

NBER WORKING PAPER SERIES

LONG-TERM CARE UTILITY AND LATE IN LIFE SAVING

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

Working Paper 20973
<http://www.nber.org/papers/w20973>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2015

This research is supported by a program project grant from the National Institute of Aging P01-AG026571. The Vanguard Group Inc. supported the data collection through the VRI (Vanguard Research Initiative). Vanguard's Client Insight Group and IPSOS SA were responsible for implementing the VRI survey and provided substantial input into its design. The design of the VRI benefited from the collaboration and assistance of Wandi Bruine de Bruin (Leeds), Alycia Chen (CMU), Brooke Helppie McFall (Michigan), Minjoon Lee (Michigan), Ann Rodgers (Michigan), and Mi Luo (NYU) as part of the program project, from Annette Bonner (Vanguard), and Wendy O'Connell (IPSOS SA). The views expressed here are those of the authors and not those of their respective institutions. For documentation of the VRI, including a dynamic link to the survey instrument, see <http://ebp-projects.isr.umich.edu/VRI/>. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by John Ameriks, Joseph S. Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Long-Term Care Utility and Late in Life Saving

John Ameriks, Joseph S. Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti
NBER Working Paper No. 20973

February 2015

JEL No. D91,E21,H31,I10,J14

ABSTRACT

Older wealthholders spend down assets slowly. To study this pattern, the paper introduces health dependent utility into a model in which different preferences for bequests, expenditures when in need of long-term care (LTC), and ordinary consumption combine with health and longevity uncertainty to determine saving behavior. To help separately identify motives, it develops Strategic Survey Questions (SSQs) that elicit stated preferences. The model is estimated using new SSQ and wealth data from the Vanguard Research Initiative. Estimates of the health-state utility function imply that motives associated with LTC are significantly more important than bequest motives in determining late in life saving.

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

Joseph S. Briggs
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
jsb493@nyu.edu

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
Stanford University
Graduate School of Business
655 Knight Way
Stanford, CA 94305-7298
tonetti@stanford.edu

The Vanguard Research Initiative (VRI) website is available at:

<http://ebp-projects.isr.umich.edu/VRI>

Online Appendix: SSQs is available at:

<http://ebp-projects.isr.umich.edu/VRI/appendices/VRI-LTC-U-Appendix-SSQs.pdf>

Online Appendix: Modeling is available at:

<http://ebp-projects.isr.umich.edu/VRI/appendices/VRI-LTC-U-Appendix-Modeling.pdf>

Online Appendix: Estimation is available at:

<http://ebp-projects.isr.umich.edu/VRI/appendices/VRI-LTC-U-Appendix-Estimation.pdf>

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. Current explanations involve either bequest motives, precautionary motives associated with high late in life health and long-term care (LTC) expenses (see Kotlikoff (1988)), or both. Yet estimates of the importance of these motives range widely. Kopecky and Koreshkova (2014) and Lockwood (2014) find LTC expenses to be significant drivers 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 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 (2014) 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.

We provide new estimates of the relative importance of bequest and precautionary motives. We find precautionary motives associated with LTC to be significantly more important than bequest motives as drivers of late in life saving behavior. Saving motives driven by LTC are active for individuals with approximately less than \$50,000 in annual income and wealth less than \$400,000 (a large majority of the U.S. population). By contrast, our estimated bequest utility parameters suggest that the corresponding motive contributes only 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), 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 an 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. Allowing this additional flexibility reflects the distinctive nature of the spending options and desires when in need of help with the

¹Soto, 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.

activities of daily living. Appendix Figure E.1, using data from the Genworth (2013) survey and available at <http://www.longtermcare.gov>, documents that costs for a part-time home health aid range from \$35,000 to \$65,000 a year, while a private room in a nursing home ranges from \$55,000 to \$250,000 annually. We impose a minimum cost of private LTC that captures the reality that this state is associated with large and lumpy costs. We also model the option for individuals to utilize the publicly provided insurance against LTC and health risks. While clearly (as shown in Hubbard, Skinner, and Zeldes (1995) and Scholz, Seshadri, and Khitatrakun (2006)) 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 as shown in Brown and Finkelstein (2008), Braun, Kopecky, and Koreshkova (2013), and Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011). We incorporate these social insurance programs as means-tested consumption floors, with a separate provisions for LTC and non-LTC health states.

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 (2014), and Gourinchas and Parker (2002). There are two novel aspects of the computation and estimation of the model. First, the individual's value function is non-concave in wealth due to the interaction between free public care, public-care aversion, health-state utility, and minimum LTC expenditure levels. The methodology that we develop to efficiently compute optimal policies builds on the endogenous grid method of Fella (2014). Without such computational efficiency gains, estimation of the model would not be feasible. Second, 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. While our favored specification leverages both types of data and estimates the model jointly combining wealth and SSQs moments, we also provide separate estimates using moments of the wealth distribution alone and SSQs alone.² To our knowledge, ours are the first estimates

²Our estimation procedure is related to other strategies when 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 of 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 (2013), who also 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.

of a state-dependent utility function explicitly for the LTC state.³

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 HRS. There is little reason to believe that the results would be different in the broader population of similarly-situated older Americans. 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.

Ultimately, we examine the implications of the estimated preferences for savings and expenditure profiles. These estimates suggest that spending when in need of help with the activities of daily living is highly valued on the margin, and show the relatively greater importance of LTC-related than bequest-based saving motives.⁴ 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.

The paper is organized as follows. Section 2 develops the model, Section 3 describes the data in the VRI with a focus on the strategic survey questions, and Section 4 describes the estimation methodology that allows us to estimate the structural life cycle model without (and with) data on observed behavior. Section 5 presents our baseline parameter estimates obtained by matching both wealth and SSQ moments and examines the resulting behavioral implications of the estimated preferences. Section 6 compares our baseline estimates to those obtained by exclusively targeting wealth moments *or* SSQ moments to disentangle the relative contribution of the SSQs. Section 7 compares our baseline estimates to those found in the literature. Section 8 concludes.

2 The Model

2.1 Consumers

Consumers are heterogeneous over wealth ($a \in [0, \infty)$), income age-profile ($y \in \{y_1, y_2, \dots, y_5\}$), age ($t \in \{55, 56, \dots, 108\}$), gender ($g \in \{m, f\}$), health status ($s \in \{0, 1, 2, 3\}$), and health cost ($h \sim H_g(t, s)$ with support $\Omega_H(t, g, s)$). 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 corresponds with age 55 and T corresponds

³In 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 (2013) for an alternative method using Euler equations to estimate the effect of health on the marginal utility of consumption. Another method more closely related to ours is developed in Koijen, Van Nieuwerburgh, and Yogo (2015), who also estimate a health-state dependent utility function to analyze its effect on insurance demand. Given their sample selection, their “sick” state may be interpreted as being in need of LTC.

⁴In contrast to our findings, Koijen, Van Nieuwerburgh, and Yogo (2015), 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 (2014) does match moments of the cross-sectional wealth distribution and LTC insurance demand by cohort and finds strong bequest motives.

with age 108. Each period, consumers choose ordinary consumption ($c \in [0, \infty)$), savings (a'), expenditure when in need of long-term care ($e_{LTC} \in [\chi, \infty)$), and whether to use government care ($G \in \{0, 1\}$). The model groups consumers into five income groups with deterministic age-income profiles.⁵ Each consumer has a perfectly foreseen deterministic income sequence and receives a risk free rate of return of $(1 + r)$ on his savings. The only uncertainty an individual has is over health/death.

2.2 Government

The consumer 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 as described below. $G = 1$ if the consumer chooses to use government care and $G = 0$ if the consumer chooses not to use government care.

2.3 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. 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. h is a stochastic health expenditure that must be paid—essentially a negative wealth shock. Each period the consumer has to pay this health cost, h , where, $h \sim H_g(t, s)$.

If a consumer chooses to use government care when he does not need LTC (i.e., when $s = 0, 1$), then the government provides a consumption floor, $c = \omega_G$, that is designed to represent welfare.

A consumer needs LTC if he needs 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. If a consumer needs LTC ($s = 2$), then he 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, i.e., $e_{LTC} \geq \chi$ for those not using government care.⁶ In the model, government-provided care is loosely based on the institutions of Medicaid. If a consumer needs LTC and uses government care, the government provides $e_{LTC} = \psi_G$. The value ψ_G parameterizes the consumer's value of public care, since that parameter essentially determines the utility of an individual who needs LTC and chooses to use government care.

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 ($s = 2$) is valued differently than spending when a consumer does not need LTC. Utility when in need of LTC associated with

⁵The model abstracts from labor supply decisions, including retirement. These labor market decisions are taken into account through the exogenous income profiles.

⁶We treat all empirical heterogeneity in LTC expenditure as deriving from voluntary additional spending, as opposed to heterogeneous necessary expenditure. In future survey work, we are collecting information on the subjective expectations of the cost of LTC.

expenditure level e_{LTC} is

$$\theta_{LTC} \frac{(e_{LTC} + \kappa_{LTC})^{1-\sigma}}{1-\sigma}. \quad (1)$$

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:

$$v(b) = \theta_{beq} \frac{(b + \kappa_{beq})^{1-\sigma}}{1-\sigma}. \quad (2)$$

Utility Functions. When an individual is healthy or sick, his 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. 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. Increases in θ increase the marginal utility of a unit of expenditure, while increases in κ indicate the expenditure is valued as more of a luxury good. Negative κ can be interpreted as the expenditure being a necessity.

2.4 The Consumer Problem

The consumer takes r as given and chooses a' , c , e_{LTC} , and G to maximize utility. The consumer problem, written recursively, is,

$$\begin{aligned} V(a, y, t, s, h, g) = & \max_{a', c, e_{LTC}, G} \mathbb{I}_{s \neq 3} (1 - G) \{U_s(c, e_{LTC}) + \beta E[V(a', y, t + 1, s', h')]\} \\ & + \mathbb{I}_{s \neq 3} G \{U_s(\omega_G, \psi_G) + \beta E[V(0, y, t + 1, s', h')]\} + \mathbb{I}_{s=3} \{v(b)\} \\ \text{s.t.} \\ a' = & (1 - G)[(1 + r)a + y(t) - c - e_{LTC} - h] \geq 0 \\ e_{LTC} \geq & \chi \text{ if } (G = 0 \wedge s = 2) \\ e_{LTC} = & \psi_G \text{ if } (G = 1 \wedge s = 2) \\ c = & \omega_G \text{ if } (G = 1 \wedge (s = 0 \vee s = 1)) \\ b = & \max\{(1 + r)a - h', 0\} \\ U_s(c, e_{LTC}) = & \mathbb{I}_{s \in \{0,1\}} \frac{c^{1-\sigma}}{1-\sigma} + \mathbb{I}_{s=2} \theta_{LTC} \frac{(e_{LTC} + \kappa_{LTC})^{1-\sigma}}{1-\sigma} \\ v(b) = & \theta_{beq} \frac{(b + \kappa_{beq})^{1-\sigma}}{1-\sigma}. \end{aligned}$$

The value function has three components, corresponding to the utility plus expected continuation value of a living individual who does not use government care, that of one who does choose to use government care,

and the warm glow bequest utility of the newly deceased individual.⁷ Note that 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 health costs. 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 Online Appendix: Modeling.⁸

The option to use means-tested government care induces similar behavior to that studied in Hubbard, Skinner, and Zeldes (1995). Roughly speaking, 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 whose actions are predominantly affected by precautionary savings 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 along the life-cycle path, 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 5.2.

