The Long-Term-Care Insurance Puzzle: Modeling and Measurement*

John Ameriks
The Vanguard Group, Inc.

Joseph Briggs
Federal Reserve Board of Governors

Andrew Caplin
New York University and NBER

Matthew D. Shapiro
University of Michigan and NBER

Christopher Tonetti
Stanford GSB and NBER

June 2018

Abstract

Individuals face significant late-in-life risks, prominently including the need for long-term care (LTC). Yet, they hold little long-term care insurance (LTCI). In this paper we use a structural model and a purpose-designed dataset to understand the determinants of insurance demand. We distinguish between a fundamental lack of desire to insure, crowd out from existing insurance, and unmet demand due to poor products available in the market. The model features individual-specific non-homothetic health-state-dependent preferences over normal consumption, consumption when in need of long-term care, and bequests, which are estimated using strategic survey questions. To account for differences between the modeled and measured insurance products, we study not only individuals' holdings of LTCI, but also their stated demand for an idealized product that mirrors that in the model. We find that many individuals would purchase LTCI and receive a large consumer surplus if it were a better product, while many others do not want to purchase even high-quality actuarially fair LTCI due to the values of their heterogeneous state-dependent preferences, their demographics, and their financial situation.

JEL classification: D14, E21, G22, I13, J14

*An earlier draft of this paper circulated under the title "Late-in-Life Risks and the Under-Insurance Puzzle." This research is supported by a program project grant from the National Institute on 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 Survey and provided substantial input into its design. The design of the VRI benefited from the collaboration and assistance of Wandi Bruine de Bruin, Alycia Chin, Brooke Helppie McFall, Minjoon Lee, Mi Luo, and Ann Rodgers, as part of the program project, from Annette Bonner of the Vanguard Group, Inc., and from the research staff at IPSOS SA. We thank Norma Coe, Mariacristina De Nardi, Rebecca Diamond, Sebastian Di Tella, Eric French, Amanda Starc, Víctor Ríos-Rull, and Moto Yogo for useful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of The Vanguard Group, Inc. or the Federal Reserve Board of Governors. For documentation of the VRI, including a dynamic link to the survey instrument, see `http://ebp-projects.isr.umich.edu/VRI/`.

1 Introduction

Long-term care is expensive and the need for it pervasive. One in three 65 year old Americans will eventually enter a care facility, with high-quality care potentially costing \$100,000 or more. Put starkly, there is about a 1 in 6 chance of needing at least three years of long-term care (LTC). In this case, the resulting \$300,000 needed to self-insure would be larger than the financial wealth of three out of four older American households. Hence, it is striking that only a small fraction of elderly Americans hold long-term care insurance (LTCI) and that these policies account for only 4 percent of aggregate LTC expenditure.[1]

Why is purchase of LTCI so low? Is there room for improved LTCI that could substantially improve welfare? It could be that low ownership rates reflect a fundamental lack of desire to insure against this health realization. Alternatively, there might be demand for insuring this health risk that is already being met, e.g., via public provision of care that could be crowding out private insurance. There could be, however, substantial unmet demand, with low LTCI ownership reflecting poor quality insurance products available in the market. Understanding the source of the observed low LTCI holdings is critical to determining the value of potential changes—via government policies or private sector products—to insuring late-in-life risks. In our sample of older Americans with enough wealth to potentially self-finance LTC, only 22 percent of individuals own LTCI. Our main finding is that many individuals would purchase LTCI and receive a large consumer surplus if it were a better product, while many others do not want to purchase even high-quality actuarially fair LTCI due to the values of their heterogeneous state-dependent preferences, their demographics, and their financial situation.

To quantify the factors that generate measured LTCI ownership, we estimate demand for insurance against needing long-term care using a structural life-cycle model with stochastic health and death and incomplete markets. The model features individual-specific non-homothetic health-state-dependent preferences over consumption when healthy, consumption when in need of long-term care, and bequests. A main contribution of the paper is that we discipline the model by using purpose-designed survey data to estimate these preferences. Estimation does not use insurance ownership data, so the heterogeneous preferences we introduce are not free parameters.

We compare modeled demand to measured ownership of LTCI and define the *LTCI puzzle* as the difference. Part of the puzzle could be due to the fact that insurance in the model and insurance products available in the market are not the same. Therefore, we collect a measure of stated demand for an insurance product that mirrors that in the model.

We model insurance as a state-contingent asset that pays when LTC is needed, which we call Activities of Daily Living insurance (ADLI). Our model predicts 59 percent ADLI ownership. When offered the opportunity to buy the idealized ADLI product in the survey, 46 percent of the population reveals positive stated-demand. As a point of reference, a homogeneous preference model predicts 78 percent ownership. Thus, incorporating heterogeneous preferences and comparing similar products dramatically shrinks the LTCI puzzle from 56 (i.e., 78-22) to 13 (i.e., 59-46) percentage points. We shrink the puzzle from above and below by lowering modeled demand and raising measured demand: Modeling individuals who are heterogeneous in financial situations, demographics, and preferences together with incomplete markets substantially reduces modeled demand relative to a complete insurance benchmark; using stated demand for the same idealized insurance product as in the model substantially raises measured demand. Furthermore, we document that consumer surplus is large for many individuals, suggesting significant welfare gains from improved insurance against needing LTC.

[1] See Brown and Finkelstein (2008) for the likelihood of needing care, Brown and Finkelstein (2011) for LTCI ownership and aggregate expenditures, and Ameriks, Caplin, Lee, Shapiro, and Tonetti (2014) for wealth statistics. Genworth (2016) calculates \$92,378 as the average cost in the U.S. for one year in a private nursing home room.

The remaining 15 percentage point difference in ADLI demand represents model misspecification or measurement error unrelated per se to the difference between LTCI and the better ADLI product. To better understand the remaining puzzle, we perform a statistical analysis of the gap between stated and modeled ADLI, providing guidance for particular features that might improve model fit in the future (e.g., issues related to family). Of course, our analysis leaves open the question of why desired products are not offered in the market to meet demand, but we view this paper as removing *lack of demand* for good insurance as a reason for the *lack of supply*.[2]

We conduct various exercises that support our main findings. We show robustness of the results to various changes in model parameters and subsamples of the population. We compute the intensive margin of demand, showing that those who buy ADLI purchase a sizable amount and that consumer surplus is large. We analyze demand and ownership patterns in the cross-section, showing that the LTCI puzzle is spread across the wealth distribution, while the gap between modeled and stated ADLI demand is concentrated in the high wealth quintiles. Overall, these exercises support our main finding that there exists significant unmet demand for improved LTCI, even though not everyone would buy it.

The remainder of the paper proceeds as follows. The introduction concludes with a literature review. Section 2 provides an overview of the long-term care insurance market. Section 3 presents the model, discuses the motives that generate ADLI demand, and details the calculation of modeled ADLI demand (given individual financial, demographic, and health characteristics and preferences). Section 4 introduces the VRI and the key financial, demographic, and health data that are the state variables in our model. Section 5 discusses strategic survey questions, including parameter identification, the estimation strategy, the resulting individual-specific preference parameter estimates, and an economic interpretation of the estimated preferences. Section 6 presents the model-based estimates of ADLI demand, including information on the cross-section of ADLI demand and the ADLI demand function. Section 7 presents stated ADLI demand. Section 8 concludes.

1.1 Relation to the Literature

Long-Term Care and Insurance. As noted in Brown and Finkelstein (2011), the need for long-term care is one of the largest uninsured risks facing the elderly and understanding the reasons for non-insurance of this risk is a first-order issue in improving household welfare and the economic and health security of elderly Americans. Research suggests that LTCI may be a difficult product for insurers to offer. Furthermore, research suggests that individuals want to insure against needing LTC, but also that they may not want to purchase private LTCI due to crowd out from publicly provided or informal insurance or because of particular features of LTCI available in the market. In contrast to previous research, in this paper we focus on older Americans who have the finances to potentially self-fund their LTC needs and quantitatively estimate their demand and associated consumer surplus for insuring the state of the world in which they need LTC, independent from their desire to purchase LTCI products currently available in the market and accounting for the option to use Medicaid.

It is well understood that there could be supply-side limitations on the provision of LTCI. Cutler (1996) discusses the difficulties of insuring inter-temporal risk. Finkelstein and McGarry (2006), Brown and Finkelstein (2007), and Hendren (2013) document evidence of adverse selection. Koijen and Yogo (2015) and Koijen and Yogo (2016) show that financial frictions and statutory regulations affect the profitability of insurance companies more generally.

There may also be significant demand-side reasons explaining the low holdings of LTCI, including crowding

[2] Our paper is a demand-side compliment to the literature on the supply-side of the insurance market. The supply-side literature points to aggregate risk, adverse selection, and government crowd-out as factors that limit the supply of insurance. See, e.g., Cutler (1996), Hendren (2013), Koijen and Yogo (2015), and Braun, Kopecky, and Koreshkova (2017a).

out from government provided care (Pauly (1990), Brown and Finkelstein (2008)) with means tested programs having effects on both low wealth and affluent households (De Nardi, French, and Jones (2016), Braun, Kopecky, and Koreshkova (2017b)). Medicaid is likely a very important determinant of LTCI purchase in America. Medicaid is means-tested, which imposes a high implicit tax on self-insurance via saving and it is a secondary payer, which generates a high implicit tax on LTCI (as discussed in Braun, Kopecky, and Koreshkova (2017a)). In this paper, we model Medicaid as a means-tested program that pays for LTC and we use SSQs to measure the subjective value individuals place on care received by Medicaid when in need of help with ADLs (similar to Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011)). Hackmann (2017) documents that the low reimbursement rates of Medicaid actually contribute to lower quality and less attentive care in nursing homes, resulting in worse health outcomes. Thus, there is ample reason to believe that people may have a desire to purchase high quality convenient care if they can afford it.

In the model, Medicaid generates a discontinuity in the savings policy of individuals, with those in a region of the state space mostly associated with low wealth choosing not to self-insure or buy LTCI, while those with more wealth save and buy LTCI with the intention to purchase private LTC when needed. Since our sample is wealthier than one representative of the U.S., we have more people in the wealthy region who buy LTCI, and thus Medicaid plays less of a role in determining LTCI than it likely does for the U.S. in general. We view our sample as representing those who are most likely to potentially self-finance their LTC needs, and thus a valuable sample to study the potential demand for improved insurance.

