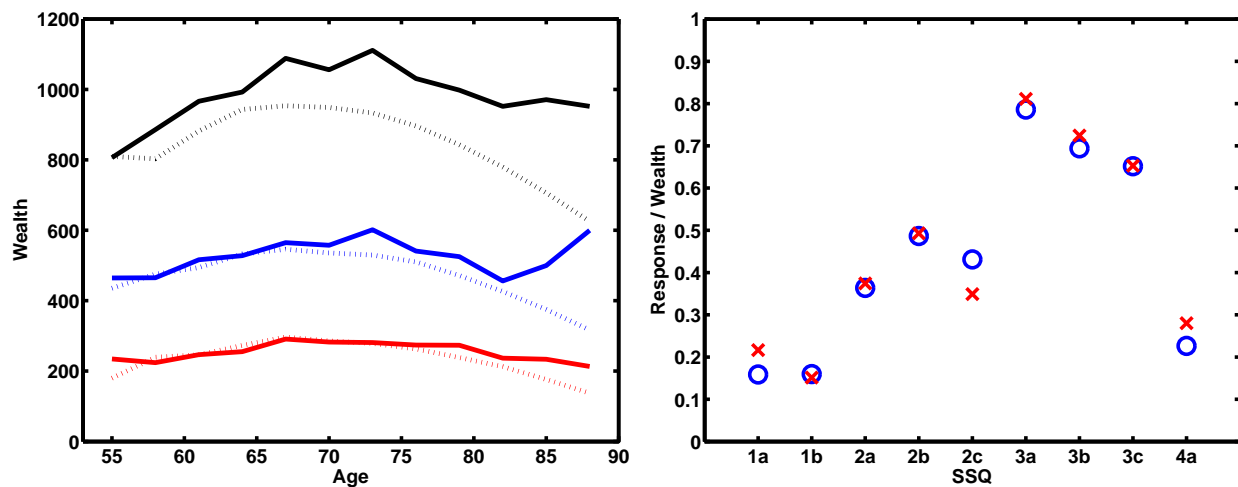

Figure B.5: Distribution of Responses to SSQs

C Estimation Results Using the Optimal Weighting Matrix

(a) 75p, 50p, and 25p Wealth Moments in Model (dashed) (b) SSQ Means in Model (blue circle) and Data (red x) and Data (solid)

Figure C.1: Model Fit When Jointly Targeting Wealth and SSQ Moments
Using the Optimal Weighting Matrix

Table C.1: Estimated Parameters: Alternative Weighting Matrices

Joint Estimation: Baseline Model					
γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
5.27	0.67	-37.44	1.09	7.83	77.43
(0.05)	(0.37)	(0.22)	(0.25)	(0.48)	(9.49)

Joint Estimation: Optimal Weighting Matrix					
γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
4.69	1.00	-38.71	1.08	15.32	78.54
(0.05)	(0.08)	(0.98)	(0.02)	(0.24)	(2.21)

This table presents parameter estimates for the estimation targeting jointly both sets of moments for the cases in which we use the baseline weighting matrix and the optimal weighting matrix. Standard errors are reported in parentheses.

D   Estimation Technicalities

The technicalities presented here, especially those related to the wealth moments, are adapted from the detailed appendix provided in De Nardi, French, and Jones (2010). The conditional moment conditions are defined as

$$\mathbb{E}\left[\mathbb{I}_{\{a_i < a_x^p(\Xi,\Theta,X)\}} - p | x_i \in x\right] = 0 \tag{D.1}$$

and

$$\mathbb{E}\left[s_m(\Theta) - z_{i,m} | x_i = x\right] = 0. \tag{D.2}$$

Converting the above conditional moment into an unconditional moment for estimation requires an application of the law of iterated expectations. Following Chamberlain (1992) the unconditional moment conditions are written using indicator functions. These moment conditions are presented here as the expectations

$$\mathbb{E}\left[\left(\mathbb{I}_{\{a_i < a_x^p(\Xi,\Theta,X)\}} - p\right) \times \mathbb{I}(x_i \in x)\right] = 0 \tag{D.3}$$