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

⁷Technically, 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.

⁸The 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.




Figure 1: Saving Policy Discontinuity (Assets a , \$1000s)

aside from the three different wealth levels: \$64,000 in the red dashed line, \$68,000 in the solid black line, and \$72,000 in the dotted blue line.

Thus, for low wealth individuals, the value of saving a small amount is higher than the value of saving a higher amount. The opposite is true for higher wealth individuals as shown in the top line. As presented in the middle line, there exists a wealth level for which the global maximum value jumps from the lower to the higher savings local maxima. It is around this wealth level where there will be a discrete jump in the optimal savings policy (although the value function will remain continuous).

These 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. Savings early in the person's life could be made 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.

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 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.




Figure 2: Wealth and Expenditure Profiles for Healthy Male

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 70 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. Consumption is almost smoothed over the life cycle, near constant with a slight positive trend. The modest 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.




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: Wealth and Expenditure Profiles with Publicly Provided LTC

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 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. For a discussion of the level of public-care expenditure (ψ_G), see Section 5.2. 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 \$700,000 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

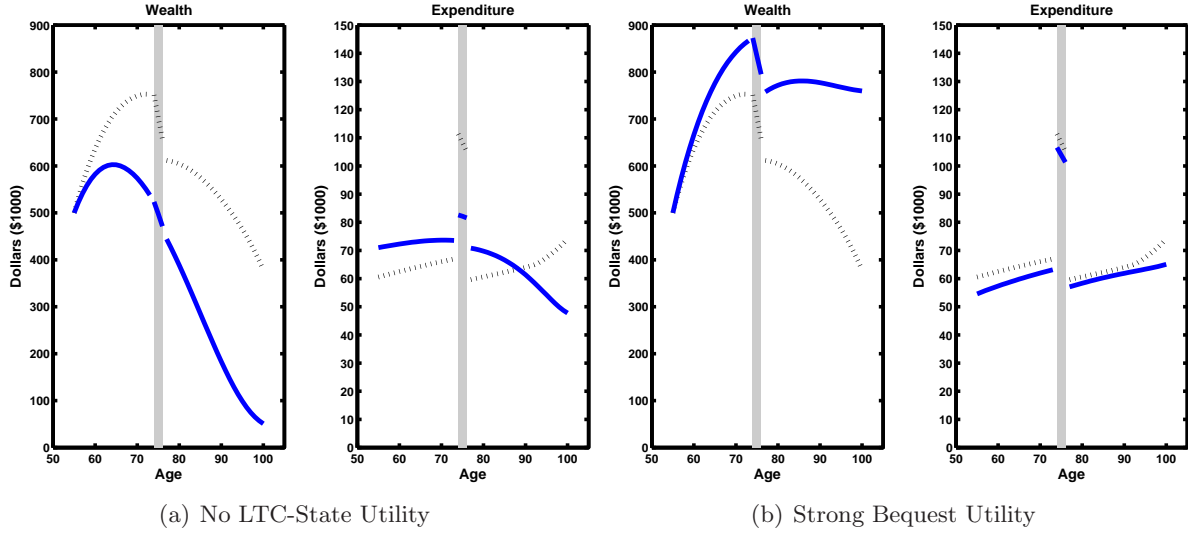


Figure 5: Wealth and Expenditure Profiles with Baseline (black dotted line) and Alternative (blue solid line) Preferences

just as spending when healthy ($\theta_{LTC} = 1$, $\kappa_{LTC} = 0$). 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 precautionary savings motives 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 ($s = 2$). This major change in expenditure patterns foreshadows that our estimated health-state utility function induces higher marginal utility from 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 parameters 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.

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 Data

In order to examine late in life wealth patterns, it is essential to have data on a population with resources large enough to face a saving decision. This paper draws on the newly developed Vanguard Research

Initiative (VRI) that combines survey and administrative account data. In this section we briefly describe the VRI, highlighting the advantages of the sample population for addressing the question at hand, and also documenting the strategic survey questions that we developed that are used to estimate the preference parameters of our model.

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 by construction 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 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 in various health states.

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 1,241 singles with no missing survey responses to mandatory questions.

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

	N	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

Wealth is measured as financial assets including defined-contribution retirement accounts less non-mortgage debt. The sample are single households meeting VRI sample screens: age 55 years and older; assets of at least \$10,000, and Internet access. 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.

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 broadly similar to these subsets of the 2012 HRS and 2013 SCF, representing individuals in roughly the upper half of the wealth distribution. Tables 1 and 2 compare 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 that the VRI sample contains significantly higher-wealth individuals compared to the HRS or SCF. In Table 2 we see that 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.

Table 2: Income Distribution Across Surveys: VRI-Eligible Single Households

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

Income is total household income (excluding distributions from defined-contribution pension plans). See previous table for sample.

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, health status, and, crucially, responses to SSQs that were specifically designed to identify parameters of the model just developed.

3.1 Strategic Survey Questions

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_{LTC} . Thus, if we want to develop direct measures of individual preferences, we need the survey respondent to provide us with information that identifies preference parameters. 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 choices 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 doing so, 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 than would otherwise be needed to identify these parameters.

SSQs build on stated preference methodology, although our applications differ in design and use. Though necessarily incomplete per se, our scenarios are significantly more detailed than those typically designed, and the parameters identified with these questions are not those of a random utility model, as is often the case. 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, and parameters are of deterministic utility functions. SSQs ask the respondent to comprehend and imagine complex scenarios. To make these tasks as easy as possible, we pay close attention to the presentation of the material. To ease respondent comprehension these questions are presented in four parts, with the implementation detailed below. See Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2015) for a detailed discussion of the SSQ design process and further analysis of individual responses to SSQs.

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. This design was refined with the assistance of Wandu Bruine de Bruin, a psychologist with expertise in survey design. The survey is carefully worded to not lead the respondent towards any specific answer. Next, 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. Although ultimately the model is estimated on responses from four types of SSQs, we illustrate the key features of SSQs by detailing a particular SSQ related to LTC (SSQ 2).

3.2 LTC SSQ

In the LTC SSQ, 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 purchase a portfolio of Arrow securities (x_1, x_2) given a relative price of x_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 optimal allocation that they choose, that solves the following problem, identifies preference parameters:

$$\begin{aligned} \max_{x_1, x_2} \quad & \pi \frac{x_1^{1-\sigma}}{1-\sigma} + (1-\pi) \frac{\theta_{LTC}(x_2 + \kappa_{LTC})^{1-\sigma}}{1-\sigma} \\ \text{s.t.} \quad & x_1 + p_2 x_2 \leq W \\ & x_1, x_2 \geq 0; \quad x_2 \geq -\kappa_{LTC}. \end{aligned} \tag{3}$$

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 (x_1, x_2) such that they are making a choice that we know corresponds to that in the optimization problem, but in a format in which they are capable of doing so.

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.

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).⁹ 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.

- 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 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.

⁹Although not presented here, 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.”

- 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 (2015), the vast majority of individuals answered almost all verification questions correctly before recording their responses to the strategic survey questions.

[Click here for complete scenario](#)

Please make your decision on splitting money into Plan C and Plan D by clicking on the scale below. To put more money in Plan C, move the slider to the left. To put more money in Plan D, move the slider to the right. The numbers in the boxes will change as you move the slider to let you know how much you will receive if you need long term care and if you do not.

Please move the slider to see how it works. When you are ready, place the slider at the split you want and click NEXT to enter your choice.



Figure 6: LTC SSQ 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 6—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 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. The implementation of the slider 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.

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

3.3 Overview of 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 savings motives. Like the previously presented SSQ, these questions present situations which individuals may be unlikely to ever face. However, the decisions 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 full text of the SSQs is available in Online Appendix: SSQs.

The first type of SSQ posed is a modified version of the Barsky, Juster, Kimball, and Shapiro (1997) style 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. 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 fractions they would be willing to risk, and then prompted to provide a point estimate of the largest fraction for which they would choose the 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

$$\frac{W^{1-\sigma}}{1-\sigma} = 0.5 \frac{(2W)^{1-\sigma}}{1-\sigma} + 0.5 \frac{((1-\xi)W)^{1-\sigma}}{1-\sigma}, \quad (4)$$

with the choice of ξ identifying σ . This question is then repeated for an individual with \$50,000 in income.

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 do 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 and \$200,000 of wealth. SSQ 3 maps to the following optimization problem:

$$\begin{aligned} \max_{x_1, x_2} \quad & \frac{\theta_{LTC}(x_1 + \kappa_{LTC})^{1-\sigma}}{1-\sigma} + \frac{\theta_{beq}(x_2 + \kappa_{beq})^{1-\sigma}}{1-\sigma} \\ \text{s.t.} \quad & x_1 + x_2 \leq W \\ & x_1, x_2 \geq 0; \quad x_1 \geq -\kappa_{LTC}; \quad x_2 \geq -\kappa_{beq}. \end{aligned} \quad (5)$$

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, SSQ 4 asks respondents to provide the W that satisfies:

$$\frac{\theta_{LTC}(\psi_G + \kappa_{LTC})^{1-\sigma}}{1-\sigma} + \frac{\theta_{beq}(W + \kappa_{beq})^{1-\sigma}}{1-\sigma} = \frac{\theta_{LTC}(x_1 + \kappa_{LTC})^{1-\sigma}}{1-\sigma} + \frac{\theta_{beq}(W - x_1 + \kappa_{beq})^{1-\sigma}}{1-\sigma},$$

where x_1 is the optimal policy when there is no public care, as calculated in SSQ 3.

3.4 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 3. In Online Appendix: SSQs, we document the text of the four SSQ questions and the survey responses to all nine SSQ variants.




Figure 7: Responses to Bequest SSQs

In Figures 7(a), 7(b), and 7(c) we observe how individuals trade off leaving money as a bequest and having wealth when in the ADL state, across wealth levels of \$100,000, \$150,000, and \$200,000, respectively, representing the three variants of the bequest SSQ 3. The figures are histograms showing the amount that the individual would allocate towards the ADL state. Here, we clearly see that individuals react to the wealth level, as we note that many respondents allocate almost all of their portfolio to the ADL state when wealth is \$100,000. When wealth is \$150,000, many reveal that they were severely restricted in the amount they desire to have in the ADL state, as evidenced by the large mass of individuals responding with allocations to the ADL state above \$100,000. Similar patterns repeat when wealth is \$200,000, with significant response mass above \$150,000. Furthermore, even at \$200,000, many individuals gave zero bequest.¹⁰ This finding suggests that many individuals 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_{LTC} and κ_{beq} . It is harder to read directly how these responses translate into θ_{LTC} 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 when compared to bequests. Having provided some descriptive evidence on how the SSQ responses will inform parameter estimates, we turn to the formal estimation strategy.