Our focus on household insurance demand aligns us closely, in method and purpose, to Hong and Rìos-Rull (2007), Inkmann, Lopes, and Michaelides (2011), Hong and Rìos-Rull (2012), Lockwood (2012), Lockwood (2018), and Koijen, Van Nieuwerburgh, and Yogo (2016) who all use life-cycle models with a rich specification of preferences to estimate demand for insurance products. Since insurance is an asset that promises state-contingent payouts, data on insurance ownership can be particularly powerful for identifying state-contingent valuations, especially when the fungibility of liquid financial wealth hampers identification.

Matching measured and model-implied ownership moments is best done when insurance products sold in the market closely resemble the corresponding state-contingent model objects. For LTCI in particular, as documented in Section 2, the product does not closely resemble a simple state-contingent claim that pays without risk or effort to be used as desired in the perfectly-verifiable ADL state, as is often modeled. Thus, there is an issue of trying to match empirical moments for one asset to model-implied moments for a fundamentally different asset. This potential gap between the insurance product in the market and in the model motivates a main difference between our approach and that typically used in the literature. First, we estimate preferences using SSQs instead of insurance data. Second, we obtain a model-free measure of stated demand for the same product as modeled. It is feasible to generate zero LTCI puzzle by choosing preferences to match insurance ownership. By estimating preferences without insurance data, we allow for a puzzle to exist, and use other data to discipline its size and characteristics.

Health-State-Dependent Utility. Although health-state dependent utility is not a new concept—around since at least Arrow (1974)—this feature is increasingly being incorporated into quantitative evaluations of household decision-making. Estimates vary on whether poor health increases or decreases the marginal utility of consumption (see Finkelstein, Luttmer, and Notowidigdo (2009) for an overview). Even so, there is a limit to the applicability of previous measures that use a general poor-health state to our study, since estimates may be highly contextual and LTC is a distinct health state that occurs at older ages and is associated with specific care needs, maladies, behaviors, and desires. Similar to our findings, work by Hong, Pijoan-Mas, and Rios-Rull (2015) uses panel data and Euler equations to estimate that lower health gives higher marginal utility at older ages. Most closely related to our approach is Brown, Goda, and McGarry (2016), who use a related survey methodology to document the degree to which there exists health-state dependent utility and find evidence of state dependence and significant heterogeneity in preferences. As models have developed to have richer heterogeneity in demographic and financial states, there are still many features of the data that are difficult to explain. Preference heterogeneity is a natural candidate for unobservable heterogeneity that drives behavior. SSQs provide independent information to estimate preference parameters so that they are not just free parameters used without discipline to match any puzzling behavior.

Life-Cycle Models and Saving Motives. The determinants of LTCI demand are similar to the forces driving late-in-life saving behavior. Thus, our work is closely related to the literature that uses life-cycle models to study the dynamics of savings in old age. Many recent models that explain the observed slow spend down of wealth in later life allow for both bequest motives and precautionary motives associated with high late-in-life health and long-term care (LTC) expenses. Laitner, Silverman, and Stolyarov (2015) and Barczyk and Kredler (2015) provide analytically tractable models that cleanly highlight the impact of different motives on saving decisions. Late-in-life health risks induce precautionary saving much like income risk does for workers (e.g., Zeldes (1989), Carroll (1997)). Despite early work by Hubbard, Skinner, and Zeldes (1994) and Palumbo (1999) suggesting that health expenses contribute only slightly to late in life saving, more recent studies find such expenses to be of greater importance. For example, Gourinchas and Parker (2002) provide a decomposition that identifies the role of precautionary saving in wealth accumulation. Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011), Kopecky and Koreshkova (2014), and Lockwood (2018) all model LTC expenses explicitly and De Nardi, French, and Jones (2010) and Koijen, Van Nieuwerburgh, and Yogo (2016) allow for a health expense risk that includes LTC, with all finding that health expenses introduce a significant precautionary saving motive.

In previous work, Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2017) also examined long-term care risk. In contrast to the environment studied in this paper, that paper studies saving dynamics using a model with homogeneous preferences in which there is no insurance available to purchase, while this one estimates how preferences differ individual-by-individual and studies insurance demand.

2 The Long-term Care Insurance Market

This section provides an overview of the LTCI market as background for the model and analysis. A main point of this paper is that it is possible to both want to insure the ADL-health state and not want to buy existing LTCI products available in the market. Thus, data on the low demand for LTCI does not necessarily indicate low demand to insure needing help with long-term care. We summarize below some features of the private LTCI market, highlighting that the existing products are far from a full set of simple state-contingent assets available for purchase to all. We find in this paper that the mismatch between modeled and existing insurance products is a significant contributor to the LTCI puzzle.

First, private LTCI policies are expensive, both in the eyes of consumers and relative to their actuarially fair value. Brown, Goda, and McGarry (2012) survey consumers and find that the cost of LTCI was the most commonly given reason households decide not to purchase a policy, cited by 57 percent of people in open-ended responses and with 71 percent of people expressing concern about being able to afford premiums in the future. This perceived high cost has basis in reality: Brown and Finkelstein (2011) note that a typical LTCI policy has a load of 32 cents on the dollar, well above loads typical in other insurance markets. In addition, Brown and Finkelstein (2011) estimate that a "typical" policy purchased at age 65 and held until death would only cover about two-thirds of the total expected present discount value of LTC expenditures.

The high cost of available LTCI policies can not alone explain the small market size however. Brown and Finkelstein (2007) note that, as of that time, existing policies did not differentiate prices by sex, resulting in better than actuarially fair policies (average load of -6 cents per dollar) for females. Nevertheless, coverage is approximately the same for males and females, suggesting that other factors are also likely important in accounting for the small market size.

A number of other potential factors were raised in Rubin, Crowe, Fisher, Ghaznaw, McCoach, Narva, Schaulewicz, Sullivan, and White (2014). For example, while most policies are guaranteed renewable, LTCI policy holders are subject to the important risk of an increase in required premium rates to maintain continuing coverage. If they cannot pay higher rates, they can lose their coverage. Insurers cannot raise premiums on individual LTCI policies in isolation, but, subject to regulatory approval, they can increase (and in several well-publicized changes have increased) rates for groups or classes of policyholders to reflect, among other factors, errors in actuarial underwriting assumptions. Moreover, policy benefit triggers, especially for tax-qualified LTCI policies, can be restrictive. Stallard (2011) finds that about half of [the elderly] disabled population does not meet the eligibility requirements for tax qualified LTC insurance policies due to not satisfying either the Health Insurance Portability and Accountability Act's ADL trigger definitions or its cognitive impairment trigger. In addition, Rubin, Crowe, Fisher, Ghaznaw, McCoach, Narva, Schaulewicz, Sullivan, and White (2014) cite current coverage portability and non-forfeiture provisions as limiting policy-holder options. Furthermore, consumer perceptions of market features, real or perceived, are likely important. Brown and Finkelstein (2007) note that "limited consumer rationality—such as difficulty understanding low-probability high-loss events. . . —may play a role" in the small size of the market, while Brown, Goda, and McGarry (2012) find that LTC coverage is highly correlated with beliefs regarding counterparty risk.

Another potentially undesirable feature of available LTCI policies is mismatch between expenses households would like to insure and those covered. Typical policies provide for institutional care and home care with a maximum daily benefit (on average $153 in 2010) for a maximum benefit period of 1 to 5 years (Brown and Finkelstein (2007)). On one hand, restrictions on use of funds may discourage demand. For example, some individuals might prefer to have a family member provide care (Brown, Goda, and McGarry (2012)), an option that is not possible in many policies. Additionally, restrictions on the benefit period may discourage private insurance purchases. Most policies have a deductible of 30 to 100 days of out of pocket care before benefit payments can begin.[3] Longer stays that exceed the maximum benefit period, which could occur in cases of cognitive decline, dementia, and Alzheimer's disease, are not covered. Thus, most existing policies neither insure the most common nor most expensive stays in nursing homes.

Furthermore, the private LTCI market actually appears to be shrinking. Following substantial growth of the market during the 1980s and 1990s, between 2003 and 2010 individual policy sales declined by 9 percent per year and the number of firms selling "meaningful policies" decreased from 102 to approximately a dozen. This significant retraction was driven by decisions to stop issuing new policies, with exiting firms citing high capital requirements, poor profits, regulatory hurdles surrounding rate increases, and difficulty mitigating investment risk as reasons for exit (Cohen, Kaur, and Darnell (2013)). While private LTCI policies are still available for purchase, this rapid retraction in market size likely does not instill confidence in consumers.

A number of policies have in recent years attempted to expand the private LTCI market. A limited federal subsidy was offered beginning in 1997 and between 1996 and 2008 the number of states offering tax incentives for private LTCI purchase had increased from 3 to 24 (Goda (2011)). The Community Living Assistance Services and Supports

[3] Medicare pays for the first 20 days and subsidizes the next 80 days of a stay at a skilled nursing facility or home health care in certain instances.

(CLASS) Act created a publicly funded federal LTCI program designed to make LTCI available to individuals who private insurance companies would not underwrite, but this law was repealed in 2013. In addition, the National Association of Insurance Commissioners (NAIC) implemented LTCI Model Regulation to protect consumers from unexpected premium increases in 2000 and voted to require greater justification for proposed rate changes in 2014. Overall, these efforts appear to have been at best modestly effective in growing the private LTCI market.

In summary, from the consumer perspective LTCI may not be attractive because of high prices, an adversarial claims process with uncertainty around the ability to successfully claim, limited contract coverage options, and the risk of increased premiums. Many firms have reported that they do not find LTCI to be an attractive product to sell, referencing capital requirements, regulatory hurdles, and difficulty in hedging associated risks. Although it is beyond the scope of this paper to determine why the private LTCI market seems under-developed, adverse selection or public crowding-out are commonly cited reasons for the market failure (see, e.g., Cutler (1996), Hendren (2013), Koijen and Yogo (2015), and Braun, Kopecky, and Koreshkova (2017a)).

How to Proceed when LTCI in the Market is not the Same as Insurance in the Model. The LTCI market does not offer a full set of simple state-contingent assets available for purchase to all. Thus, the distinction between existing products and modeled products should be accounted for when inferring preference parameters from observed LTCI holdings. There are two possible approaches: model LTCI as it exists in the market or measure demand for the simple product that is in the model. There are pros and cons to each approach. Studying the idealized product provides fundamental information on the welfare from insuring the ADL-state, but does not explain what are the features of the actual product that limit demand. Modeling the existing LTCI products is complicated by the fact that LTCI contracts are multi-dimensional and we have limited measurement on many of the relevant dimensions. As noted in Brown, Goda, and McGarry (2012), "a policy intervention that addresses only one market limitation, such as pricing, without addressing other concerns, such as counterparty risk, is unlikely to increase demand dramatically." We therefore abstract from available LTCI products and study demand for ADLI, a type of LTCI that takes the form of a simple state contingent asset without the above noted product imperfections. Focusing on ADLI allows us to quantify the fundamental demand for insuring this health realization and the value of creating such an insurance product, abstracting from the supply-side barriers to its creation and complications associated with existing LTCI contracts.