and

$$\mathbb{E}\left[(s_m(\Theta) - z_{i,m}) \times \mathbb{I}(x_i \in x)\right] = 0, \tag{D.4}$$

which can be feasibly be implemented as

$$\frac{1}{N}\sum_{i=1}^{N}\left[\left(\mathbb{I}_{\{a_i < a_x^p(\Xi,\Theta,X)\}} - p\right) \times \mathbb{I}(x_i \in x)\right] = 0 \tag{D.5}$$

and

$$\frac{1}{N}\sum_{i=1}^{N}\left[(s_m(\Theta) - z_{i,m}) \times \mathbb{I}(x_i \in x)\right] = 0. \tag{D.6}$$

This set of the sample analogue moment conditions is denoted by $g(\hat{\Xi}, \Theta, X)$. There are I individuals in the survey and we draw N individuals (with replacement) to compute the simulated moments. Define $\tau = \frac{I}{N}$. Let Ω denote the covariance matrix of these moment conditions, with $\hat{\Omega}$ denoting the empirical covariance matrix. Define the optimal weighting matrix $\hat{W} = \hat{\Omega}^{-1}$, and define

$$\hat{\Theta} = \arg\min_{\Theta} \frac{I}{1+\tau} g(\hat{\Xi}, \Theta, X) \hat{W} g(\hat{\Xi}, \Theta, X).$$

For matrix $W = plim_{N\to\infty}\hat{W}$, we know that

$$\sqrt{N}(\hat{\Theta} - \Theta_0) \to \mathbb{N}(0, \Psi),$$

where
$$\Psi = (1 + \tau)(D'WD)^{-1}(D'W\Omega WD)(D'WD)^{-1}.$$

D is the gradient:
$$D = \left.\frac{\partial g(\hat{\Xi}, \Theta, X)}{\partial \Theta'}\right|_{\Theta=\Theta_0}.$$

To obtain an expression for D, let f_a denote the density function of the empirical asset distribution and let f_z denote the density function of the empirical survey response distributions. f_a and f_z are both estimated using kernel density estimation. Then, following Pakes and Pollard (1989), Newey and McFadden (1994), and Powell (1994) the unconditional wealth moments can be represented as

$$\mathbb{P}(x_i \in x) \times \left[\int_{-\infty}^{a_x^p(\hat{\Xi},\Theta,X)} f_a(a_i|x)da_i - p\right] \tag{D.7}$$

and the unconditional strategic survey question moments can be represented as

$$\mathbb{P}(x_i \in x) \times \left[\int (s(\Theta) - z_{i,m}) \times f_z(z_{i,m}|x)dz_{i,m}\right]. \tag{D.8}$$

Then, the rows of the derivative matrix can be expressed as

$$D_V = \mathbb{P}(x_i \in x) \times f(a_x^p(\hat{\Xi}, \Theta, X))|x) \times \left.\frac{\partial a_x^p(\hat{\Xi}, \Theta, X)}{\partial \Theta'}\right|_{\Theta=\Theta_0} \tag{D.9}$$

and

$$D_S = \mathbb{P}(x_i \in x) \times \int \left.\frac{\partial s(\Theta)}{\partial \Theta'}\right|_{\Theta=\Theta_0} \times f_z(z_m^i|x)dz_m^i. \tag{D.10}$$

The above expressions are used to calculate the respective D matrices, with numerical derivatives used to calculate $\frac{\partial s(\Theta)}{\partial \Theta'}$ and $\frac{\partial a_x^p(\hat{\Xi},\Theta,X)}{\partial \Theta'}$.