4 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 (2014), Gourinchas and Parker (2002), and Laibson, Repetto, and Tobacman (2007) to estimate parameter set Γ . $\Gamma := [\Xi, \Theta]$ is divided into

¹⁰Note that there is some heaping at round numbers, as is typical in surveys.

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, Θ , consisting of preference parameters that are estimated using moments generated by simulating the structural model. After calibrating the first stage estimates at their estimated values, the second stage parameters are estimated using an MSM procedure by minimizing the distance between model implied moments and their empirical counterparts.

4.1 First Stage Estimates

4.1.1 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. 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.

4.1.2 Health Transitions

Health transitions are estimated using HRS waves 2 through 10, with the defined health states constructed 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 $s = 1$).¹¹ 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 5 activities of daily living and asks whether respondents receive help with any of the 5 activities. If the respondent answers yes to any of these questions, then we define that respondent to be in the ADL health state. Although alternative LTC/ADL state definitions, such as having spent time in a nursing home, are feasible given the available data, we choose this state definition since it is most consistent with the ADL definition presented in the VRI survey. The questions necessary to make this health state assignment are not available in the 1992 survey, so we exclude this wave from the health transition estimates. 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.

4.1.3 Health Expense

The health expense distributions are estimated from the 2010 HRS distribution. 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 estimate the mean and standard deviation of mandatory out of pocket health expenditures conditional on age, gender, and health state. For more information on the estimation of costs and the resulting estimates see Appendix A.2.

¹¹This classification follows criteria presented in the RAND HRS. 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.

4.2 Second Stage Estimates

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

Empirical moments are defined by the data and taken as given. To generate simulated moments, we first solve the model to generate optimal decision rules as a function of all relevant state variables. Next, we sample (with replacement) a large number of individuals, N , from the observed data X . Then, for each sampled 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 decisions to construct simulated population moments $m(\hat{\Xi}, \Theta, X)$. The second stage estimator with weighting matrix W is:

$$\hat{\Theta} = \arg \min_{\Theta} g(\hat{\Xi}, \Theta, X)' W g(\hat{\Xi}, \Theta, X). \quad (6)$$

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 the objective function defined in equation 6 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 stage covariance matrix $\hat{\Omega}_1$. Since many of the off-diagonal elements of the inverse of the calculated covariance matrix are computationally close to zero, we restrict the second stage weighting matrix to consist of only the diagonal elements of $\hat{\Omega}_1^{-1}$. Thus we denote the second stage weighting matrix as:

$$\hat{W} = \text{diag}(\hat{\Omega}_1^{-1}). \quad (7)$$

We then minimize equation 6 using \hat{W} to estimate the final parameter set $\hat{\Theta}$. Asymptotic properties of the estimator and a derivation of the standard errors are presented in Appendix B.

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 savings models. A second set of moments is derived from SSQ responses. In the remainder of this section we describe in detail the construction of both sets of moments and present the comprehensive moment set we target in our baseline estimation. In Section 5 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 6.

4.2.1 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 (2014)), the first moment set consists of asset percentile levels, conditional on a set of state variables x . The empirical moment conditions for wealth percentile p conditional

on state variables x are denoted a_x^p , while $a_i(\hat{\Xi}, \Theta, X_i)$ denotes simulated individual i 's wealth holdings when he has state variables X_i and the model is specified with parameters $\hat{\Xi}$ and Θ . We denote the wealth moments conditional on x as

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

an expression which can easily be converted to an unconditional expectation through the Law of Iterated Expectations. For more information on this, and formal derivation of the asymptotics, please see Online Appendix: Estimation.

As a baseline, we define the moment conditions as the 25th, 50th, and 75th percentiles (a_t^p) conditional on age t . We aggregate the age profiles into disjoint three year intervals, so that $t \in \{55-57, 58-60, \dots, 88-90\}$.¹² We thus define a_t^p as the empirical p^{th} percentile for those aged t , $t+1$, or $t+2$, with the simulated percentile defined accordingly. 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.

4.2.2 Strategic Survey Question Moments

The second moment set is constructed entirely from SSQ responses. The usefulness of non-behavioral measurements in structural estimation has been demonstrated in papers such as van der Klaauw (2012) and Blass, Lach, and Manski (2010). However, unlike many of these studies, we do not utilize subjective expectations data that require specification of an expectation formation process or use stated choices from discrete options, but instead estimate the model using stated strategies in hypothetical situations.¹³

Since the SSQ data are non-standard, we discuss in some detail how they will influence inference by describing how SSQ responses identify parameters of the model. The structural model has 8 preference parameters: relative risk aversion parameter, σ , LTC state utility parameters, θ_{LTC} and κ_{LTC} , bequest utility parameters, θ_{beq} and κ_{beq} , public-care aversion, ψ_G , healthy state consumption floor, ω_G , and intertemporal discounting, β . Using responses to the four types of SSQs described in Section 3.3 and their iterations at different wealth levels and state probabilities, we are able to identify all parameters except for the discount factor, β , and the consumption floor, ω_G .

We use the parameterized model of SSQ response to identify individual parameters. In each situation, we ask individuals to make a trade-off between certain states, and record their decisions. 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 3 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 states. Since relative risk aversion parameter σ 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 help with ADLs. This trade-off is optimally determined (abstracting from

¹²We 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.

¹³Another point of departure is that we use deterministic utility functions, as opposed to the random utility specifications frequently used in this literature to obtain a closed form likelihood function convenient for estimation.

corner solutions for the moment) by equating marginal utility in the healthy state as determined by σ with marginal utility when in need of help, as determined by σ , θ_{LTC} , and κ_{LTC} . Utilizing the observed trade-offs at different wealth levels and state probabilities, we thus are able to identify the θ_{LTC} and κ_{LTC} 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 σ , θ_{LTC} , and κ_{LTC}) and the marginal value of wealth allocated towards a final bequest (determined by σ , θ_{beq} , and κ_{beq}), while in the fourth 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). Because these trade-offs vary with wealth levels, we are able to use these questions to identify the underlying parameters that map survey responses to parameters.

As an example of this identification, recall the optimization problem presented in equation 3. The optimal decision rule of this problem is given by:

$$x_2 = \begin{cases} 0 & \text{if } W^{-\sigma} - (1 - \pi)\theta_{LTC}(k_{LTC})^{-\sigma} > 0 \\ W \frac{\left(\frac{1}{(1-\pi)\theta_{LTC}}\right)^{\frac{-1}{\sigma}}}{1 + \frac{1}{\pi}\left(\frac{1}{(1-\pi)\theta_{LTC}}\right)^{\frac{-1}{\sigma}}} - \frac{k_{LTC}}{1 + \frac{1}{\pi}\left(\frac{1}{(1-\pi)\theta_{LTC}}\right)^{\frac{-1}{\sigma}}} & \text{otherwise.} \end{cases} \quad (9)$$

In the above expressions, π and W are specified in the survey question, and we rely on varying these values in different question iterations to ensure identification. In addition, SSQ 1 provides a measure of individuals' relative risk aversion, so this question is needed primarily to measure θ_{LTC} and κ_{LTC} .

Thus, 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_{LTC} and κ_{LTC} , given σ . Combining these with the results of SSQ 1 we are thus able to identify all three parameters from responses to these questions. Finally, since the survey poses three different variants of SSQ 2 (as well as two variants of SSQ 1), we have an overidentified system suitable for estimation. The conditional identification of θ_{beq} , κ_{beq} , and ψ_G by SSQ 3 and SSQ 4 can be shown by similar logic, thus ensuring that all parameters are jointly identified from the set of SSQ responses. See Online Appendix: Modeling for a derivation of the link between survey responses and preference parameters for all SSQs.

Having demonstrated how preference parameters are identified by SSQ responses, we now present the moments that we use to estimate the model. Formally, we match the empirical mean of each SSQ variant ($m = 1...9$). To generate the simulated means, we denote simulated individual i 's response to SSQ variant m as $s_m^i(\Theta)$. We then write each moment as:

$$g_m(\hat{\Xi}, \Theta, X) = \mathbb{E} [s_m^i(\Theta) - z_m^i], \quad (10)$$

providing us with 9 SSQ moments.

Table 3: Link Between Parameters and SSQs

	Question	Motives	Scenario Parameters	Preference Parameters
SSQ 1	Lottery over spending	Ordinary consumption	(a) $W = \$100K$ (b) $W = \$50K$	σ
SSQ 2	Allocation between ordinary and ADL states	Ordinary consumption and ADL expenditure	(a) $W = \$100K, \pi = 0.75$ (b) $W = \$100K, \pi = 0.50$ (c) $W = \$50K, \pi = 0.75$	$\sigma, \theta_{LTC}, \kappa_{LTC}$
SSQ 3	Allocation between ADL and bequest states	ADL expenditure and bequest	(a) $W = \$100K$ (b) $W = \$150K$ (c) $W = \$200K$	$\sigma, \theta_{LTC}, \kappa_{LTC}$ $\theta_{beq}, \kappa_{beq}$
SSQ 4	Indifference between public and private LTC	ADL expenditure and bequest	(a) Public Care Available	$\sigma, \theta_{LTC}, \kappa_{LTC}$ $\theta_{beq}, \kappa_{beq}, \psi_G$

4.2.3 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.

More specifically, let $g_V(\hat{\Xi}, \Theta, X)$ 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. Because matching moments of different magnitude presents computational difficulties, we normalize the SSQ moments to the unit interval by dividing each SSQ variant m 's simulated and empirical responses by the maximum feasible response, W_m , as given by the budget set.¹⁴ Note that because

$$\begin{aligned}
 g_m(\hat{\Xi}, \theta, X) &= \mathbb{E} [s_m^i(\Theta) - z_m^i] \\
 &= \mathbb{E} \left[\frac{s_m^i(\Theta)}{W_m} - \frac{z_m^i}{W_m} \right] \\
 &= 0,
 \end{aligned} \tag{11}$$

this normalization is only a rescaling, and we thus use the same notation $g_S(\hat{\Xi}, \Theta, X)$ to denote normalized

¹⁴Because there is no maximum response to SSQ 4, we set W_m to be the windsorized 95th percentile of the raw response distribution.

moment conditions for the joint estimation. As a result, we denote the joint estimation’s moment set as

$$g_J(\hat{\Xi}, \Theta, X) = \begin{bmatrix} g_V(\hat{\Xi}, \Theta, X) \\ g_S(\hat{\Xi}, \Theta, X) \end{bmatrix}. \quad (12)$$

Formal derivation of the asymptotic properties of this distribution and implementation of the simulation procedure are presented in Online Appendix: Estimation.