3 The Model

This section presents the consumer choice model that will be used to predict demand for insurance products. The model is a heterogeneous-preference extension of that developed in Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2017), which studies saving and spending over the life cycle. The model is a modern incomplete market heterogeneous agent life-cycle consumption/saving problem with health and longevity risk, similar to that, e.g., in De Nardi, French, and Jones (2010) and Lockwood (2018). A key methodological contribution of the paper is identification and estimation of rich preferences at the individual level—including design and measurement of the necessary associated survey data. We embed these preferences within a model otherwise similar to those used in the recent literature.

3.1 The Individual Optimization Problem

The model considers individuals who are heterogeneous over wealth, income age-profile, age, sex, health status, and preferences. An individual's health status can be good health, poor health, needs help with the activities of daily living (ADLs), or dead. Needing help with ADLs is defined as needing significant help with activities such as eating, dressing, bathing, walking across a room, and getting in or out of bed, and is regarded as provoking need for long-term care. Health is risky and evolves according to a Markov process conditional on age, sex, and prior health status. Individuals start at age 55 and live to be at most 108 years old. We first develop a model of individuals that do not have the ability to purchase insurance, and then introduce private insurance into the model. Each period individuals choose consumption, savings, and whether to use government care. The model groups people into five income groups with deterministic age-income profiles. For tractability, we abstract from labor-income risk which, while essential to model for younger individuals, is less a determinant of behavior for older individuals out of the labor force or near retirement who have a larger ratio of financial to human wealth.[4] There is a risk free rate of return of $(1 + r)$ on savings. The risk free return is calibrated to a baseline 1 percent, although Section 6.3 shows that results are robust to allowing for a 3 percent rate. There is no borrowing and the retiree cannot leave a negative bequest.

Together $\Theta^i := \{\gamma^i, \theta^i_{ADL}, \kappa^i_{ADL}, \theta^i_{beq}, \kappa^i_{beq}, \psi^i_G\}$ define an individual's preferences over risk, expenditure in the ADL-state, and bequests. When in good or poor health, consumers value consumption according to standard CRRA preferences with parameter $\gamma^i > 0$:

$$\frac{c^{1-\gamma^i}}{1-\gamma^i}.$$

Utility associated with consumption level c when in need of help with ADLs is

$$\left(\theta^i_{ADL}\right)^{-\gamma^i} \frac{(c + \kappa^i_{ADL})^{1-\gamma^i}}{1-\gamma^i}.$$

Upon death, the individual receives no income and pays all mandatory health costs. Any remaining wealth is left as a bequest, b, which is valued with warm glow utility

$$\left(\theta^i_{beq}\right)^{-\gamma^i} \frac{(b + \kappa^i_{beq})^{1-\gamma^i}}{1-\gamma^i}.$$

The non-homothetic bequest utility function is a workhorse in quantitative life-cycle models, adopted from De Nardi (2004), that is used to capture differences in behavior as a function of wealth. Here we model the ADL-state utility symmetrically with bequest utility, also allowing it to be non-homothetic. The ADL-state and bequest utility functions are each governed by two key parameters: θ^i and κ^i. κ^i_s controls the degree to which the expenditure when in health state s is a luxury or a necessity by deviating from homotheticity. θ^i_s is the multiplier on the marginal utility of an additional dollar spent in health-state $s \in \{\text{ADL, Dead}\}$ relative to when in good health . Asymptotically, as wealth grows, κ^i has less of an influence on expenditure relative to θ^i. An increase in θ^i decreases the marginal utility of a unit of expenditure; an increase in κ^i indicates that expenditure is more of a luxury. Negative κ^i can be interpreted as the expenditure being a necessity.

The consumer has the option to use a means-tested government provided care program. The cost of using government care is that a consumer forfeits all wealth. If the consumer chooses to use government care when not in the ADL health state, the government provides a consumption floor, $c = \omega_G$. A person who needs help with ADLs has access to government-provided care that is loosely based on the institutions of Medicaid. If an individual needs help with ADLs and uses government care, the government provides $c = \psi^i_G$. The value ψ^i_G parameterizes the consumer's value of public care, since that parameter essentially determines the utility of an individual who needs help with ADLs

[4] The model abstracts from labor supply decisions, including retirement. These labor market decisions are taken into account through the exogenous income profiles. See Appendix A.2 for details.

and chooses to use government care. Additionally, to capture the fact that private LTC provision is a lumpy and costly expense, we model a minimum level of spending needed to obtain private LTC, i.e., $c \geq \chi_{ADL}$ when help with ADLs is needed and government care is not used.

Let wealth be $a \in [0, \infty)$, age be $t \in \{55, 56, ..., T = 108\}$, the income age-profile be $y \in \{y_1, y_2, \ldots, y_5\}$ with $y_k = \{y_k(t)\}_{t=55}^{T}$, sex be $g \in \{m, f\}$, health status be $s \in \{0, 1, 2, 3\}$ (0 = good health, 1= poor health, 2= needs help with ADLs, and 3 = death), the health state Markov transition matrix be $\pi_g(t, s)$, health cost be h, and $G \in \{0, 1\}$ be the government care indicator. The idiosyncratic state variables are $X^i := \{a^i, y^i, t^i, s^i, h^i, g^i\}$. Written recursively, the consumer problem is:

$$
\begin{aligned}
V^i(a, y, t, s, h, g) = \max_{a',\, c,\, G} \; & \mathbb{I}_{s \neq 3} \, (1 - G) \; \left\{ U_s^i(c) + \beta E[V^i(a', y, t+1, s', h')] \right\} \\
& + \mathbb{I}_{s \neq 3} \, G \; \left\{ U_s^i(\omega_G, \psi_G) + \beta E[V^i(0, y, t+1, s', h')] \right\} + \mathbb{I}_{s=3} \{ v^i(b) \}
\end{aligned}
\tag{1}
$$

s.t.

$$
\begin{aligned}
& a' = (1 - G)[(1 + r)a + y(t) - c - h] \geq 0 \\
& c \geq \chi_{ADL} \;\text{ if }\; (G = 0 \;\wedge\; s = 2) \\
& c = \psi_G^i \;\text{ if }\; (G = 1 \;\wedge\; s = 2) \\
& c = \omega_G \;\text{ if }\; (G = 1 \;\wedge\; (s = 0 \;\vee\; s = 1)) \\
& b = \max\{(1 + r)a - h' \,,\; 0\} \\
& U_s^i(c) = \mathbb{I}_{s \in \{0,1\}} \, \frac{c^{1-\gamma^i}}{1 - \gamma^i} + \mathbb{I}_{s=2} \, \left(\theta_{ADL}^i\right)^{-\gamma^i} \, \frac{\left(c + \kappa_{ADL}^i\right)^{1-\gamma^i}}{1 - \gamma^i} \\
& v^i(b) = \left(\theta_{beq}^i\right)^{-\gamma^i} \, \frac{\left(b + \kappa_{beq}^i\right)^{1-\gamma^i}}{1 - \gamma^i}.
\end{aligned}
$$

Individuals in this model have precautionary saving motives related to longevity and ADL risk. They also save in order to leave a bequest. Depending the value of preference parameters at the individual level (Θ^i), some people may have a strong desire to leave a bequest, some might care strongly about having large savings when in need of help with ADLs, and others might strongly prefer to spend while healthy. The bequest and ADL-related saving motives are tightly linked. If an individual ex-post over-saves for LTC because an expected health event never occurs, the strength of the bequest motive determines the cost of ex-post over-saving. Similarly, the size of a bequest reflects not only the active desire to leave a bequest, but also the sum of savings for other reasons combined with an uncertain timing of death. All together, preferences, demographic and financial variables, and the estimated health and longevity risks, determine saving behavior and demand for insurance products. A main contribution of this paper is estimation of Θ^i using new data and methods as described in Section 5.

See "Vanguard Research Initiative Technical Report: Long-term Care Model" for details on the computation of optimal decision rules. The means-tested government programs generate a non-concave value function and discontinuous optimal saving policy. To efficiently compute optimal decision rules we use a modified endogenous grid method suitable to this environment, building on insights from Fella (2014).

3.2 Calculating Activities of Daily Living Insurance Demand

We use each individuals' financial and demographic states and estimated preference parameters to calculate the model-implied demand for insurance. ADLI is modeled as a state contingent security that pays out whenever an individual is in the ADL health state ($s = 2$). To facilitate comparison to the stated-demand measure obtained in the survey, we introduce ADLI into the model by providing each individual with a one-time option to purchase ADLI at the financial and demographic states for the respondent taken from the survey. Purchasing this product entails paying a lump sum of $\$\tilde{y} \times p(t^i, s^i, g^i)$ at current age t^i in return for payout $\tilde{y}$ in each year that assistance with ADLs is needed.

The pricing function is such that the product is actuarially fair conditional on an individual's sex, age, health state, and access to a risk free outside asset with a 1 percent annual return. We show results for different loads that raise the price above actuarially fair in Section 6.3, including those typically seen in the LTCI market. Actuarially fair is defined such that the insurer selling this product makes zero expected profit (using the same health transition matrix as in the individual decision problem). ADL insurance only pays out when health state $s = 2$. Thus, ADLI that pays out $\tilde{y}$ per year when help with ADLs is needed has payout vector across health states $s \in \{0, 1, 2, 3\}$ given by the 1-by-4 vector $\tilde{y} \times [0,\ 0,\ 1,\ 0]$. Let $\vec{s}$ be an indicator vector for current health state, i.e., $\vec{s}$ is a 4-by-1 vector with elements s_j for $j \in \{0, 1, 2, 3\}$ equal to 1 if $s(t) = j$ and equal to zero for $s(t) \neq j$. Let $\pi_{g,t+k}$ be the 4-by-4 health state transition matrix for a person of sex g and age $t + k$. The insurance product is priced to equal the expected discounted stream of payments. Thus, for a person of age t, sex g, with current health status s, ADLI that pays out $\$\tilde{y}$ per year in which help with ADLs is needed costs

$$\begin{aligned} \tilde{y} \times p(t^i, s^i, g^i) &= \tilde{y} \times \sum_{\tau=0}^{T-t^i} \frac{1}{(1+r)^\tau} \left[p(s(t+\tau) = 2) | s(t) = s^i) \right] \\ &= \tilde{y} \times \sum_{\tau=0}^{T-t^i} \frac{1}{(1+r)^\tau} \left[[0, 0, 1, 0] \times \left(\prod_{k=0}^{\tau} \pi_{g^i, t^i+k} \right) \times \vec{s^i} \right]. \end{aligned} \tag{2}$$

For example, the resulting one-time cost for purchasing ADLI that pays out \$100K in each year when LTC is needed is as follows: For a healthy male, the cost is \$128K at age 55 and \$123K at age 65; for a healthy female, the cost is \$219K at age 55 and \$214K at age 65. The significantly higher cost for women reflects their longer life expectancy and higher probability of needing LTC. The slightly higher cost when age 55 reflects that the relatively small risk of needing long-term care prior to age 65 slightly outweighs the low risk-free interest rate used for discounting.