Newey (1985) proves, with the following definitions,

$$Q = I - D(D'WD)^{-1}D'W \tag{D.11}$$
$$R = Q\hat{\Omega}Q, \tag{D.12}$$

that

$$\frac{I}{1+\tau}g(\hat{\Xi}, \hat{\Theta}, X)R^{-1}g(\hat{\Xi}, \hat{\Theta}, X) \rightsquigarrow \chi_{J-M}^2. \tag{D.13}$$

Finally, noting that asymptotically $\hat{W} \to \Omega^{-1}$ then $W = \hat{\Omega}^{-1}$, $V = (D'\hat{\Omega}^{-1}D)^{-1}$ and $R = \hat{\Omega}$.

In practice, repeating the calculation of the estimated covariance matrix with the parameter set $\hat{\Theta}$ allows us to calculate standard errors using the above asymptotic distribution. In this expression, we ignore the error in the first stage estimates by treating those as fixed numbers.

D.1   Details of Simulation

This section details step-by-step the simulation procedure that is used to implement the MSM procedure.

Before beginning the process, aggregate the survey data to create the initial empirical distribution of state variables. These consist of age, health, income, wealth, and SSQ responses. In addition, estimate the first stage parameters, $\hat{\Xi}$, and treat them as constants equal to the point estimate.

1. Set the second stage preference parameters that will be used in this simulation: $\hat{\Theta} = \Theta$

2. Sample a large number of individuals (N=10,000) from the initial distribution

3. Compute optimal policies using structural model for the specified parameter set Θ

    (a) Solve for the optimal policy functions for the life cycle savings model (as detailed in VRI Technical Report: Long-term Care Model)

    (b) Solve for the optimal strategic survey question responses (as detailed in VRI Technical Report: Long-term Care SSQs)

4. For each simulated individual $i \in \{1, 2, \ldots, N\}$, using $\hat{\Xi}$, simulate a series of health, health cost, and mortality shocks that dictate their life histories

5. Conditional on each individual's life history of exogenous idiosyncratic shocks, simulate each individual's choices using the optimal policy functions, yielding a life-cycle wealth profile for each individual

6. Conditional on $\hat{\Theta}$, simulate the individual strategic survey responses according to the optimal response policies, yielding a set of strategic survey responses for each individual

7. Construct moments:

    (a) Aggregate the individual life-cycle wealth profiles to construct the wealth moments described in Section 5.2.1

    (b) Aggregate the individual strategic survey question answers to construct the SSQ moments described in Section 5.2.2

    (c) Concatenate these two sets of moments for the baseline case. Otherwise, just use the relevant moment set

8. Compare these simulated moments to their empirical counterparts and update $\hat{\Theta} = \Theta'$

9. Repeat steps 1 to 8 with the updated parameters until the minimum of the MSM objective function, equation 7, is located

Steps 8 and 9 require the specification of an optimization algorithm. We optimize in two parts. First, we conduct a global search for our initial point Θ by repeating the algorithm minus steps 8 and 9 for 5,000 points defined by a Sobol sequence over our feasible parameter space. We then implement the algorithm above starting from the point associated with the minimum of the objective function, using a parallelized pattern search algorithm. At each iteration, we evaluate the function at M new points drawn from a scaled basis set of the parameter space. We then define the new parameter values, Θ' to be the $\arg\min$ of the

objective function evaluated at these M points and the previous $\arg\min \Theta$. If Θ remains the $\arg\min$, then we scale the basis set downwards to shrink the search region. If a new $\arg\min$ is found, the basis set is scaled upwards to expand the search region. This process repeats iteratively until convergence.

E   HRS Analysis

To implement comparison to a representative population and explore how general are findings from the Vanguard Research Initiative (VRI), in this appendix we explore the implications of our findings using the Health and Retirement Study (HRS). The HRS is a longitudinal survey of a representative sample of older (age 50+) Americans. Beginning in 1992, participants in this study provide information regarding their health, labor market activity, and finances biennially. Importantly, the HRS collects health, income, and wealth measures that are needed to construct the state variables, transition probabilities, and non-SSQ moments that are used in our study of the VRI. This allows us to explore the behavioral implications of LTC and bequest motives using the life-cycle saving model presented in this paper in a sample that is very different from the VRI.