In addition, when conducting the two-step MSM estimation, we adjust the weighting matrix to control the weight that is placed on each of the moment sets $g_V(\hat{\Xi}, \Theta, X)$ and $g_S(\hat{\Xi}, \Theta, X)$ in both the first (with identity weighting matrix) and second (with diagonal of optimal weighting matrix) steps. We do so by element-by-element multiplication of the diagonal elements of the identity weighting matrix that correspond to the wealth moments with a vector λ_1 in the first stage and the diagonal elements of the second stage weighting matrix with a vector λ_2 . Because it is difficult to match exactly both sets of moments simultaneously, introducing this parameter 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 design, SSQ moments provide strong identification of the fundamental parameters, and thus are in general weighted higher than wealth moments by the optimal weighting matrix. 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 the additional parameter λ_2 to arrive at our baseline configuration.¹⁵ By introducing some degree of flexibility in the estimation algorithm, our choice of λ_2 permits us to improve the fit of wealth moments, particularly at higher wealth levels, without forcing the SSQs moments to be missed by too much.¹⁶ Furthermore, parameter estimates are quantitatively similar for a wide range of λ_2 , suggesting the implications of preferences for saving, expenditure, and bequests are robust to this choice.

5 Parameter Estimates

Before presenting the parameter estimates, we show the fit of the model to the SSQ and wealth moments.

5.1 Model Fit

Figures 8(a) and 8(b) document the model fit from our baseline estimation that jointly targets both wealth and SSQ moments. We present the SSQ moments on a $[0, 1]$ scale, by normalizing the mean response by the maximum possible response. Overall, the fit is quite good, as the 25th and 50th percentiles of the wealth distribution are matched extremely well, as are the SSQ moments. However, a general feature of the estimation results is that it is difficult for the model to jointly match the SSQ moments and generate large enough savings for high wealth individuals. We address this issue in detail in Section 6.

5.2 Estimated Parameter Values

Table 4 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

¹⁵In our baseline estimation we set λ_1 to multiply all wealth moments by 20 and set λ_2 to multiply the 25th and 50th percentile wealth moments by 20 and the 75th percentile moments by 40.

¹⁶When using the optimal weighting matrix, the only substantial difference is that the fit on the 75th percentile is worse, due to slightly lower risk aversion and lower θ parameters. See Appendix D for results using the optimal weighting matrix.

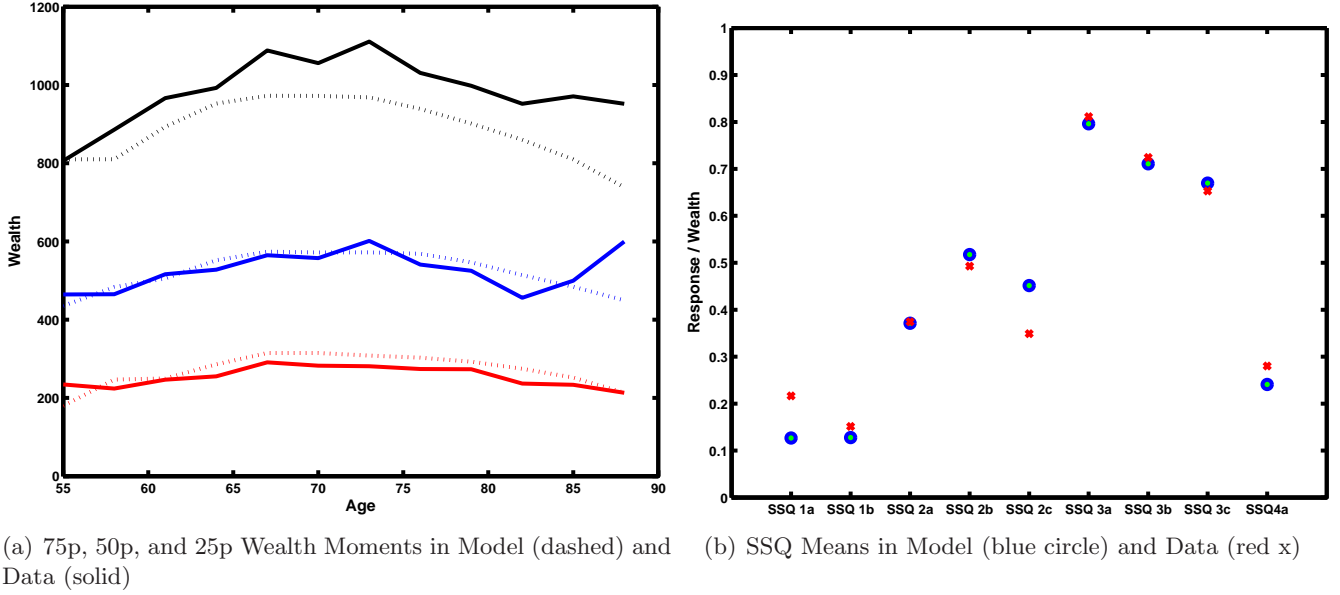


Figure 8: Model Fit When Jointly Targeting Wealth and SSQ Moments

(σ) 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_{LTC} and κ_{LTC}) presented are the first we are aware of to estimate this functional form applied to LTC. Third are the bequest parameters θ_{beq} and κ_{beq} . Finally, ψ_G controls the degree to which there is public-care aversion. We restrict $\beta = 0.97$, as the empirical strategy was not designed to estimate this parameter.¹⁷

Table 4: Parameter Estimates

Joint Estimation: Baseline Model							
σ	θ_{LTC}	κ_{LTC}	θ_{beq}	κ_{beq}	ψ_G	ω_G	β
5.85	1.57	-45.65	0.59	7.88	85.11	29.45	0.97
(0.06)	(0.13)	(0.01)	(0.71)	(0.25)	(1.74)	(36.67)	—

Standard errors are reported in parentheses.

The coefficient of relative risk aversion is estimated to be 5.85. 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 model related to ours and very different data De Nardi, French, and Jones (2010) estimate $\sigma = 3.8$, while Kimball, Sahm, and Shapiro (2008) estimate a mean σ 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 from wealth in the state of the world when an individual needs LTC and when LTC is not needed. The estimated $\kappa_{LTC} = -45.65$ is negative and large,

¹⁷ 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.

suggesting LTC is viewed as a necessary good with a spending floor of about \$45,000.¹⁸ $\theta_{LTC} = 1.57$ implies a high marginal utility from expenditure in the LTC state, especially when combined with κ_{LTC} . Note that if $\kappa_{LTC} = 0$, relative to ordinary consumption, $\theta_{LTC} < 1$ implies that LTC expenditures provide less marginal utility for each dollar spent while $\theta_{LTC} > 1$ implies a higher marginal utility for each dollar expenditure. Individuals with $\theta_{LTC} < 1$ and $\kappa_{LTC} < 0$ would view expenditure when in need of LTC as a lumpy cost and optimally desire to consume the consumption equivalent that this necessary expenditure provides, 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 $\theta_{LTC} > 1$, 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.

Compared to utility from expenditure in the LTC state, individuals receive less utility from leaving a bequest. First, $\kappa_{beq} = 7.88$, which suggest bequests are viewed as a luxury good. Together with a positive κ_{beq} , the low estimate of $\theta_{beq} = 0.59$ implies a low marginal utility from bequests, relative to LTC. This is very different from estimates in the literature like those from Lockwood (2014) and Koijen, Van Nieuwerburgh, and Yogo (2015), who find much stronger bequest motives.

Because expenditure when in need of LTC is so highly valued, $\psi_G = \$85,110$ is estimated to be large compared to previous estimates in the literature. This estimate implies an equivalent utility level to that provided by a \$35,960 government provided public-care expenditure in a model with the same risk aversion but without state dependent preferences.¹⁹ Although ψ_G is larger than estimates in previous studies, its expenditure equivalent of \$35,960 is close to the typical monetary amount provided by Medicaid. The consumption floor (ω_G) is almost irrelevant for this study, since the large majority of individuals in the sample have sufficient financial resources that welfare is an unlikely option. Consequently, the estimate of ω_G at about \$30,000 is imprecise, and also inconsequential for the simulations of the model.

5.3 Behavioral Implications of Estimated Preferences

To help further interpret the parameter values in economically meaningful ways, in this section we use the model to document behavior induced by the estimated preferences.

5.3.1 Fit for Non-Targeted Wealth Moments

To demonstrate the saving behavior across the wealth distribution implied by the baseline estimates, Figure 9 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. Hence, that preference parameters estimated by targeting the higher wealth moments successfully replicate savings patterns of the less affluent, the parameters estimated from the VRI sample are applicable to other populations. 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

¹⁸In our baseline, we set the minimum private LTC expenditure $\chi = \$40,000$. Note that the estimation results are independent of the choice of χ for any $\chi \leq -\kappa_{LTC}$, as is the case for all estimates in this paper.

¹⁹To 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: $\frac{\bar{\psi}^{1-\sigma}}{1-\sigma} = \theta_{LTC} \frac{(\psi_G + \kappa_{LTC})^{1-\sigma}}{1-\sigma}$.

institutional and market constraints—face the same economic trade-offs as members of the VRI sample. In contrast, as previously shown, the model predicts less wealth than appears in the data at the 75th percentile, and Figure 9 documents that this model undersaving is present and even stronger at the top wealth decile. 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 9: 90p, 75p, 50p, 25p, and 10p Wealth Moments in Model (dashed) and Data (solid)

5.3.2 Implied Demand for Financial Products

In addition to analyzing the implied life cycle savings patterns, we can offer agents in the model the option to purchase different financial products, with the resulting demand further illuminating the relative strength of different saving motives. By observing the demand for annuities and for insurance against the need for long-term care, we can further quantify the impact of the estimated preferences.

Demand for Annuities. We first analyze how the demand for annuities varies across the population and how different features of preferences and the economic environment interact to determine annuity demand. To do so, we calculate the amount of annuity income an individual would like to purchase (or sell) if given one-time access to an annuity at a particular age. An annuity is a financial product that provides a price to determine a mapping between wealth and riskless income streams. An individual can increase the income he receives for the rest of his life by $\$X$ per year by spending the lump sum of $\$Y$, with the annuity price determining the ratio of X to Y . Annuities are priced conditioning on the individual's age, gender, and health status, and are actuarially fair valued if an issuer of the asset would expect to make zero profit on the contract. For the purposes of this exercise, we treat individuals as if their future expected income was collateralizable, thus allowing individuals to pick their optimal ratio of wealth to income. This thought experiment allows us to determine if individuals wish to deannuitize, by decreasing their future income flow for some lump-sum wealth gain. For exposition, we focus on the annuity demand of a male in good health at age 65.