Given prices, demand for insurance is calculated as

$$\begin{aligned} D(a, y, t, s, h, g) &= \arg\max_{\tilde{y}}\ V(a - p(t, s, g)\tilde{y}, \hat{y}, t, s, h, g) \\ \hat{y} &= \{y(\tau) + \tilde{y}(s(\tau))\}_{\tau=t}^{T}, \end{aligned} \tag{3}$$

where $\hat{y}$ is the new income stream that is the sum of the original income stream plus the health-state dependent insurance payouts $\tilde{y}(s(\tau))$ and V is the value function evaluated at the new wealth level and income stream.

Who Wouldn't Buy Insurance? In this model there are two major risks that older individuals face: ADL and longevity risk. In response to these risks, people save so that they have enough money if they need help with ADLs, so that they do not run out of money and have to cut consumption due to a long life, and in order to leave a bequest. It is natural that risk-averse individuals might want to purchase insurance against these risks. Nonetheless, in this incomplete market setting, depending on preference, financial, and demographic state variables, individuals may prefer to self-insure via precautionary saving instead of locking-up their wealth in state-contingent assets. While this is especially true for low-wealth individuals for whom means-tested government care imposes a higher implicit tax, even people who do not anticipate using government care might not want to buy actuarially fair ADLI insurance.

Insurance is a way of transferring resources to specific states of the world; individuals purchase insurance based on consideration of the expected marginal value of wealth in different states. There are two main reasons other than government care that a person might not want to purchase actuarially fair ADLI. First, because of health-state-dependent preferences, the health state directly affects marginal utility by affecting the valuation of expenditures in the state. Second, the health state affects the *expected* marginal value of wealth because it changes the distribution of future states. For example, an individual is more likely to die next year if they currently need help with ADLs than if they are currently healthy.

To illustrate the direct effect of state-dependent preferences, consider a static decision of someone who needs to buy state-contingent assets before a health risk is realized, will need LTC with probability π and not need LTC with probability $1-\pi$, who has $\$X$ wealth, and preferences $\theta_{ADL}=1$ and $\kappa_{ADL}=X+Y$, with $X,Y>0$. Due to κ_{ADL}, this person would choose not to purchase any state-contingent asset that pays off if they need LTC (ADLI), precisely because they have a lower marginal utility in the ADL state than that in the healthy state even if they spent all of their wealth on the asset that pays when healthy. This simple logic extends to dynamic environments and to differences in θ. What matters for insurance purchase decisions is the relative marginal utility across states implied by an individual's state dependent preferences.

To illustrate the effect of the state-dependent distribution of future states, consider a person who values consumption in ordinary times more strongly than either leaving a bequest or spending when in need of LTC. This person might choose to hold onto their wealth in lieu of buying ADLI because needing LTC is associated with a higher probability of death and, thus, an expectation of fewer future periods of highly-valued ordinary consumption. Because ADLI transfers wealth into a state that is associated both with ADL-utility and higher mortality risk, the relative values of ADL and bequest utility can drive someone to not want to bring wealth into the ADL health state. The less valued is a bequest, the less someone would like to buy ADLI due to the higher probability of death in the ADL state.

If markets were complete, there were no government provided care, and people could hold long or short positions in insurance against health risks, then individuals in the model would generally take non-zero positions in actuarially fair ADL, death, and longevity insurance. See Koijen, Van Nieuwerburgh, and Yogo (2016) for a related analysis in the complete markets case. Very few people in our sample—and in the U.S.—own sizable amounts of such insurance (six percent of the VRI sample own a private annuity) and it is not possible to take short positions in these products. Thus, we choose to study the incomplete market setup in which we study ADLI demand in an environment in which the only other available insurance is self-insurance via saving in a risk free bond. Quantitatively, the main insurance market abstraction we make is that we do not account for life insurance ownership when computing ADLI demand.

4 Sample and Data Overview

This section presents the financial, demographic, and health data that describe individuals' circumstances and that are the state variables in the model. Section 5 presents estimates of the preferences, including a description of the survey data used in estimation. Section 6 uses the estimated model with these financial and demographic state variables and preferences to compute demand for ADLI for each person in the sample.

4.1 Financial, Demographic, and Expectations Data

This paper draws on the newly-developed Vanguard Research Initiative (VRI), a panel study of Vanguard clients aged 55 and older who had at least $10,000 in Vanguard accounts (see `http://ebp-projects.isr.umich.edu/VRI.html` for a complete description of the VRI, including all surveys and studies using this sample). The VRI has been stratified across two of Vanguard's major lines of business—individual accounts and retirement accounts through employers. In this paper we focus on single respondents. The sampling procedure and comparison of the VRI to the broader U.S. population is detailed in Ameriks, Caplin, Lee, Shapiro, and Tonetti (2014). Overall, the VRI sample is wealthier, more educated, more married, and healthier than the representative Health and Retirement Study (HRS) sample. Differences diminish, however, either when comparing to the HRS sample that meets the VRI age and wealth criteria or restricting focus to employer-based VRI members. By focusing on a sample that is wealthier than the typical American, crowd-out from Medicaid is likely less important in our study than it is in the U.S. population as a whole. Our study of moderately wealthy Americans, however, allows us to focus on those who are most likely to have demand for LTCI, since they could plausibly self-finance some of their care if they needed help with ADLs, but are not so wealthy that LTC poses a relatively small financial risk.

VRI respondents participated in three surveys used in this paper that were administered between June 2013 and August 2014.[5] VRI Survey 1 measures all of the demographic and financial state variables of the model for each respondent (wealth, income, age, sex, and health status), using novel methods for measuring household portfolios of assets and debts. Survey 2 has at its center both the key SSQs that identify preferences and the stated preference questions. Survey 3 gathers information on family structure as well as within-family intervivos transfers. Our final sample consists of single respondents who completed all three surveys and provided answers to all necessary survey questions. Table 1 gives summary statistics for the sample used in this paper.

	Wealth and Income					
	Mean	**10p**	**25p**	**50p**	**75p**	**90p**
Financial Wealth	715,655	115,000	271,731	545,935	1,021,443	1,602,000
Income	62,990	17,155	33,725	56,000	85,000	119,019

	Demographics							
	Age			**Health**			**Sex**	
	55-64	**65-74**	**75+**	**Good**	**Poor**	**ADL**	**Male**	**Female**
N=1,086	36.4%	43.2%	20.4%	94.6%	4.4%	1.0%	44.4%	55.6%

Table 1: Summary Statistics on Wealth, Income, Age, Health, and Sex: This table presents the marginal distributions of wealth, income, and demographic characteristics of the sample used in this paper. Individuals in this sample completed all three surveys and answered all necessary survey questions. This sample is composed of single (unmarried) households, so it is a subset of the VRI. Financial wealth is the sum of IRA, employer sponsored retirement, checking, saving, money market, mutual fund, certificate of deposit, brokerage, and educational related accounts plus the current cash value (if any) of life insurance and annuities. Income is defined as the sum of labor income, publicly and privately provided pensions, and disability income.

In addition to the SSQs and stated demand questions, respectively detailed in Sections 5 and 7, we use VRI measures of subjective longevity and health expectations, including the probability of needing help with ADLs in the future. We also use a measure of insurance holdings, the perceived quality of public long-term care relative to a typical private nursing home, and the expected cost of a year of care in a typical private nursing home in their community. Regarding family, we measure intervivos family transfers, number of children, and the probability that a family member would be the main caregiver if LTC were needed.

[5]Participants in this study receive a small incentive for participation in each survey in the form of sweepstakes for prizes, as well as a small monetary payment for completing all three surveys. Respondents also indicated a willingness to respond in order to aid and participate in a scientific endeavor. A set of initial respondents was designated as the pilot sample. A pilot version of each survey was fielded to this sample to test all aspects of the design and implementation.

4.2 Health and Mortality Estimation

In addition to using the financial and demographic data from the VRI, we estimate age and sex specific Markov transition matrices across health states using longitudinal data from an HRS subsample that meets the VRI age and wealth criteria. Individuals are defined as in good health if they report their health is good, very good, or excellent, and are defined to be in poor health if they report their health is poor or fair. A person is classified as needing help with ADLs if they list that they need significant help with at least one ADL and if they also receive help with that task. See Appendix A.1 for details on the mapping of states from the data to the model and the estimation procedure. See Appendix Figures A.1 and A.2 for the health-state transition matrices for males and females for all ages and health states.

To highlight the magnitude of LTC risk, Figure 1 presents estimated survival functions. The figure plots the probability of needing help with ADLs for more than $X \in \{0..10\}$ years for healthy men and women at various ages. The figures have several striking features. First, although most individuals will need help with ADLs at some point in their life, approximately 50 percent of males and 40 percent of females will never need help with ADLs. Second, there is substantial risk of spending extended time in need of help with ADLs. For men, approximately 23 percent will spend three or more years, 16 percent will spend four or more years, and 11 percent will spend five or more years needing help with ADLs. For women this risk is even larger, as approximately 31 percent will spend three or more years, 23 percent will spend four or more years, and 17 percent will spend five or more years needing help with ADLs. This substantial probability of needing care for many years highlights the large magnitude of LTC risk.