Since the HRS does not have the SSQs necessary for sharp identification of the model's parameters, we use parameters estimated from the VRI to simulate the model and compare the predicted saving profiles with those actually observed in the HRS. Our main conclusion from this analysis is that our model and estimated saving motives perform extremely well at predicting non-targeted HRS wealth moments in nearly all analyses we conduct, and the only case where our baseline model does not perform well can easily be explained by weak identification of the normal health consumption floor. This provides strong indication that our main finding that strong LTC motives are significant drivers of late in life saving is a general finding that holds outside of the VRI sample we focus on in this study.

E.1   Data and Estimated Inputs to Model

E.1.1   HRS Data

The Health and Retirement Study (HRS) is a longitudinal study of Americans aged 50+ administered and maintained by the Survey Research Center (SCR) at the Institute for Social Research (ISR) at the University of Michigan since 1992. The study interviews approximately 20,000 respondents biennially and records information on health, income, labor market activity, and finances. The HRS is composed of 7 distinct cohorts, four of which are used in this appendix's analysis. We provide an overview of each here, including date of birth, and survey years in Table E.1.

Differences in income profiles, inability to identify time/age/cohort effects, and notable differences in wealth profiles lead us to conclude (as do other studies that use the HRS as a primary data source) that it is preferable to analyze cohorts separately. We thus focus on the original HRS, AHEAD, and only pool the most recent cohorts (WB/EBB) when analyzing saving motives.

To maintain comparability to the results in the paper, we restrict our analysis to individuals who are over age 50, are not married or partnered in any observed waves, and for which we observe both wealth and income at the time of entry into the sample. Table E.2 presents summary statistics by cohorts after imposing these sample restrictions.

Table E.1: Summary of Cohort Structure

	Cohort			
	Original HRS (HRS)	**Asset and Health Dynamics Among the Oldest Old (AHEAD)**	**War Babies (WB)**	**Early Baby Boomers (EBB)**
Birth year:	1931-1941	pre-1924	1942-1947	1948-1953
First Surveyed:	1992	1993	1998	2004

E.1.2 Estimated Inputs to Model

Health and Survival Transition Probabilities & Health Costs. Health transitions, survival probabilities, and health costs are estimated using the procedure described in Section 5.1 using the HRS data without imposing VRI-eligible screens.

Income. To esitmate HRS income profiles, we follow De Nardi, French, and Jones (2010). Specifically, we first define permanent income as the cohort-specific percentile of average income observed across survey waves. We then sort individuals into income quintiles and model income (y_t) as a function of age, gender, permanent income

$$\ln y_{i,t} = \beta X_{i,t} + \Gamma_i + \eta_{i,t}$$

where $X_{i,t}$ includes age, gender, permanent income quintile, and polynomial and interaction terms, Γ_i is an individual fixed effect, and $\eta_{i,t}$ is an error term. For the AHEAD sample respondents are observed at very old ages, and income can thus be estimated at all ages using the AHEAD sample alone. For the HRS and WB/EBB samples, income is not observed at older ages. Thus, when studying these cohorts, income profiles are estimated by combining these samples with the AHEAD sample and including a cohort indicator variable in $X_{i,t}$. The resulting income profiles that are used in our model are presented in Table E.3.

Table E.2: Summary of State Variables

HRS *(N = 2258)*

	Mean	**10p**	**25p**	**50p**	**75p**	**90p**
Wealth:	128,105	1,000	1,000	35,737	127,729	323,246
Income:	29,633	2,511	9,437	22,439	40,798	64,335

				Health		
Avg. Age		**Male**		**Good**	**Poor**	**LTC**
55.3		32%		67%	33%	0%

AHEAD *(N = 3765)*

	Mean	**10p**	**25p**	**50p**	**75p**	**90p**
Wealth:	143,526	1,000	4,743	61,957	164,526	340,382
Income:	16,073	7,364	9,302	12,806	18,623	27,445

				Health		
Avg. Age		**Male**		**Good**	**Poor**	**LTC**
79.		21%		56%	28%	16%

WB/EBB *(N = 1373)*

	Mean	**10p**	**25p**	**50p**	**75p**	**90p**
Wealth:	143,871	1,000	1,000	32,415	149,186	395,585
Income:	34,044	234	9,244	25,478	49,144	76,056

				Health		
Avg. Age		**Male**		**Good**	**Poor**	**LTC**
53.0		36%		65%	29%	5%

Note: Values are measured in 2010 dollars during the year the individual entered the HRS.