Figure 10 presents demand for an annuity that has a 10 percent premium (load) above the actuarially

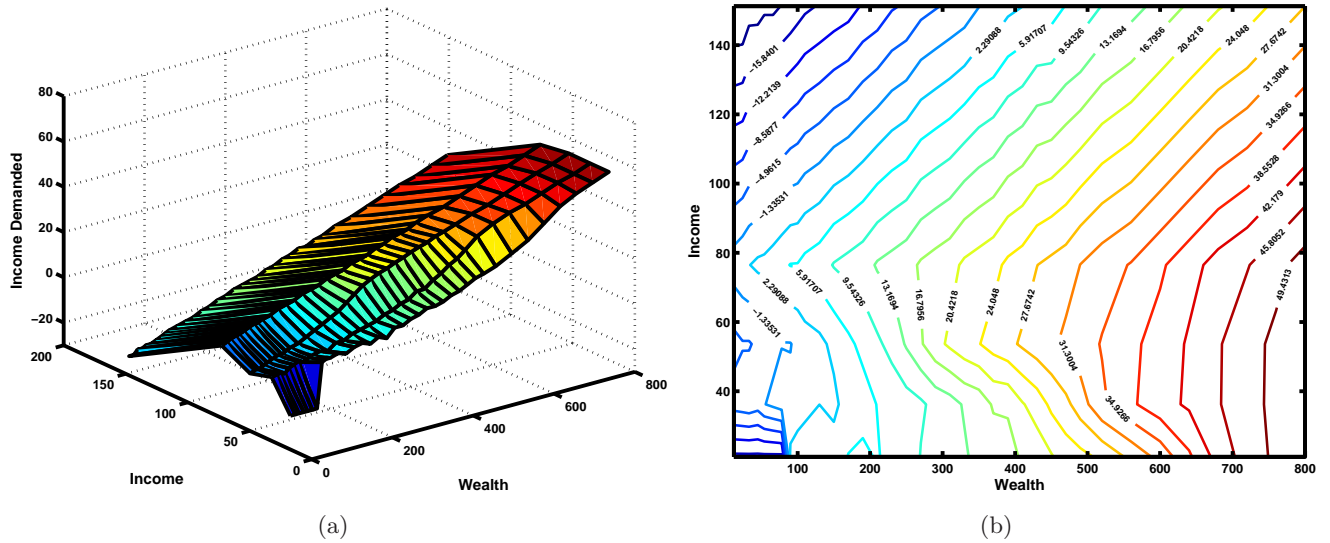


Figure 10: Demand for Annuities (10% Premium)

fair price, which is a representative load for the U.S. private annuity market documented by Brown (2007). Calculation of annuity prices and annuity demand are described in Appendix C.²⁰ The horizontal plane spans the income-wealth space, while annuity demand is graphed on the vertical axis. (The panel on the right is a contour plot that highlights where the demand shows the greatest change.) The wealth axis denotes the amount of wealth an individual has at age 65 and the income axis denotes the income an individual is receiving at age 65. The vertical axis denotes the amount of annual income an individual would purchase. First note the standard results that annuity demand is increasing in wealth. Higher wealth individuals have more wealth to annuitize, so even if the fraction of wealth annuitized was constant across wealth levels, annuity demand would be increasing in wealth. The slope of the annuity demand function with respect to wealth is a measure of the differences in the fraction of wealth annuitized across wealth levels. For high wealth and high income individuals, annuity demand is decreasing in income, as an annuity is a perfect substitute for income (given that we do not model stochastic income). Annuity demand is noticeably lower for the low wealth and low income individuals. These are individuals with strong precautionary savings motives induced by the potential need for LTC and the strong utility from spending when in need of LTC.²¹ The existence of this sharp decrease in annuity demand is driven by the same incentives that cause the optimal savings policy of an individual to be discontinuous in wealth levels.

In a model without LTC utility and a public care option, there would not be this chasm of noticeably reduced annuity demand for lower wealth and income individuals. Since in the baseline parameter set $\theta_{LTC} > 1$ and κ_{LTC} is negative and large, the health-state utility function increases the marginal value of expenditure when the individual needs LTC, and thus decreases annuity demand for those lower income and wealth individuals most affected by precautionary saving motives. Thus, the potential need for LTC, together with the minimum expenditure and the differential utility in the LTC state, drives annuity demand and late in life saving behavior for a large fraction of the population. This is true even for people with

²⁰Given prices, the amount of annuitized wealth can easily be calculated from the amount of yearly income purchased. See Kojien and Yogo (2015) for a discussion of why insurance may not be actuarially fair in practice.

²¹See Laitner, Silverman, and Stolyarov (2014) for an analytically tractable model that highlights many of these features.

moderate income and high wealth (e.g., people with \$50,000 in income and \$400,000 in wealth).

From the perspective of the estimated model, a very large fraction of the U.S. population has near zero or negative demand for a typical annuity available in the market, while the wealthier individuals still desire to annuitize a large fraction of wealth. We view the results as suggesting that LTC and the utility derived from expenditures when in need of LTC contribute substantially to the lack of demand for annuities in a large fraction of the population. We believe there are other motives missing from the model that are driving the observed behavior of the very wealthy who also avoid purchasing annuities in practice.

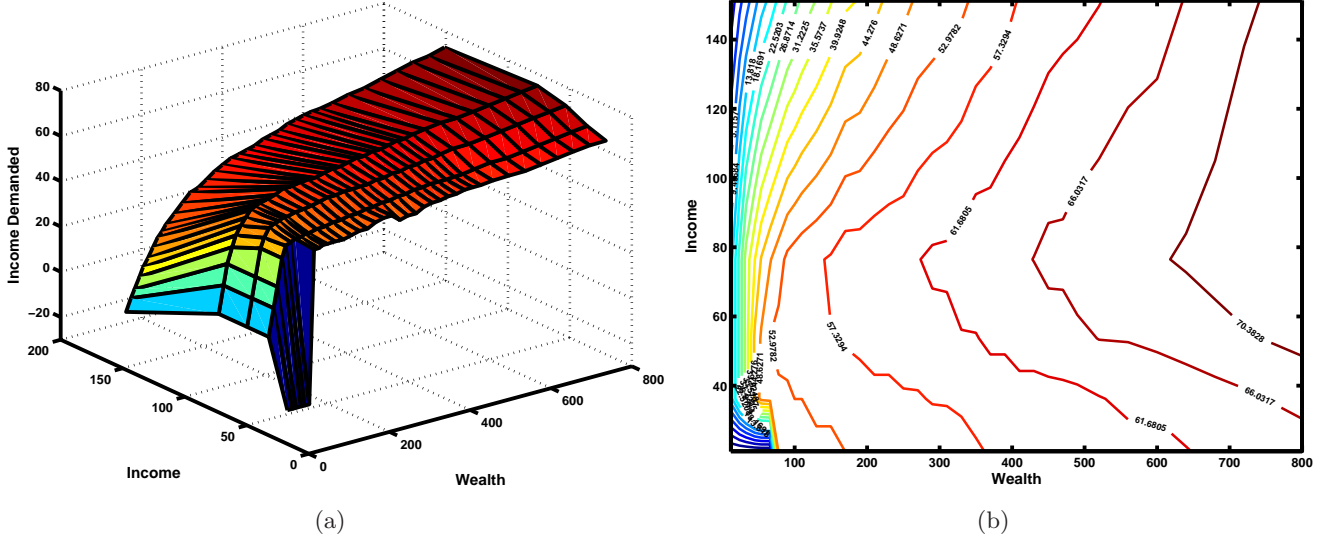


Figure 11: Demand for ADLI Insurance (10% Premium)

Demand for Activities of Daily Living Insurance. We can do a similar analysis for an Arrow security that provides income when an individual needs help with the activities of daily living ($s = 2$). To distinguish this from long-term care insurance (LTCI), which is a product currently offered in the market with many characteristics that make it less desirable than the product we are modeling, we call this activities of daily living insurance (ADLI). ADLI simply delivers income in states of the world in which the purchaser needs LTC. Since the differences between LTCI and ADLI are vast, one should not use moments related to LTCI holdings in the population to estimate parameters of the model or to use the small fraction of LTCI holdings as evidence that a strong precautionary savings motive can not be prevalent in the U.S. population.

Figure 11 presents the amount of ADLI income a 65 year old healthy male would purchase if the ADLI was priced at a 10 percent premium above actuarially fair value. See Appendix C for a description of how ADLI demand and prices are determined. One striking feature is that there is substantial demand for insurance against the LTC health state, generally with higher levels of purchased income compared to annuities across the wealth and income distribution. Furthermore, the fact that demand is roughly stable across wealth and income outside of the low wealth region of negative demand is evidence that there is strong demand for ADLI and that there are strong precautionary savings motives induced by LTC risk and LTC utility. Roughly speaking, the flat demand for ADLI at a high level suggests that there is a desired level of insurance that all individuals spend to achieve, even those with not much wealth. The negative demand for lowest-wealth individuals is also illuminating. These are the people who know it is likely that they will

use government care when they need LTC. Thus, they are trading off increased wealth at age 65 with an increasing chance that they will give up the ability to obtain private long-term care and use means-tested public care. The higher probability of using public care is somewhat offset by the lower level of appropriable resources available when means-tested public care is used. Overall, the strong demand for ADLI confirms that the estimated preferences quantitatively imply a strong impact of late in life health risks on financial behavior.


6 Alternate Estimates Targeting Wealth or SSQ Moments Separately

In our preferred 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 all preference parameters solely from the SSQ responses. Thus, unlike other studies, 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.

One advantage of splitting the data into SSQ and wealth data is that it permits out of sample model verification. That is, we can examine the saving behavior implied by SSQ responses and the SSQ responses implied by saving behavior by analyzing the fit of the model to these various targeted and non-targeted moments.

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 12: Model Fit When Exclusively Targeting Wealth Moments

The first set of figures documents model fit when only wealth moments are targeted. Figure 12(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. The wealth profiles at the 25th, 50th, and 75th percentiles are matched extremely well, with some undersaving generated by the model at older ages. In the last column of Table 5 we present a test statistic for a Hansen J -test of over-identification, which does not reject the null hypothesis that the over-identifying restrictions are true. Yet, as displayed in Figure 12(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 13(a) and 13(b), the SSQ moments are hit almost exactly. Similar to the case in which we targeted wealth moments exclusively, an over-identification test fails to reject the model. The success in fitting these moments should not be surprising, given that SSQs were designed to ensure identification. Yet, 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.