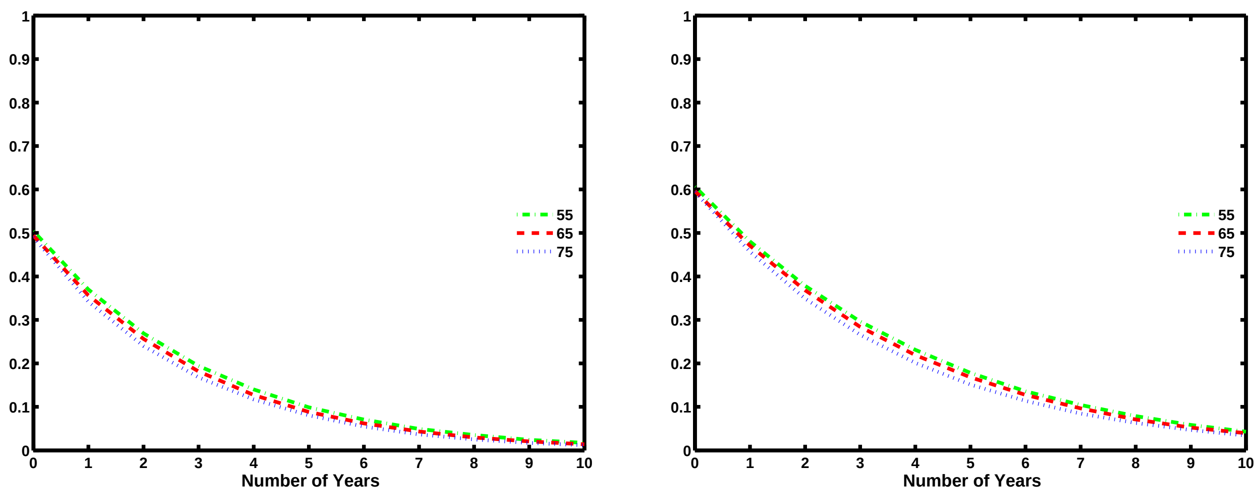

(a) Male ADL Survival Function

(b) Female ADL Survival Function

Figure 1: The Probability of Needing Help With ADLs for More Than $X \in \{0..10\}$ Years: Panel (a) presents the survival function of the total number of years of life spent needing help with ADLs for a 55, 65, and 75 year old male currently in good health according to the estimated health transition matrix. Panel (b) presents the corresponding figure for females.

5 Estimating Preferences using Strategic Survey Questions

To understand an individual's demand for LTC insurance products, it is essential to know that individual's preferences over related allocations. In this paper, we focus on preferences over consumption in ordinary times, over expenditures when in need of help with activities of daily living, and over bequests. A barrier to estimating such preferences is that available behavioral data does not provide enough information. The key idea behind strategic survey questions is that there are some choices that individuals might never face, but that would be very revealing of preferences if only such choice data were observed. SSQs ask respondents to make such choices hypothetically, by placing respondents in theoretically motivated scenarios that are significantly more detailed and structured than those in typical stated-preference questions. This paper makes use of nine variations of four SSQs asked to each survey respondent.[6] Since these preferences are at the heart of this paper, we dedicate this section to detailing how SSQ data identifies individuals' preference parameters (by construction), how we estimate preferences, and an analysis of the resulting estimated preferences.

To illustrate how SSQs work, imagine we want to know a person's coefficient of relative risk aversion (CRRA). One approach would be to directly ask "What is your coefficient of relative risk aversion?" on a survey, but that is obviously unlikely to be fruitful. The task for a survey designer is to write a question that is precise enough to elicit quantitative information about the respondent's CRRA, but in a simple enough format that the respondent can understand. Since the CRRA has strong implications on the trade-off between risky lotteries and certain amounts of wealth, recording an individuals choice when offered a lottery or a set amount of money would be informative about the value of their CRRA. This is a way of phrasing a question that non-economists can answer that still provides a direct map from response to structural parameter of interest. The SSQs in this paper adapt this logic to more complicated scenarios: when the difference between outcomes is not just the realization of a random variable, but also the utility function associated with different states of the world, when the choice is not contemporaneous, but would be made in the future with accompanying details about the state of the individual in that future, and when the scenario places restrictions on the choice set that are not likely to be faced in reality.

Each SSQ is designed first as a well-defined optimization problem, such that the optimal policy is a mapping from preference parameters to an allocation.[7] Then, an internet-based survey instrument is designed to present this choice problem in verbal form such that it is easy for respondents to understand the question and easy for them to report their choice. Appendix B details the design and implementation of the SSQ survey module, with information on how the survey questions were developed to help respondents understand the situation and choice while trying to make the verbal problem adhere as closely as possible to the math problem. Each SSQ begins with a statement of the scenario and rules, then provides comprehension-verification questions to ensure that respondents understand the most salient features of the scenario. A purpose-designed slider visualization is used for recording responses (see Appendix Figure B.1). Although the SSQs are complex, the slider neatly encapsulates the relevant trade-off and provides the resulting allocations and their economic interpretation in real time on-screen.

A number of features of the VRI data give us confidence in the credibility and quality of the survey measures. To this end, we present an analysis of SSQ responses in Appendix C. Most importantly, we report the raw data of the SSQ responses in a histogram for each SSQ variant. We then summarize objective and subjective measures of respondent comprehension. Finally, we discuss measures of the internal and external coherence of SSQ responses, describing patterns in the data both within individuals across SSQs and between SSQ responses and other variables of interest (e.g., financial, demographic, and subjective expectations data).

[6]Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2017) also use these data. That paper uses only the mean response for each SSQ variant together with wealth data to estimate homogeneous preferences, while this paper estimates individual-specific preference parameters using each person's responses to each SSQ variant. That paper also has no insurance market, so its question and findings do not address the central issues of this paper, which is the demand for insurance.

[7]While the SSQs were designed with specific functional forms in mind and while we use these functional forms to produce estimates of preference parameters, they provide valid information about preferences much more generally.

5.1 Identification

Identification of utility function parameters is achieved by matching survey responses with optimal responses to a mathematical representation of the SSQ questions. Since the optimal policies are functions of preference parameters, we map the SSQ responses to parameters by inverting the optimal policy function. The text of the third type of SSQ (SSQ 3) asks individuals to split wealth W between spending on self in the last year of life when help with ADLs is needed versus leaving a bequest. While estimation is done jointly for all parameters using all SSQs, for exposition, we sketch the identification argument for θ^i_{beq} and κ^i_{beq} using SSQ 3 assuming that γ^i, θ^i_{ADL}, and κ^i_{ADL} are known. A mathematical representation of SSQ 3 is the following optimization problem, in which each type of expenditure is valued with the state-specific utility function:

$$\begin{aligned} \max_{z_1,z_2} \quad & \left(\theta^i_{ADL}\right)^{-\gamma^i} \frac{\left(z_1+\kappa^i_{ADL}\right)^{1-\gamma^i}}{1-\gamma^i} + \left(\theta^i_{beq}\right)^{-\gamma^i} \frac{\left(z_2+\kappa^i_{beq}\right)^{1-\gamma^i}}{1-\gamma^i} \\ \text{s.t.} \quad & z_1 + z_2 \leq W \\ & z_1 \geq 0;\ z_2 \geq 0. \end{aligned} \tag{4}$$

The optimal allocation rule is given by

$$z_1 = \begin{cases} 0 & \text{if } \left(\theta^i_{beq}\left(W+\kappa^i_{beq}\right)\right)^{-\gamma^i} - \left(\theta^i_{ADL}\,\kappa^i_{ADL}\right)^{-\gamma^i} > 0 \\ W & \text{if } \left(\theta^i_{ADL}\left(W+\kappa^i_{ADL}\right)\right)^{-\gamma^i} - \left(\theta^i_{beq}\,\kappa^i_{beq}\right)^{-\gamma^i} > 0 \\ \frac{\theta^i_{beq}(W+\kappa^i_{beq})-\theta^i_{ADL}\kappa^i_{ADL}}{\theta^i_{ADL}+\theta^i_{beq}} & \text{otherwise.} \end{cases} \tag{5}$$

Conditional on γ^i , θ^i_{ADL}, and κ^i_{ADL}, the interior response is linear in wealth, and thus θ^i_{beq} and κ^i_{beq} are identified by two interior responses to the question posed with different wealth levels. Because SSQ 3 is fielded for variants at three different wealth levels (and these parameters also impact the response to SSQ 4), the system is overidentified. Identification of other parameters from the remaining SSQs follow a similar argument, mapping survey responses to the optimal responses of the mathematical representation of the SSQ question. All together, these responses identify the preference parameters of the model.[8]

In addition to SSQ 3 described above, we posed three other SSQs. A brief summary of the SSQs and their variants is presented in Table 2. Expanding on the material in the table, the key trade-offs posed in the SSQs are as follows.

1. SSQ 1 asks about willingness to take a risky bet over annual consumption. It compares consumption at a given level (W) to a lottery with a 50 percent chance that consumption is doubled and a 50 percent chance consumption shrinks by fraction λ^*. It finds the maximum loss (λ^*) that an individual is willing to face and prefer the risky lottery.[9]
2. SSQ 2 asks individuals facing uncertain future health to allocate wealth to the state of the world in which they are healthy or the state in which they need help with ADLs. With given probability (π) the person will be healthy and with probability (1-π) they will need help with ADLs. The person faces a portfolio allocation problem in which, before the health risk is realized, they must allocate a given amount of wealth (W) into state-contingent assets that pay only when healthy (z_1) or only when in need of help with ADLs (z_2).

[8] The parameters ω_G and β are not identified by any of the SSQs, and thus are calibrated to standard values from the literature. We perform sensitivity analysis around these values.

[9] This is similar to the survey questions and identification strategy as those developed in Barsky, Juster, Kimball, and Shapiro (1997) and Kimball, Sahm, and Shapiro (2008).

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

Table 2: Link Between Parameters and SSQs: The first column briefly summarizes the tradeoffs, while the second lists the underlying optimization problem. The third column lists how question parameters were changed for different variations of each SSQ, where W is wealth and $1-\pi$ is the probability of needing LTC. The z_1 in SSQ 4 is the optimal z_1 function calculated in SSQ 3. The fourth column lists the parameters that determine optimal responses in the model. This table is reproduced from Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2017).

3. SSQ 3 places individuals in the last year of their life when they need help with ADLs for certain and asks them to allocate a given amount of wealth (W) to either spending on self (z_1) or to a bequest (z_2).
4. SSQ 4 asks individuals how much wealth they would need to have (W^*) in order to purchase private LTC instead of using government provided care.

In "Vanguard Research Initiative Technical Report: Long-term Care Strategic Survey Questions" we present the text for each SSQ, as well as the optimization problem and optimal allocations (as a function of preference parameters) corresponding to each SSQ.