Table E.3: Income Profiles by Cohort and Quintile

	WB/EBB					HRS				
Quintile	_const._	_Age_	Age^2	_Male_	$Male \times Age$	_const._	_Age_	Age^2	_Male_	$Male \times Age$
1	-12.125	.518	-.003	1.570	-.017	-28.599	.972	-.006	1.620	-.019
2	2.300	.164	-9.6e-4	1.126	-.013	-7.161	.419	-.003	1.883	-.024
3	8.404	.026	-1.5e-4	.955	-.011	1.186	.213	-.001	1.981	-.026
4	11.9713	-.042	2.0e-4	.314	-.004	8.689	.033	-2.03e-4	.906	-.012
5	11.890	-.022	3.9e-5	.152	-.003	8.578	.066	-5.3e-3	.427	-.006

	WB/EBB				
Quintile	_const._	_Age_	Age^2	_Male_	$Male \times Age$
1	11.030	-.062	4.0e-4	.937	-.010
2	7.500	.046	-2.9e-4	.038	8.5e-5
3	8.197	.036	-2.4e-4	.051	-2.0e-4
4	7.630	.062	-4.1e-4	-.189	.002
5	8.578	.066	-5.3e-4	.427	-.006

E.2 Model-implied Wealth Profiles for the HRS Population

Our simulation procedure is as follows: We initialize each individual in our sample at the same state variables presented in Table E.2. We then calculate the saving policy function implied by our model, simulate a set of shocks for each individual, and simulate their asset holdings forward for as long as they remain alive. We then calculate the $25/50/75^{th}$ percentile of the aggregate asset holdings in each wave that the cohort is observed. We plot these simulated asset percentiles against the actual wealth percentile observed in the wave indicated. We thus do not seek to match the plotted empirical wealth profiles after the initial wave in which we start our simulation.

E.2.1 AHEAD

In landmark paper, De Nardi, French, and Jones (2010) estimate the saving dynamics of the AHEAD sample. Thus, to evaluate our model's performance in a familiar cohort and benchmark our results to a literature standard, we conduct the analysis presented above in the AHEAD sample. Following De Nardi, French, and Jones (2010) we split the sample into five birth-year cohorts to ensure we are comparing similar individuals. The resulting fit for the first two cohorts is presented in Figure E.1.
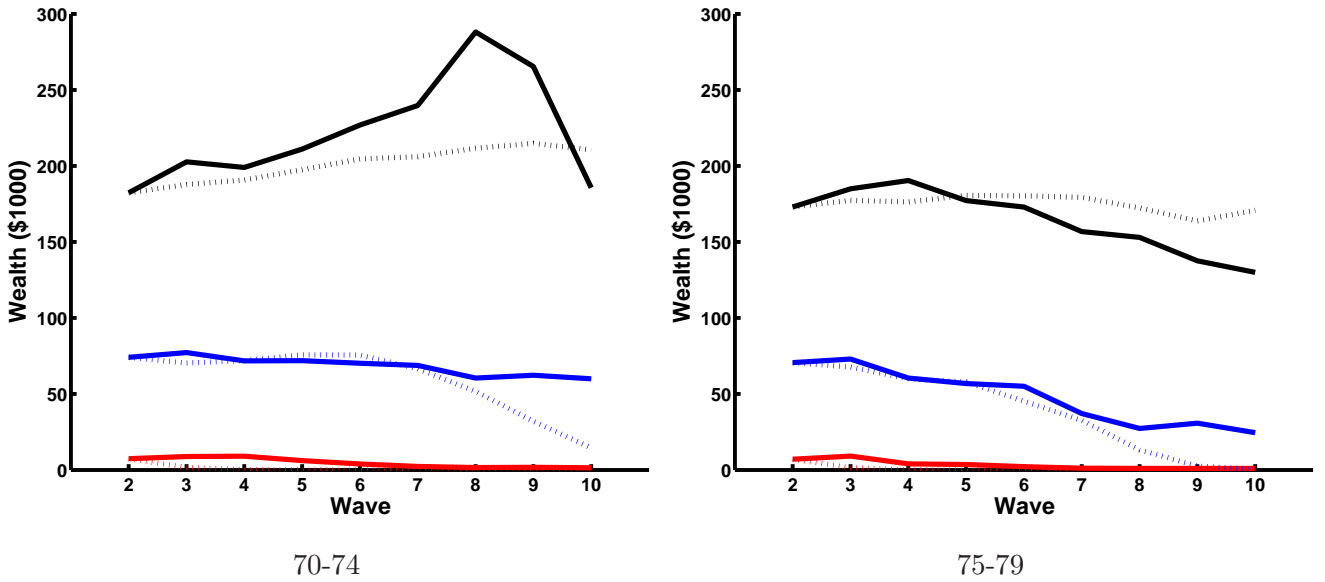