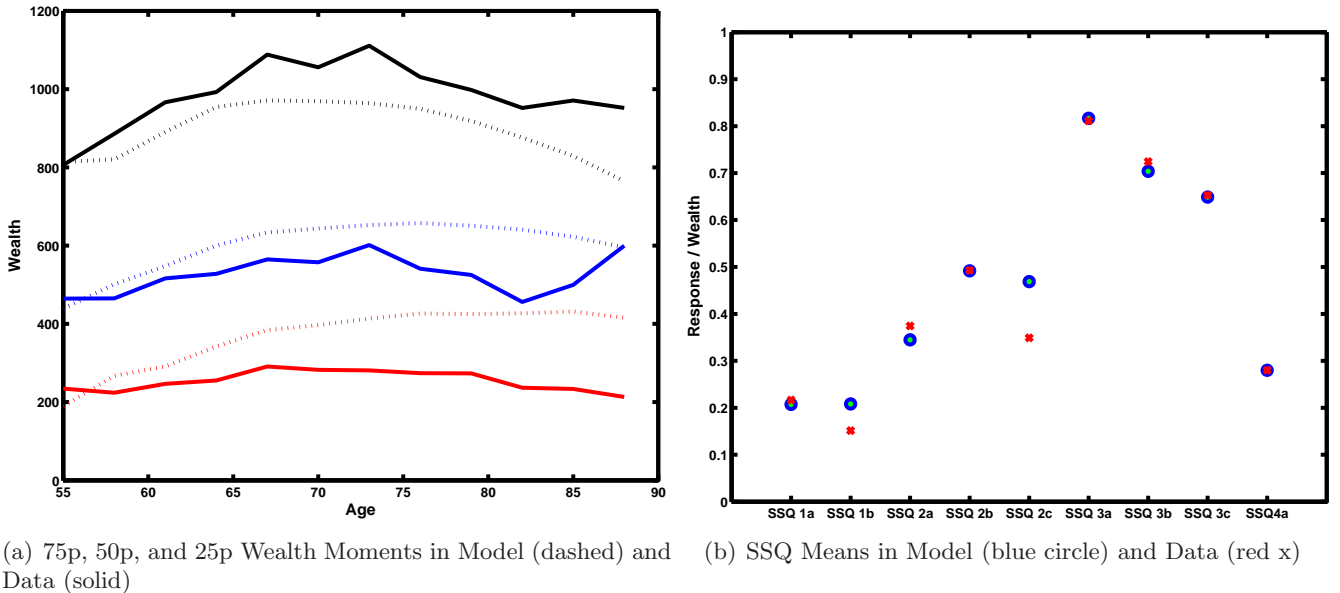


Figure 13: Model Fit When Exclusively Targeting SSQ Moments

Recall, Figures 8(a) and 8(b) document the model fit from the joint baseline estimation. Overall, the fit to both wealth and SSQ moments was very good. It is difficult, however, to both match the SSQ moments and have the model generate large enough savings at the top of the wealth distribution. This finding suggests that there may be motives missing from the model that predominantly affect the very affluent, such as inter-vivos family transfers and expenditure on luxury consumption goods. The tension in matching both sets of moments illustrated visually in Figures 8(a) and 8(b) is supported statistically by the over-identification

test statistic: The model’s ability to satisfy the joint moment conditions is rejected at the 1% level.²²

The statistical rejection of the joint estimation is not altogether surprising, given the distinct characteristics of the moments we target. The standard procedure in the literature is to treat the wealth-moment estimation as the baseline, with the difference of 46.58 between this case’s over-identification statistic and that of the joint estimate leading us to reject the inclusion of the SSQ moments at a 1% level. This result is not unique to this study, as other models of late in life saving that match the wealth distribution would have similar trouble in matching the survey responses we collected while maintaining such a tight fit to the wealth moments. One could similarly treat the SSQ-only estimation as baseline, resulting in a similar rejection of inclusion of wealth moments. We believe both sets of moments contain important information regarding saving motives, and thus choose to treat the baseline estimate as that which best matches both.

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 accurately represented. 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 data sets simultaneously highlights both the need for further model development (as we have started here by introducing the health-state utility function) and that the information contained in the SSQs provides a source of identification for richer models.

6.2 Estimated Parameter Values from Alternative Moments

Table 5: Estimated Parameters: Alternative Estimates

Joint Estimation: Baseline Model								
σ	θ_{LTC}	κ_{LTC}	θ_{beq}	κ_{beq}	ψ_G	ω_G	β	J -stat
5.85	1.57	-45.65	0.59	7.88	85.11	29.45	0.97	76.76
(0.06)	(0.13)	(0.01)	(0.71)	(0.25)	(1.74)	(36.67)	–	$\chi^2(39)$
Wealth Moments Only								
σ	θ_{LTC}	κ_{LTC}	θ_{beq}	κ_{beq}	ψ_G	ω_G	β	J -stat
7.80	17.43	-53.04	0.50	16.52	79.63	48.25	0.97	27.44
(6.15)	(33.72)	(10.91)	(2.80)	(12.27)	(9.98)	(2175.9)	–	$\chi^2(30)$
Strategic Survey Questions Only								
σ	θ_{LTC}	κ_{LTC}	θ_{beq}	κ_{beq}	ψ_G	ω_G	β	J -stat
3.56	0.13	-63.97	0.18	0.34	91.682	–	–	0.43
(0.04)	(0.01)	(1.15)	(0.02)	(1.44)	(2.24)	–	–	$\chi^2(3)$

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. The distribution of the J -stat is chi-squared, with degrees of freedom presented in parentheses.

Table 5 documents the parameter estimates that result from exclusively targeting wealth or SSQ mo-

²²This test statistic is calculated using the parameters estimated from the baseline weighting matrix to evaluate the objective function with the efficient weighting matrix. Thus, by construction, it is larger than the test statistic associated with the parameter set from the optimal weighting matrix estimation, but by a very small amount (76.76 vs. 74.02).

ments. 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).

In the estimates targeting wealth moments only, the model is able to better match the 75th wealth percentile by estimating both a higher risk aversion coefficient and a higher marginal utility when in need of long-term care. In the wealth estimation we find that most parameters are estimated with 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.

When matching the SSQ moments alone, the estimation procedure delivers a lower σ , θ_{LTC} , and θ_{beq} . It is difficult to compare the economic interpretation of estimated parameters in isolation because the value of σ affects the interpretation of the θ and κ parameters, just as the value of κ_{LTC} affects the interpretation of the value of θ_{LTC} . As a result, we present the following exercise, which illustrates the relative strength of the different expenditure motives induced by the different parameter values.

Comparing Alternative Parameter Sets: An Illustrative Static Choice Problem. In essence, our strategy to compare across sets of parameters is to map parameter values to expenditure behavior. To do so, we solve for optimal allocations in a simple synthetic problem, in which an individual regards regular consumption, LTC, and bequests as three different goods that can simultaneously be purchased and are valued according to the corresponding estimated utility functions. This problem is not one that maps to a situation ever faced by individuals. It is nonetheless a convenient illustration device to present the marginal utility of expenditures by expenditure type associated with each estimated parameter set. Specifically, consider the synthetic problem:

$$\begin{aligned} \max_{x_1, x_2, x_3} \quad & \frac{(x_1)^{1-\sigma}}{1-\sigma} + \frac{\theta_{LTC}(x_2 + \kappa_{LTC})^{1-\sigma}}{1-\sigma} + \frac{\theta_{beq}(x_3 + \kappa_{beq})^{1-\sigma}}{1-\sigma} \\ \text{s.t.} \quad & x_1 + x_2 + x_3 \leq W \\ & x_1, x_2, x_3 \geq 0; \quad x_2 \geq -\kappa_{LTC}; \quad x_3 \geq -\kappa_{beq}. \end{aligned} \tag{13}$$

Figures 14(a), 14(b), and 14(c) plot the resulting optimal allocations for the parameter sets that result from targeting different moments. The most striking feature of these figures is that, across all estimates, LTC expenditure is allocated the majority of wealth up to rather high wealth levels. Because LTC expenditure is a necessary good at least $-\kappa_{LTC}$ is always allocated to this motive, giving rise to an initial allocation share of 100% 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.

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




Figure 14: Expenditure Fractions Across Parameter Estimates

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 14(b) by the slow decline in LTC expenditure as wealth increases. The SSQ estimates indicate very low θ_{LTC} and θ_{beq} . The estimates imply a low marginal value of wealth in these states, and is reflected by the steeper allocation profile of ordinary consumption in panel 14(c). Finally, note that the lines in panel 14(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.

Summary of Estimated Preferences. 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_{LTC} is always estimated to be negative and large in magnitude. Often θ_{LTC} is larger than θ_{beq} , and in the SSQ-moment case when they are almost equal, the large negative κ_{LTC} 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 an 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.

7 Comparison of Preferences Across the Literature

Different preference parameters can induce radically different saving and spending behavior. Unfortunately, there is no consensus in the literature on the values for these parameters, as there are substantial differences in parameter estimates across leading papers. 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 15(a) and 15(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 (2014), as these are some of the most closely related papers to ours. We translate the parameters from other papers

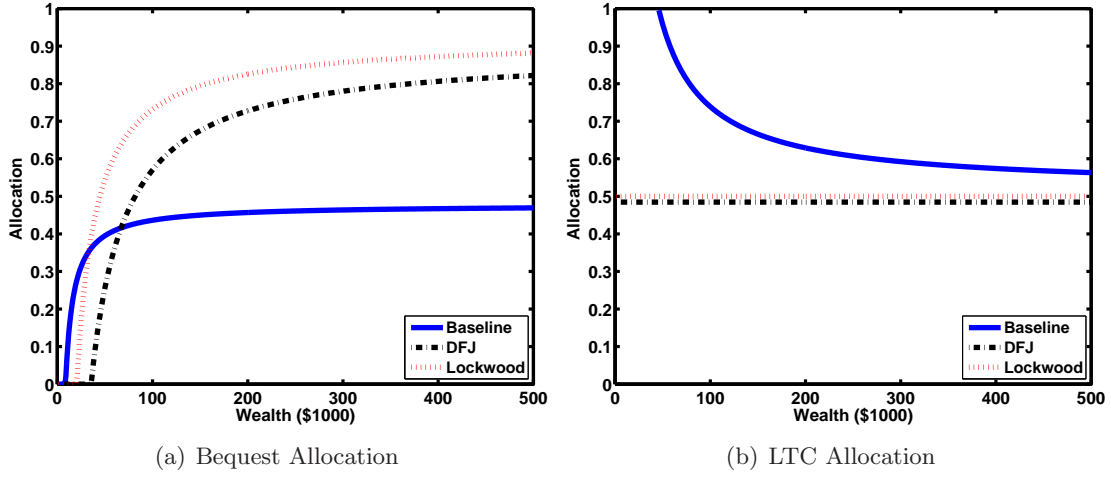


Figure 15: Expenditure Levels Across Studies

to be compatible with our utility function specification, yielding DFJ parameters of $\sigma = 3.81$, $\theta_{LTC} = 0.79$, $\kappa_{LTC} = 0$, $\theta_{beq} = 2360$, and $\kappa_{beq} = 273$ and Lockwood parameters of $\sigma = 3$, $\theta_{LTC} = 1$, $\kappa_{LTC} = 0$, $\theta_{beq} = 1460$, and $\kappa_{beq} = 231$. In a similar analysis to that presented in Figure 14, in Figure 15(a) we present the optimal bequest allocation (x_2) from the following optimization problem when calibrated according to each paper's baseline estimate of risk aversion and bequest parameters:

$$\begin{aligned} \max_{x_1, x_2} \quad & \frac{(x_1)^{1-\sigma}}{1-\sigma} + \frac{\theta_{beq}(x_2 + \kappa_{beq})^{1-\sigma}}{1-\sigma} \\ \text{s.t.} \quad & x_1 + x_2 \leq W \\ & x_1, x_2 \geq 0; \quad x_2 \geq -\kappa_{beq}. \end{aligned} \tag{14}$$

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 De Nardi, French, and Jones (2010) and 92% for Lockwood (2014). 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. An interpretation of De Nardi, French, and Jones (2010) is that bequests are found to play a small role in late in life savings patterns, while Lockwood (2014) seems to suggest that highly valued bequests can explain much of late in life wealth holdings. This difference between the estimated parameters of the two studies 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 lower bequest motives is that our modeling of the LTC state reduces the need for such a strong bequest motive to match wealth patterns. In Figure 15(b) we present the x_2 allocation for the problem presented in equation 14, with the bequest function replaced by the LTC state utility function. Lockwood (2014) does not allow for state dependent utility, and hence would suggest an equal allocation across states. De Nardi, French, and Jones (2010) does allow for a differential marginal utility multiplier when unhealthy, but no marginal utility 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 16: Model Fit to Wealth Moments for Alternative Parameters

Figure 16 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 Modigliani environment (see Modigliani and Brumberg (1954)), 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 ($\sigma = 3$, $\theta_{LTC} = 1$, $\kappa_{LTC} = 0$, $\theta_{beq} = 0$, $\kappa_{beq} \approx \infty$, $\psi_G = 0$, $\omega_G = 0$, 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 parameters

look most similar to those of Lockwood (2014), with DFJ and especially Modigliani 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.