5.2 Maximum Likelihood Estimation of the Preference Parameters

This section presents estimates of the individual preference parameters that best match the SSQ data. These preferences are essential individual characteristics at the core of the modeled ADLI demand exercise. As summarized in Table 2 there are four types of SSQs, some asked multiple times at different scenario parameters, resulting in nine SSQ variants in total. We denote each individual i's set of responses to the 9 SSQ variants as $\hat{Z}_i = \{\hat{z}_k^i\}_{k=1}^9$. Recall each individual i's set of preference parameters is $\Theta^i = \{\gamma^i, \theta^i_{ADL}, \kappa^i_{ADL}, \theta^i_{beq}, \kappa^i_{beq}, \psi^i_G\}$. To estimate Θ^i, we assume that the recorded survey response is the true response that reflects preferences plus some error. For each individual we assume a response process that permits an analytical likelihood function and then use the 9 SSQ variants to estimate the parameter set that generated each individual's responses by maximum likelihood estimation.[10]

To derive the likelihood function, we denote the true response to the k^{th} SSQ as $z_k(\Theta^i)$, which we calculate by solving the optimization problem associated with the SSQs, e.g., equation 4 as the solution to the problem in 5. We

[10] We use MLE for tractability, but GMM using SSQ moments is equally feasible. A SMM approach using both SSQ moments and moments of the wealth distribution, as done in Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2017) for homogeneous preferences, is computationally infeasible because the dynamic program would need to be solved for all of the individuals who have their own preferences at each iteration of the candidate parameters.

assume each individual's response is reported with normally distributed errors. That is, let measured responses be

$$\hat{z}_k^i = z_k(\Theta^i) + \epsilon_k^i, \tag{6}$$

where $\epsilon_k^i \sim \mathbb{N}(0, \sigma_{k,i}^2)$ and ϵ_k^i denotes the realization of individual i's response error to SSQ variant k. For robustness, in Section 6.3 we show results for a multiplicative error structure that assumes $\log \hat{z}_k^i = \log z_k(\Theta^i) + \epsilon_k^i$.

For the six preference parameters to be identified at an individual level from 9 questions, the error distribution must be a function of no more than three free parameters. This is satisfied by specifying $\sigma_{k,i}^2$ to be a function of a question specific and an individual specific component. Specifically, we assume that the standard deviation of the response error to question k is linear in the maximum feasible response W_k and individual scaling factor σ^i, so that $\sigma_{k,i} = \sigma^i \times W_k$. The idiosyncratic component of the error distribution accounts for differences in the precision with which different individuals report answers. The question specific component takes into account the different scales of the nine SSQ variants and thus normalizes the error standard deviation according to the feasible response size. This roughly corresponds to assuming response errors are proportional to the level of responses. Note that W_k is naturally defined in each question by the budget constraint, except in SSQ 4. In estimation using SSQ 4, W_k is set to the 95^{th} percentile of the survey responses, resulting in \$500,000 as the factor that scales σ^i.[11]

This specification yields the following closed form expression for the likelihood of observing a response to each question as a function of $\left(\Theta^i, \sigma^i\right)$:

$$\mathcal{L}_k(\Theta^i, \sigma^i | \hat{z}_k^i) = \begin{cases} F_{\sigma_{k,i}^2}(-z_k(\Theta^i)) & \text{if } \hat{z}_k^i = 0 \\ f_{\sigma_{k,i}^2}(\hat{z}_k^i - z_k(\Theta^i)) & \text{if } 0 < \hat{z}_k^i < W_k \\ 1 - F_{\sigma_{k,i}^2}(W_k - z_k(\Theta^i)) & \text{if } \hat{z}_k^i = W_k. \end{cases} \tag{7}$$

The boundary cases account for the possibility that survey response error causes the response to violate the budget constraint, and $F_{\sigma_{k,i}^2}$ and $f_{\sigma_{k,i}^2}$ denote the mean-zero normal cdf and pdf with variances $\sigma_{k,i}^2$. Independence of survey response errors yields a multiplicatively separable likelihood function for individual i's full response set $\hat{Z}^i$,

$$\mathcal{L}(\Theta^i, \sigma^i | \hat{Z}^i) = \prod_{k=1}^{9} \mathcal{L}_k(\Theta^i, \sigma^i | \hat{z}_k^i). \tag{8}$$

The MLE estimates individual-level parameter sets,

$$\left\{\hat{\Theta}^i, \hat{\sigma}^i\right\} = \arg\max \mathcal{L}(\Theta^i, \sigma^i | \hat{Z}^i).$$

Identification is achieved via multiple responses to SSQ variants at different scenario parameterizations.[12] Parameters

[11] Results are not sensitive to large variation in W_k for SSQ 4.

[12] This is in contrast to Barsky, Juster, Kimball, and Shapiro (1997) and Kimball, Sahm, and Shapiro (2008), which use multiple responses to the same question across time, although we share the same additive normal error structure. There are two main differences between the estimation approach of this paper and that of Barsky, Juster, Kimball, and Shapiro (1997) or Kimball, Sahm, and Shapiro (2008). First, these previous studies estimate a log-normal population distribution of preference parameters to accommodate the discrete cutoffs that are built into the design of the HRS questions. Having continuous responses allows us to treat the population distribution of preference parameters non-parametrically. Second, this study estimates multiple preference parameters for each individual, whereas these previous studies focus on estimating only the risk aversion parameter for each individual.

are identified for those with few boundary responses, specifically fewer than three boundary responses in total and fewer than two boundary responses on the three SSQ 3 variants. All subsequent analysis is restricted to the 89 percent of respondents that satisfy this condition.

The parameter sets are estimated with reasonable precision. Since the same parameter appears in the solution to multiple questions, there are cross-equation restrictions that parameters must satisfy. The individual component of the error, σ_i, is a measure of how much response error is required to bring the survey responses in line with the functional forms imposed in the theory. For the large majority of individuals, the response error is very low: σ_i has a median value of 0.06. This implies that when individuals have \$100,000 to allocate, the median response error has a standard deviation of \$6,000. Furthermore, σ_i is less than 0.17 for over 95 percent of the population. The full distribution of σ_i is presented in Figure 2.

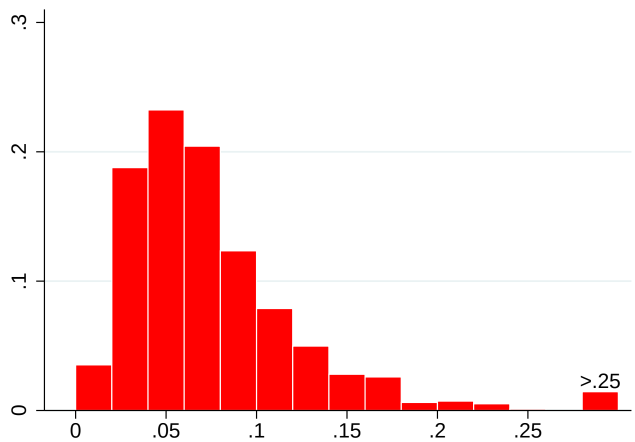

Figure 2: Distribution of Individual Response Error Standard Deviation σ_i

Before presenting the resulting estimated preferences, one technical note on how we will use these estimates is in order. To account for uncertainty around estimated parameter values when later calculating model-implied insurance demand, we resample five parameter sets for each individual from their distribution of estimates and calculate the demand for each parameter set. That is, using the parametric assumptions, we perform a wild bootstrap by adding different error realizations to the point estimates. Taking the average of these demand measures integrates out error in modeled demand caused by parameter uncertainty. For the remainder of the paper, all reported baseline results reflect these bootstrapped estimates.[13]

5.3 Analyzing the Estimated Preference Parameters

The result of the estimation procedure is the joint distribution of 6 parameters per person by 963 people. Since it is difficult to display such a high dimensional object, we provide the marginal distribution for each parameter and, in lieu of the copula, the correlation between parameters. Table 3 presents percentiles of the marginal distributions for the estimated population parameter distribution. κ is interpreted as thousands of dollars of expenditures and θ is the asymptotic marginal utility multiplier relative to ordinary consumption (as wealth grows large and the effect of κ dissipates). The median marginal estimates imply a relative risk aversion parameter $\gamma = 4.45$, ADL expenditure as a necessity ($\kappa_{ADL} < 0$) with high marginal valuation ($\theta_{ADL} < 1$), bequests as a significant luxury ($\kappa_{beq} > 0$) with a

[13] For the robustness exercises in Section 6.3, due to computational run time limitations, we use one sample of preference parameters that is held constant across exercises.

Marginal Distribution of Parameters

	γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
10%	2.03	.27	-83.66	.16	-41.22	19.98
25%	2.99	.44	-51.77	.26	6.96	39.49
50%	4.45	.90	-12.12	.54	98.05	59.16
75%	6.52	2.26	39.45	1.89	286.13	97.77
90%	9.65	6.62	130.74	7.11	643.96	166.25

Correlations of Parameters

	γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
γ	1.00					
θ_{ADL}	-.18	1.00				
κ_{ADL}	-.09	-.10	1.00			
θ_{beq}	-.17	.53	-.10	1.00		
κ_{beq}	-.21	.01	.27	-.10	1.00	
ψ_G	.07	.00	-.31	.02	-.22	1.00

Table 3: Estimated Parameter Distributions: The marginal distributions of each parameter are presented in the top panel table above. Note that each column is the marginal distribution of the specified parameter and that the parameter values in any given row do not correspond to any individual's preferences. The final line presents the parameters estimated from the same model with homogeneous preferences matched to SSQ and wealth distribution moments. Correlations of estimated parameter values are presented in the bottom panel.

high marginal valuation ($\theta_{beq} < 1$), and a public long-term care dollar equivalent of $\$59,160$ (ψ_G). For exposition, using the median parameter values, the dollar equivalent of public long-term care corresponds to an expenditure of \$41,063 in a model without state dependent preferences.[14]

For a point of comparison, when restricting preferences to be homogeneous, we estimate $\gamma = 3.45, \theta_{ADL} = 0.86, \kappa_{ADL} = -37.61, \theta_{beq} = 1.77, \kappa_{beq} = 18.47,$ and $\psi_G = 54.07$. Given that the presented marginals do not account for the correlation between parameters, there is no simple mapping from the marginals to the homogeneous preference case. Comparison does, however, show consistency in qualitative patterns, with the homogeneous parameters contained between the $25^{th} - 75^{th}$ percentiles of the estimated heterogeneous parameter distribution.