70-74   75-79

Figure E.1: Model results for AHEAD cohort

Among poorer households (25^{th} percentile) we successfully predict that households will not accumulate any wealth. Our model also performs reasonably well at matching the saving decisions of wealthier households, although our model predicted 75^{th} percentile is slightly more flat than that which we observe in the data. Our model does slightly slightly over-predict the spend down of the median wealth level in both figures, but the difference between model and data is only pronounced during the last 3 waves.

E.2.2 WB/EBB

We next examine the model's prediction for the War Babies and Early Baby Boomers. We isolate these cohorts because they are the most recent and thus make saving decisions at a similar time and environment

to the VRI. The resulting fit is presented in Figure E.2.

Again, the model does a reasonably good job matching the data. It does underpredict wealth holdings in Wave 10 (2010), which could potentially be explained by our not modeling equities and therefore missing capital gain losses experienced during the experience the financial crisis.
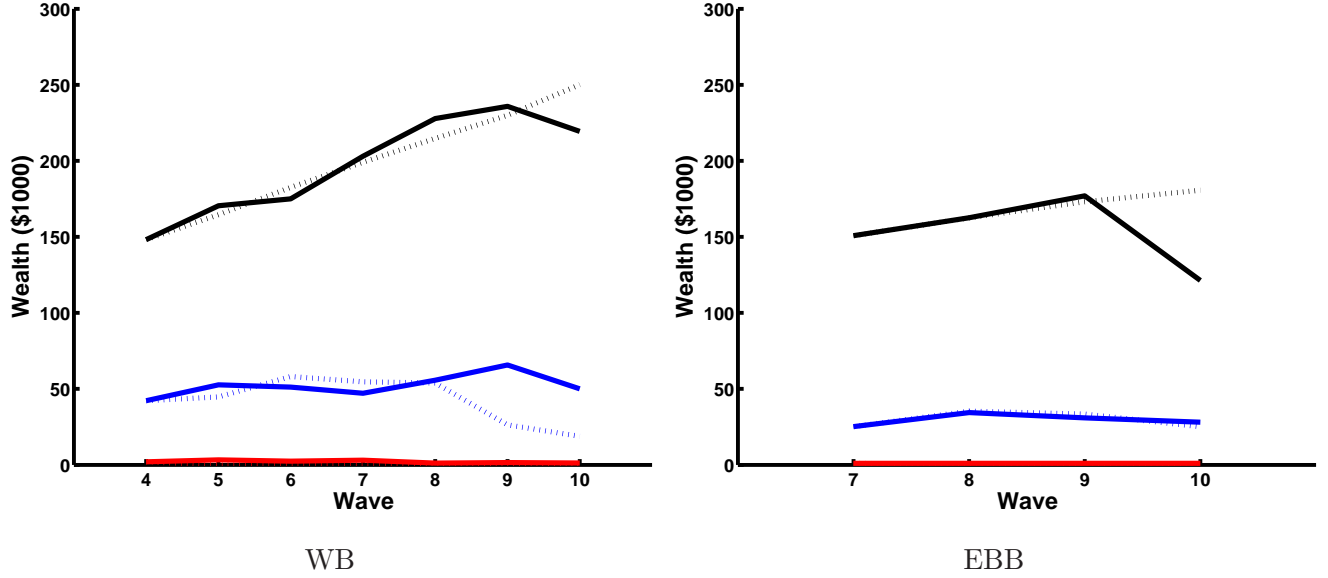