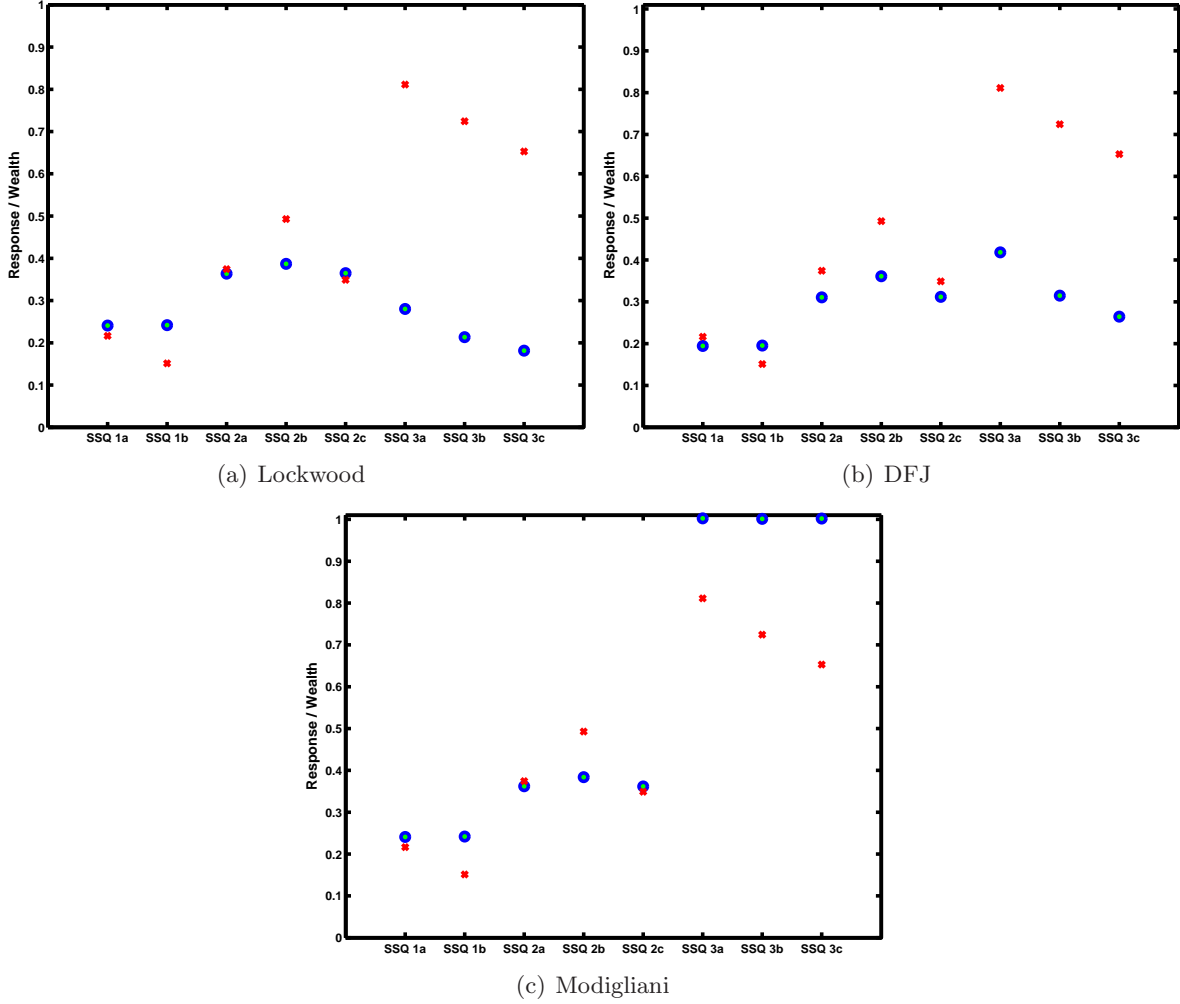


Figure 17: SSQ Means in Model (blue circle) and Data (red x) for Alternative Parameters

Figure 17 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 private LTC and leaving a bequest. The parameters estimated by Lockwood and DFJ are completely at odds with the SSQ responses, as they imply a much too high propensity to allocate towards bequests given their estimated θ_{beq} and κ_{beq} . In contrast, since the Modigliani 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 16 and 17 show that models can be consistent with the preference information provided by SSQs without sacrificing fit with the empirical wealth-age distribution. Furthermore, SSQ

²³Many papers in the literature, including De Nardi, French, and Jones (2010), 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.

methods can be used as a powerful parameter identification device.

Given that these parameters represent such different preferences and induce very different saving behavior in different contingencies, SSQs provide extra information that help 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.

8 Conclusion

In this paper, 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 can choose the amount to spend when in need of LTC on the intensive margin. We develop strategic survey questions (SSQs) that identify preference parameters using a novel application of stated-preference methodology. The model is estimated using data from the newly created Vanguard Research Initiative, using moments from the wealth distribution alone, SSQs alone, and both wealth and SSQs. A robust finding is that 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 patterns—tending to increase savings both across the wealth distribution and over the life cycle—for a large fraction of the U.S. population.

References

- AMERIKS, J., J. BRIGGS, A. CAPLIN, M. D. SHAPIRO, AND C. TONETTI (2015): “Long Term Care Insurance, Annuities, and the Under-Insurance Puzzle,” *Vanguard Research Initiative Working Paper*.
- 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,” *Scandinavian 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): “Using 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 (2013): “Old, Sick, Alone and Poor: A Welfare Analysis of Old-Age Social Insurance Programs,” *working paper*.
- BROWN, J. (2007): “Rational and Behavioral Perspectives on the Role of Annuities in Retirement Planning,” *NBER Working Papers*, 13537, 1–32.
- 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*, pp. 1083–1102.
- BROWN, J. R., G. S. GODA, AND K. MCGARRY (2013): “State-Dependent Utility and Insurance Purchase Decisions,” *working paper*.
- DE NARDI, M. (2004): “Wealth Inequality and Intergenerational Links,” *Review of Economic Studies*, 7, 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.
- DYNAN, K., J. SKINNER, AND S. ZELDES (2002): “The importance 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 (2013): “Genworth 2013 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 (2013): “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), 360–399.
- 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, 1028–1038.
- KOIJEN, R. S. J., S. VAN NIEUWERBURGH, AND M. YOGO (2015): “Health and Mortality Delta: Assessing the Welfare Cost of Household Insurance Choice,” *Journal of Finance*, forthcoming.
- KOIJEN, R. S. J., AND M. YOGO (2015): “The Cost of Financial Frictions for Life Insurers,” *American Economic Review*, forthcoming.
- 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, 29–72.
- KOTLIKOFF, L. J. (1988): “Intergenerational Transfers and Savings,” *Journal of Economic Perspectives*, 2(2), 41–58.
- 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,” Discussion paper, National Bureau of Economic Research.
- LAITNER, J., D. SILVERMAN, AND D. STOLYAROV (2014): “Annuitized Wealth and Post-Retirement Saving,” *working paper*.
- 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. (2014): “Bequest Motives and the Choice to Self-Insure Late-Life Risks,” *working paper*.
- 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.

- MODIGLIANI, F., AND R. H. BRUMBERG (1954): “Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data,” in *Post-Keynesian Economics*, ed. by K. K. Kurihara, pp. 388–436. Rutgers University Press.
- 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.
- 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 KLAUW, W. (2012): “On the Use of Expectations Data in Estimating Structural Dynamic Choice Models,” *Journal of Labor Economics*, 30(3), 521–554.
- VAN DER KLAUW, 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.

Appendix

A External Estimates

A.1 Income




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 10th, 30th, 50th, 70th, and 90th 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,

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

We then estimate β_k using a maximum likelihood estimator, and use these estimates to construct the corresponding cells in the health transition matrices.

Figures A.2 and A.3 display the estimated health state transition probabilities ($\pi_g(s'|t, s)$). Section 4.1.2 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 an 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 log out-of-pocket medical expenditures on age, gender, health state, and interaction terms. 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, $\sigma^2(t, g, s)$. Discretizing the error term $\epsilon_t \sim N(0, 1)$ into separate health cost states determines the medical expense process.

Figure A.4 plots the mandatory health costs spent over the life cycle by men 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. Overall, out of pocket health costs are much smaller than LTC expenditures and thus contribute little to the overall precautionary savings motive.