Interpreting Preference Parameters Using Simple Synthetic Choice Problems. Given that it is hard to interpret preference parameters in isolation, both because parameter values are inherently difficult to interpret and because the interpretation of any one parameter depends on the values of other parameters, we describe the estimated distribution of preferences by analyzing choices implied by the preferences. The idea is to represent the strength of the spending motives implied by the different utility functions by showing implied expenditures in simple-to-understand choice problems before using the estimated preference parameters in the full structural model. Thus, the following illustrative figures do not represent spending predicted by the full model nor that expected to be observed in data.

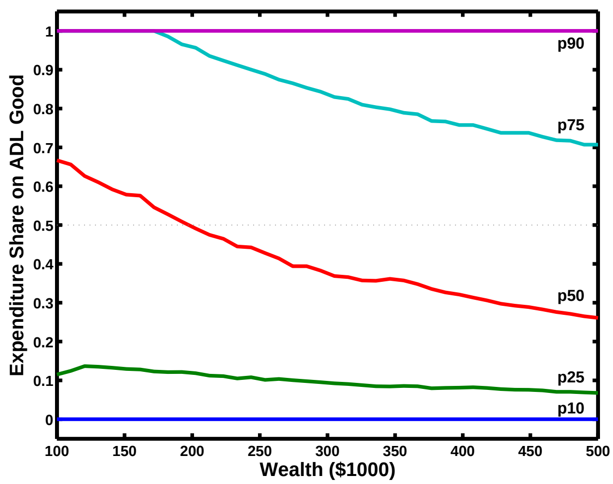

Figure 3: Distribution of SSQ 3 Responses Implied by Estimated Parameters: This figure plots quantiles of the distribution of allocations to the ADL state in response to SSQ 3, for different levels of wealth W, that are implied by the estimated distribution of preference parameters. The problem is not well defined for those with negative κ s.t. $-\kappa > W$, so the horizontal axis starts at \$100K and people are added into the figure as their problem becomes defined.

In Figure 3, we present the 10th/25th/50th/75th/90th percentiles of allocations to the ADL state in response to SSQ 3 (ADL Spending vs. Bequest) that are implied by the estimated distribution of preference parameters.[15] Higher

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

[15]Specifically, given the estimated distribution of preference parameters, we plot quantiles of the distribution of z_1 that solve the following last-year-of-life allocation problem for different levels of wealth W:

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

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

Since the problem is not well defined for those with negative κ s.t. $-\kappa > W$, we start the horizontal axis at \$100K and add people into the figure as their problem becomes defined. Since almost all ill-defined problems are because κ_{ADL} is too negative, the figure provides a rough lower bound on the strength of the ADL saving motive in the population.

percentiles allocate more to expenditures when needing help with ADLs at the cost of a lower bequest.[16] Due to the κ parameters, the allocations are not invariant to the wealth level. At \$100K, those in the 50th percentile of allocations to the ADL state spend \$65K on own expenditure when needing help with ADLs and leave \$35K as a bequest. The 75th percentile person leaves no bequest, while the 25th percentile person leaves \$90K as a bequest. This demonstrates an enormous heterogeneity in estimated preferences. At lower levels of wealth, much of the behavior is driven by the large differences in κ_{ADL} and κ_{beq}. At \$200K, some still so highly value spending on self when in need of help with ADLs relative to leaving a bequest that they leave no bequest, but the median person almost splits the money evenly. At all wealth levels, there are always some people for whom the bequest motive dominates. As wealth increases, the θ terms are more important than the κ terms in determining spending patterns and people shift towards leaving large bequests. This is exactly what the concept of bequests as a luxury good captures.

In summary, there is a large degree of heterogeneity in estimated preferences, with spending on self when in need of help with ADLs typically viewed as a necessity and leaving a bequest viewed as a luxury. At lower levels of wealth, spending on self when needing help with ADLs dominates leaving a bequest for most individuals in the sample even though many display a strong desire to leave a bequest. At high levels of annualized wealth (e.g., \$400K), the bequest motive dominates for most people, with the median person leaving 70 percent of wealth as a bequest. Nonetheless, even at \$400K, there is still a large fraction of the population spending most of their money on self (e.g., the 75th percentile spends 70 percent on self).

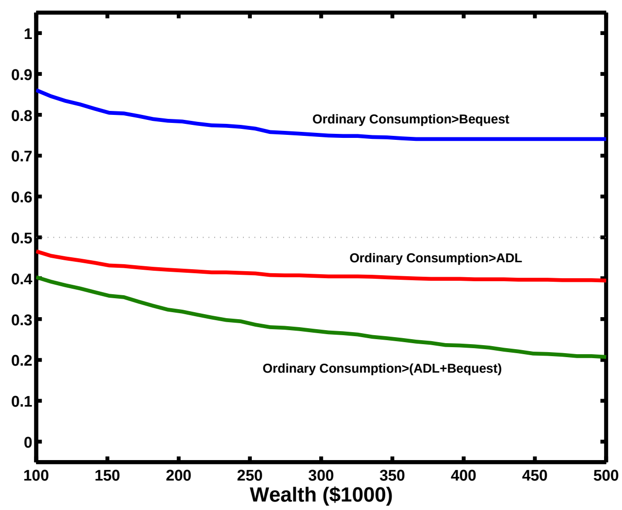

Figure 4: Expenditure in the three good synthetic choice problem: The above figures present statistics on expenditure in the three-good synthetic choice problem in which we treat the utility function associated with ordinary health, needing help with ADLs, and bequests as the utility function associated with three goods purchased contemporaneously. The figure presents the fraction of the population spending more on the "Ordinary Consumption good" than on the "ADL good," the fraction of the population spending more on the "Ordinary Consumption good" than on the "Bequest good," and the fraction of the population spending more on the "Ordinary Consumption good" than on the sum of the "ADL good" and "Bequest good." The problem is not well defined for those with negative κ s.t. $-\kappa > W$, so the horizontal axis starts at \$100K and people are added into the figure as their problem becomes defined.

[16]Here we treat bequest spending and spending on self when needing help with ADLs as two different goods valued with different utility functions in a simple allocation problem. In the full structural model, bequest utility is a one time payoff upon death, while ADL utility represents an annual flow utility. Compared to the full dynamic model, this 1-period representation makes the ADL-related saving and insurance demand motive weaker because ADL utility can be active for multiple periods, but makes it stronger because death is reached with certainty while the ADL state is stochastic and may never be reached.

Figure 4 presents statistics on expenditure in a three-good synthetic choice problem in which we treat the utility function associated with good health, needing help with ADLs, and bequests as the utility function associated with three goods purchased contemporaneously. The figure presents the fraction of the population spending more on the "Ordinary Consumption Good" than on the "ADL good," the fraction of the population spending more on the "Ordinary Consumption Good" than on the "Bequest good," and the fraction of the population spending more on the "Ordinary Consumption Good" than on the sum of the "ADL good" and "Bequest good" when solving the following problem:

$$\max_{x_1,x_2,x_3} \quad \frac{(x_1)^{1-\gamma}}{1-\gamma} + (\theta_{ADL})^{-\gamma}\,\frac{(x_2+\kappa_{ADL})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma}\,\frac{(x_3+\kappa_{beq})^{1-\gamma}}{1-\gamma} \tag{9}$$
$$\text{s.t. } x_1+x_2+x_3 \leq W$$
$$x_1,\ x_2,\ x_3 \geq 0;\ x_2 \geq -\kappa_{ADL};\ x_3 \geq -\kappa_{beq}.$$

About half the sample has preferences such that spending from the ordinary utility function is stronger than that from the ADL function, while the other half spends more on the ADL good than ordinary consumption. Although this result uses the joint distribution of preferences, this is captured roughly by the median κ_{ADL} being negative but not too far from 0 and the median θ_{ADL} being a bit less than one. The large majority of individuals, around 75 to 85 percent of the population, have a stronger per period spending motive from ordinary consumption than from the bequest motive, reflecting the large positive estimated κ_{beq} for most people. There is, however, a non-trivial 15–25 percent of the population with stronger bequest motives. To get a sense that spending on self when healthy is very important to most people, even relative to ADL and bequest motives, at \$100K 40 percent of the population would spend more on the ordinary consumption good than on the bequest and ADL good combined.

In summary, the distribution of estimated parameters suggest there is significant preference heterogeneity with regards to spending in ordinary times, when in need of help with ADLs, and as a bequest. Nonetheless, there are clear patterns present for many people in the data. Most survey respondents have positive but moderate risk aversion, a strong desire to spend in ordinary times, view spending when in need of help with ADLs as a necessary good that is valued highly on the margin, and view bequests as a luxury good that requires a large outlay before eventually being highly valued on the margin. As we show in Section 6, it is exactly these patterns in preferences that largely determine the substantial modeled demand for activities of daily living insurance while also generating heterogeneity in demand that shrinks the LTCI puzzle relative to the homogeneous preference case.

6 Results: The Long-Term-Care Insurance Puzzle

We now predict demand for Activities of Daily Living insurance for each person in the sample. The optimal policy that solves equation 3 (for each individual's observed financial/demographic data and estimated preference parameters) determines demand for ADLI if the surveyed individual were to be offered the opportunity to purchase the product at the time of the survey.

Our paper's main result is that our baseline estimation predicts 59 percent of respondents would have positive demand for ADLI, while only 22 percent of the VRI sample actually own any private LTCI. Thus, the long-term care insurance puzzle—that modeled insurance demand is larger than observed LTCI holdings—is sizable.

In the remainder of this section we show that the extensive margin of the purchase decision is significantly determined by preference heterogeneity; that incorporating heterogeneous preferences significantly reduces the size of the

LTCI puzzle relative to a homogeneous preference benchmark; that the puzzle is present across the cross-section of the wealth distribution; that modeled demand is robust to various model specifications and in different subsamples; and that the intensive margin of demand is large for many individuals, as is the associated consumer surplus. We conclude that most of our sample would substantially gain from purchasing high quality LTCI if it were available, but there are many people who would not.

6.1 LTCI Demand and the Importance of Heterogeneous Preferences

Our prediction that 59% of our sample has positive ADLI demand implies that a majority of individuals assign a high valuation to—and want to insure—wealth in the ADL state but do not reveal this demand in the market, presumably due to low quality LTCI products. However, there is a substantial minority for whom purchasing is not predicted to be attractive. Majority interest is thus not built into our result, but rather a result of financial/demographic states and desires as inferred from the responses to SSQs.