Figure E.2: Model results for WB/EBB cohort

E.2.3   HRS

Finally we turn to the longest observed cohort, the original HRS. Because we match initial wealth percentiles by design, considering saving patterns over the longest possible time-frame provides the largest opportunity for us to reject our model as a reasonable approximation of the data generating process. We consider two cohorts: HRS respondents between ages 50-54 and 55-59 at the time of entry into the HRS. The results of this exercise are presented in Figure E.3.

In this exercise we fair relatively poorly. While we again match the lack of wealth accumulation among poorer households, our model significantly overpredicts the spend down of wealth for the median households. Furthermore, at longer horizons we overpredict saving at the 75^{th} percentile, which again is likely partly explained by our failure to capture capital losses during the financial crisis.

In our main analysis our normal health consumption floor was very weakly identified as this parameter had no effect on the saving decisions of households in the VRI. Because of this weak identification, and because this consumption floor is likely relevant for households in the HRS, we next explore whether we can improve the empirical fit of our model by choosing this parameter differently. As an extreme, we set the normal health consumption floor equal to $1,000, effectively making government care a non-attractive options for non-ADL state individuals, and redo our simulations. These results are presented in Figure E.4.

Here we see that our model fit is much improved, and the median wealth level predicted by our model is flat and very close to its empirical counterparts at all waves. Although our model still overpredicts wealth accumulation of wealthy households, this again can be attributed to unmodeled capital losses. Thus, if we allow ourselves the freedom to freely calibrate the one parameter that is unidentified in the VRI sample and
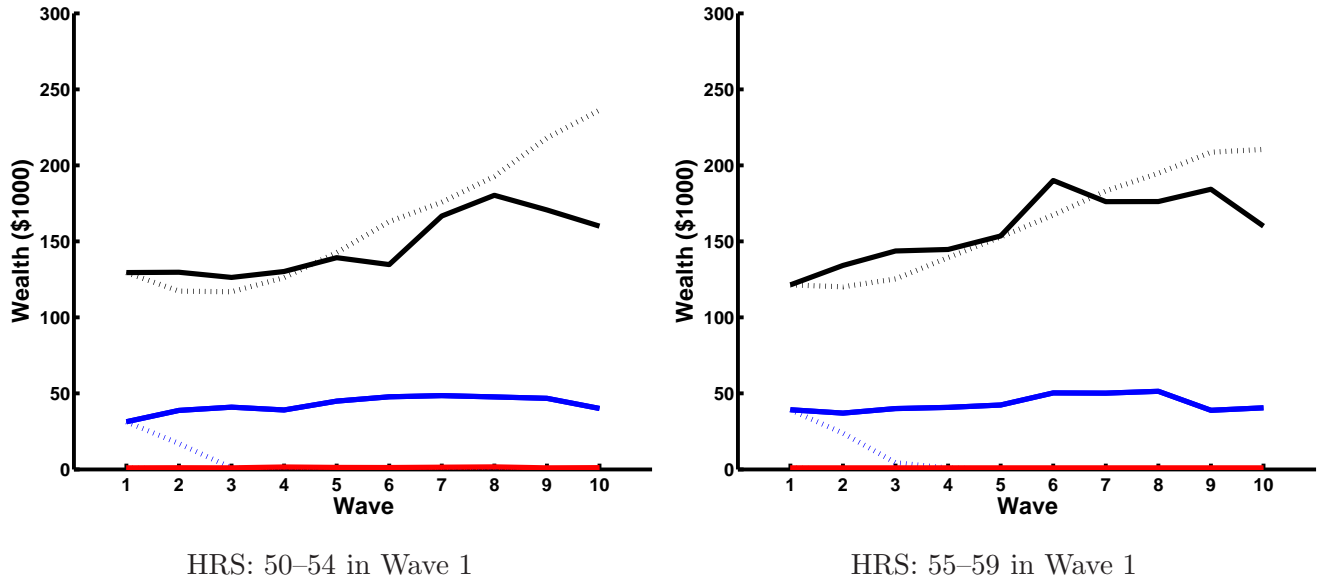
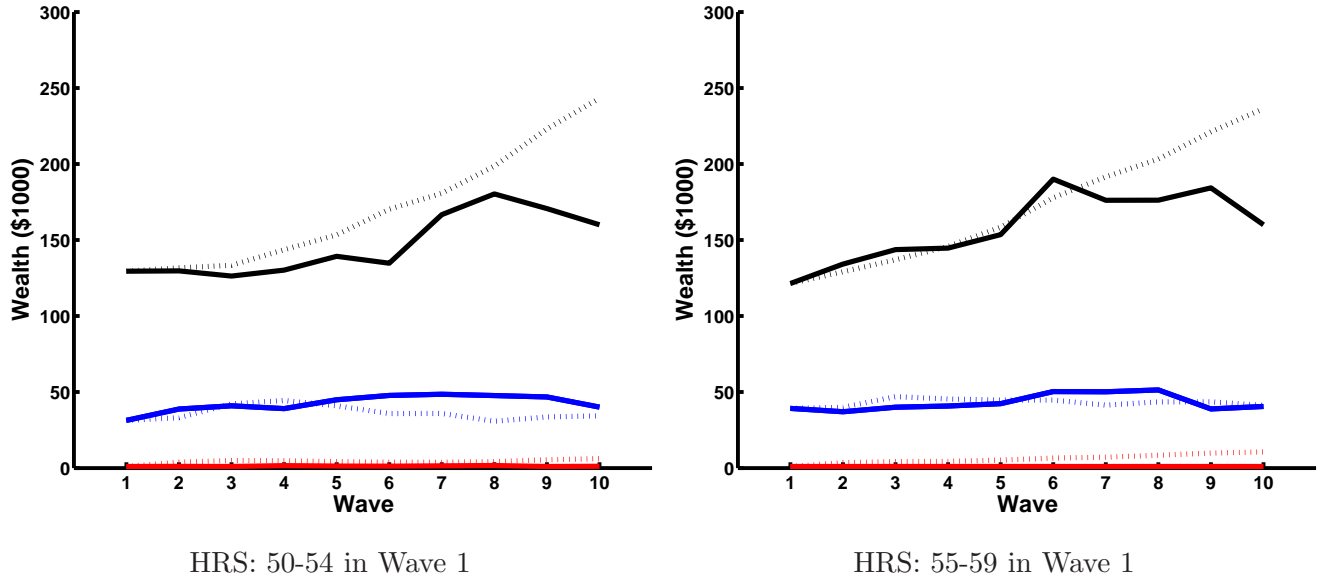Figure E.3: Model results for original HRS cohort

Figure E.4: Model results for original HRS cohort with consumption floor (ω_G) equal to $1,000.

acknowledge we do not capture large capital losses between 2008 and 2010 HRS waves, we again conclude that our estimated saving motives predict wealth holdings in line with those observed in the HRS.

E.3 Conclusion

The above analysis provides suggestive evidence that the saving motives estimated in the VRI are capable of explaining the wealth accumulation patterns we observe in the HRS, a representative US sample. In addition, this analysis suggests that the normal-health consumption floor parameter that we estimate to be inconsequential for the VRI is potentially important for households with less significant financial resources.