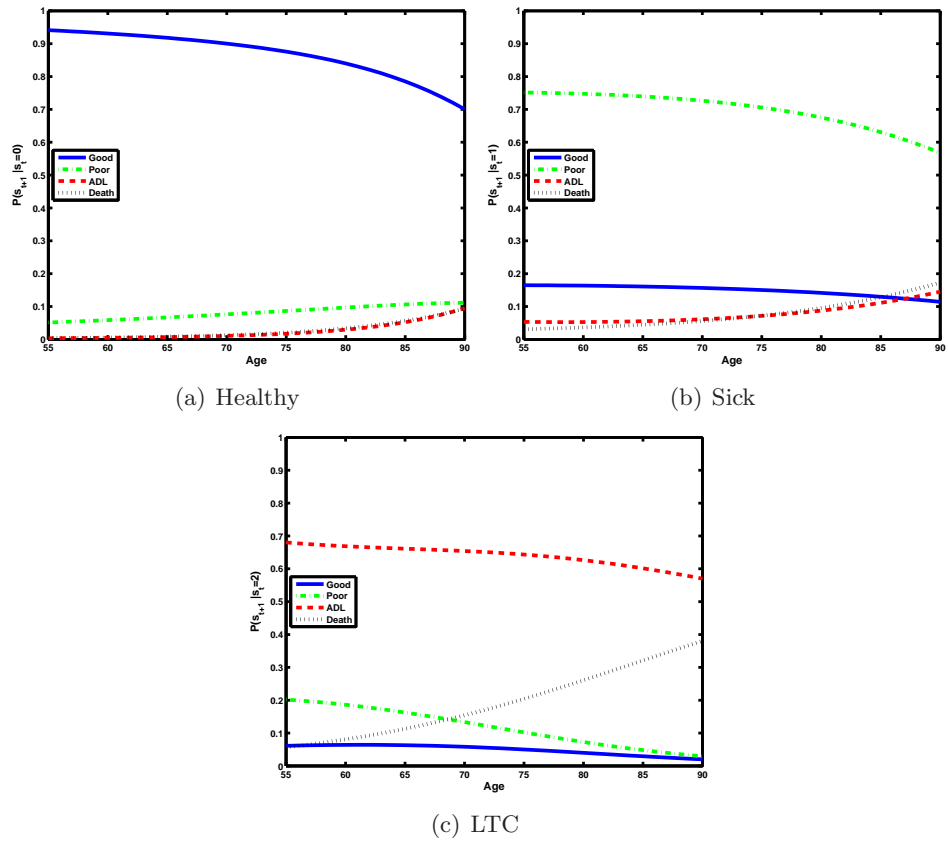


Figure A.2: Male Health State Transition Profile

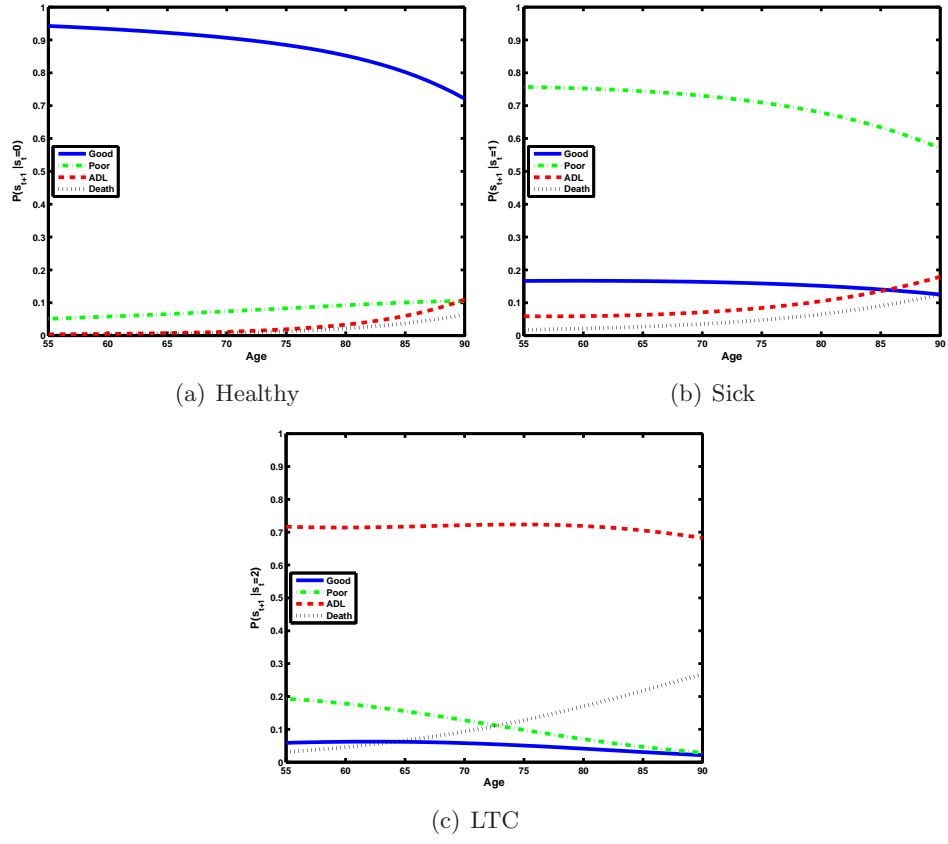


Figure A.3: Female Health State Transition Profile




Figure A.4: Median Health Cost Profile

B Estimation

As is standard, asymptotically,

$$\begin{aligned}\sqrt{N}(\hat{\Theta} - \Theta_0) &\rightarrow \mathbb{N}(0, \Psi) \\ \Psi &= \left(\frac{N+I}{N}\right) (D'WD)^{-1} (D'W\Omega WD) (D'WD)^{-1} \\ D &= \left. \frac{\partial g(\hat{\Xi}, \Theta, X)}{\partial \Theta'} \right|_{\Theta=\Theta_0},\end{aligned}$$

with Ω defined as the empirical covariance matrix. This result holds generally for every estimation exercise we conduct in this paper.

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. Due to computational constraints, in place of Monte Carlo methods we use numerical derivatives to obtain standard errors. Specifically, the gradient of the moment conditions is calculated numerically. Standard errors are then reported as the square root of the diagonal elements of the asymptotically normal covariance matrix.

We then denote the implied first stage covariance matrix as:

$$\begin{aligned}\hat{\Omega}_1 &= \mathbb{E} \left[g(\hat{\chi}, \hat{\Theta}_1, X)' g(\hat{\chi}, \hat{\Theta}_1, X) \right] \\ &= \frac{1}{N} \sum_{n=1}^N g(\hat{\chi}, \hat{\Theta}_1, X_n)' g(\hat{\chi}, \hat{\Theta}_1, X_n).\end{aligned}$$

C Insurance Product Demand and Prices

Insurance products are priced conditional on age, health status, and gender. The price of an insurance product is denoted by $p(t, s, g)$, such that spending \$ $\tilde{y} \times p(t, s, g)$ purchase payout \tilde{y} per year when the insurance-relevant states are realized. In the case of annuities, income is paid every year, and in the case of ADLI it is paid only in years when $s = 2$.

Taking prices as given, demand for insurance is calculated as

$$\begin{aligned}D(a, y, t, s, h, g) &= \arg \max_{\tilde{y}} V(a - p(t, s, g)\tilde{y}, \hat{y}, t, s, h, g) \\ \hat{y} &= y(t) + \tilde{y},\end{aligned}\tag{C.1}$$

where \hat{y} is the income stream including insurance payouts and V is the value function evaluated at the new wealth level and income stream. Note that without insurance, income is deterministic and only age-dependent. Purchasing insurance makes the new income stream (\hat{y}) stochastic through its dependence on age and health. To calculate $D(a, y, t, s, h, g)$ consumer optimal policies are computed over a grid of \tilde{y} and the \tilde{y} that maximizes the value function is obtained by interpolation.

Prices. Prices are first calculated to be actuarially fair conditional on age, health, and gender. Actuarially fair is defined to be the price such that the agency selling the product makes zero profit in expected value.

The realized period payouts for annuities and ADL insurance depend on health state s . An annuity pays out while $s = 0, 1$ or 2 , while ADLI pays out while only when $s = 2$. Thus, the vector of period payouts across health states $s \in \{0, 1, 2, 3\}$ for annuities is

$$\vec{y} = [\tilde{y}, \tilde{y}, \tilde{y}, 0]',$$

while for ADLI it is,

$$\vec{y} = [0, 0, \tilde{y}, 0]'$$

Let \vec{s} be an indicator vector that has elements s_i for $i \in \{0, 1, 2, 3\}$ equal to zero for $s \neq i$ and equal to 1 if $s = i$. The insurance product is priced to equal the expected discounted stream of payments. Thus, an insurance product that pays out \vec{y} per period for a person of age t , gender g , with current health status s has price

$$p(t, s, g) = \vec{s} \times \left[\sum_{i=0}^{T-t} \frac{1}{(1+r)^i} \prod_{k=0}^i \pi_g(s'|t+k, s) \right] \times \vec{y}. \quad (\text{C.2})$$


In practice, we price the insurance income in units of $\tilde{y} = \$10,000$ per year. In Section 5.3.2 we present demand for annuities and ADLI for a 65 year old male in good health. The corresponding lump-sum prices before the 10 percent load are \$131,578 for annuities and \$8,503 for ADLI.

D Estimation Results Using the Optimal Weighting Matrix

Table D.1: Estimated Parameters: Alternative Weighting Matrices

Joint Estimation: Baseline Model								
σ	θ_{LTC}	κ_{LTC}	θ_{beq}	κ_{beq}	ψ_G	ω_G	β	J-stat
5.85	1.57	-45.65	0.59	7.88	85.11	29.45	0.97	76.76
(0.06)	(0.13)	(0.01)	(0.71)	(0.25)	(1.74)	(36.67)	—	$\chi^2(39)$
Joint Estimation: Optimal Weighting Matrix								
σ	θ_{LTC}	κ_{LTC}	θ_{beq}	κ_{beq}	ψ_G	ω_G	β	J-stat
4.45	0.41	-44.49	0.11	3.79	83.95	4.59	0.97	74.02
(0.001)	(0.001)	(0.015)	(0.001)	(0.03)	(0.71)	(2111.7)	—	$\chi^2(39)$

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. The distribution of the J -stat is chi-squared, with degrees of freedom presented in parentheses.

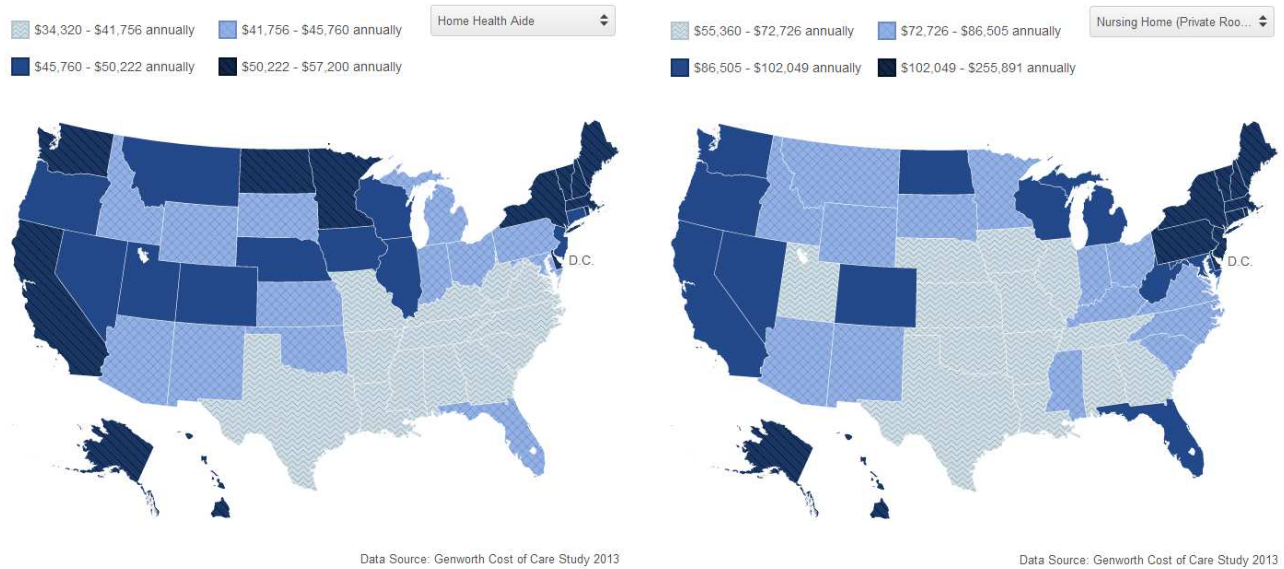


(a) 75p, 50p, and 25p Wealth Moments in Model (dashed) and Data (solid)

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

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

E LTC Expenditure



(a) Home Health Aide

(b) Nursing Home Private Room

Figure E.1: Heterogeneous Expenditure on LTC