	Average Preference Parameters Conditional on Modeled ADLI Demand					
ADLI Demand	γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
Positive (59%)	5.57	4.01	-11.33	2.07	118.13	76.77
Zero (41%)	4.31	23.11	33.86	22.69	278.84	77.42

	Average State Variables Conditional on Modeled ADLI Demand				
ADLI Demand	**Age**	**Income Quintile**	**Wealth**	**Percent Male**	**Percent Good Health** ($s = 0$)
Positive (59%)	69.2	3.24	821,129	45	95
Zero (41%)	67.1	2.84	517,129	40	94

Table 4: Who Purchases ADLI? This table presents averages of demographic, financial, and preference variables for two groups: the 41 percent of the sample with zero modeled ADLI demand and the 59 percent of the sample with positive modeled ADLI demand.

To disentangle the respective roles of financial/demographic state variables and preferences in predicting demand, Table 4 compares the mean parameter values and state variables for individuals with positive modeled ADLI demand versus zero modeled ADLI demand. Regarding financial states, we observe that purchasers are on average wealthier and in a slightly higher income quintile, suggesting that ADLI is a normal good. Because prices are actuarially fair conditional on all demographics, we observe little difference in average demographic states across the two group.

We observe significant and intuitive differences in average preferences between purchasing and non-purchasing individuals. ADLI purchasers are more risk averse than non-purchasers. They also have a much stronger preference for expenditure when in the ADL state. The average κ_{ADL} of purchasers is negative yet positive for non-purchasers, so that purchasers value ADL-state expenditure as a necessity while non-purchasers perceive it as a luxury. Furthermore, the average marginal utility multiplier θ_{ADL} of non-purchasers is over five times larger than that of purchasers, representing a higher utility of wealth on the margin in the ADL-state for purchasers even if κ_{ADL} were the same. The theoretical implications of bequest motives are less clear-cut. On one hand, bequest motives decrease the desire to spend on self when needing help with ADLs by increasing the desire to hold on to bequeathable wealth. However, ADLI insures bequests against being depleted by potentially large expenditures when in the ADL state, which could occur if spending in the ADL state was highly valued. That is, ADLI transfers wealth to a state with higher mortality risk and, on the margin, people who place higher value on leaving a bequest find this more worthwhile. We find that purchasers of ADLI have a lower θ_{beq} (value bequest more on the margin) and κ_{beq} (view bequests as less of a luxury) than non-purchasers, suggesting that the second motive is dominant. Overall, preferences appear to play an important role in determining ADLI demand, with purchasers being more risk averse, having a greater desire to spend on self when healthy, and having stronger bequest motives than nonpurchasers.

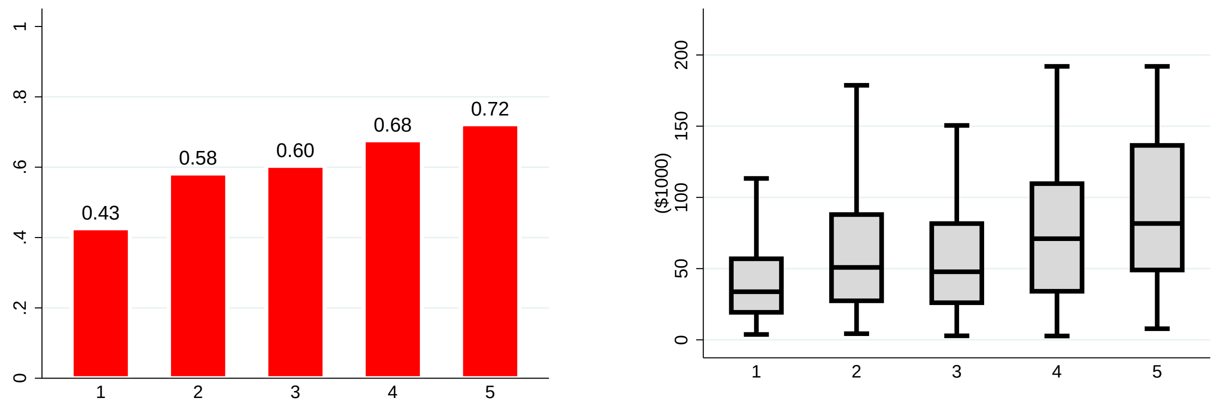

(a) Fraction Buying ADLI by ADL-Preference Quintile **(b)** Quantity of ADLI purchased by ADL-Preference Quintile

Figure 5: Heterogeneous Preferences and ADLI Demand: This figure presents ADLI demand by quintile according to a rank of individuals by the amount they spend in the synthetic choice problem on the "ADL Good" as described in equation 9 and Figure 4. The left panel presents the fraction of the population with positive modeled ADLI demand. The right panel presents the quantity of ADLI demanded for those with positive demand.

Figure 5 provides further evidence of the importance of preference heterogeneity in determining ADLI demand. Individuals are ranked by the strength of their ADL-utility function and then ADLI demand is plotted by quintiles of this rank. To rank individuals, we use the amount they spend on the "ADL Good" in the three-good synthetic choice problem described in equation 9 and Figure 4. The left panel of Figure 5 shows a steep gradient in the extensive margin of demand by ADL-utility rank. Comparing the highest and lowest ADL-utility rank quintiles, there is a 29 percentage point difference in the fraction of the population that wants to buy ADLI, which represents a 67 percent increase in people who want to buy ADLI. The right panel of Figure 5 presents box-plots of purchased income for individuals with positive ADLI demand, again by ADL-utility rank quintiles. On average, those in the lowest preference quintile buy insurance that pays around \$30K per year when needing LTC and those in the highest quintile buy around an \$80K per year payout, an over 150 percent increase in average demand. There is also substantial demand heterogeneity within ADL-utility rank quintile, representing within-quintile preference heterogeneity as well as heterogeneity in financial and demographic state variables. Such significant differences in the extensive and intensive margins of ADLI demand across ADL-utility rank quintiles suggests that individual differences in preferences is quantitatively important when predicting the distribution of ADLI demand.

The above analysis highlights the role of preference heterogeneity in determining ADLI demand. To show the implications of ignoring this source of heterogeneity, we predict demand using our estimate of homogeneous preferences presented in Section 5.3. Under this specification 78 percent of individuals purchase ADLI, a 19 percentage point increase over our 59 percent baseline. This higher modeled demand—and movement further away from observed LTCI holdings—derives from the average person's preferences assigning high values to spending when in need of care despite there being a substantial fraction of individuals who assign relatively low values to such spending (see Section 5.3). Failing to account for such differences in tastes for spending across states thus leads to a quantitatively significant overprediction of ADLI demand.

In summary, we find that preference heterogeneity is an important determinant of ADLI demand, that demand varies across people with different preferences in meaningful ways, and accounting for this heterogeneity significantly reduces the LTCI puzzle relative to a homogeneous preference reference.

6.2 The Long-Term-Care Insurance Puzzle is Present throughout the Cross-Section

The large difference between actual LTCI holdings and modeled ADLI holdings is not just concentrated in the higher-wealth individuals in the VRI sample, but is also present for those with savings similar to many older Americans. Figure 6 compares actual LTCI ownership and modeled ADLI demand conditional on wealth and income quintiles. The smallest wealth quintile has median wealth of \$115,000 and the smallest income quintile has median annual income of \$17,000, not dissimilar to the broader U.S. population. Both observed holdings and model predictions of ADLI demand are increasing in wealth and income. The difference between modeled and observed holdings is large and significant at all quintiles, confirming the robustness of the puzzle. We therefore conclude that the LTCI puzzle exists across all financial groupings in our sample and is not driven just by the very wealthy.

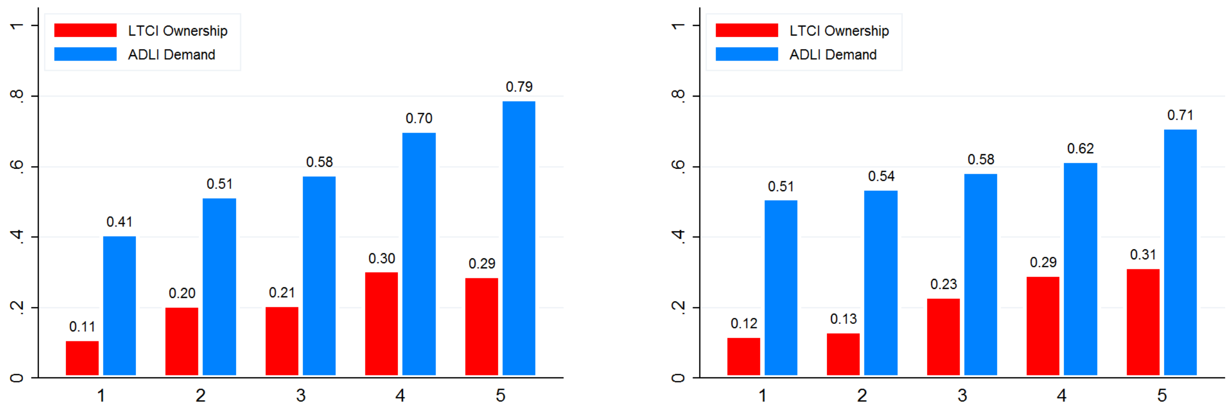

(a) by Wealth Quintile **(b)** by Income Quintile

Figure 6: LTCI Ownership vs. ADLI Demand. The above figures present ownership of LTCI/ADLI by wealth and income quintiles. The red bars on the left show the fraction of the population in a given quintile who own LTCI, while the blue bars on the right are the corresponding model predictions.

6.3 The Long-term-care Insurance Puzzle is Robust: Model and Subsample Sensitivity Analysis

To document the robustness of the LTCI puzzle, we present in Figure 7 ADLI demand calculated for different model specifications and for different subsamples. For reference, the top three rows show the 22 percent measured LTCI ownership rate, the 59 percent ownership rate of the baseline model, and the 78 percent ownership rate of the model with homogeneous preferences.[17]

First, we show that the existence of the LTCI puzzle is not sensitive to reasonable increases in the price of ADLI. To document the sensitivity of demand to prices, we compute demand when ADLI is priced with either a 10 percent, 20 percent, or 30 percent load above the actuarially fair price. Thus, if low observed insurance holdings were driven by high loads, the model under this specification should predict substantially lower demand. On the extensive margin,

[17] Although 22 percent of our sample currently owns LTCI, some may purchase it in the future. In the HRS, LTCI ownership is flat after age 65 and increases by about one-third from age 55 to 65. Assuming similar patterns in the VRI, accounting for future purchases would increase our baseline ownership number only modestly, by around 2 percentage points.

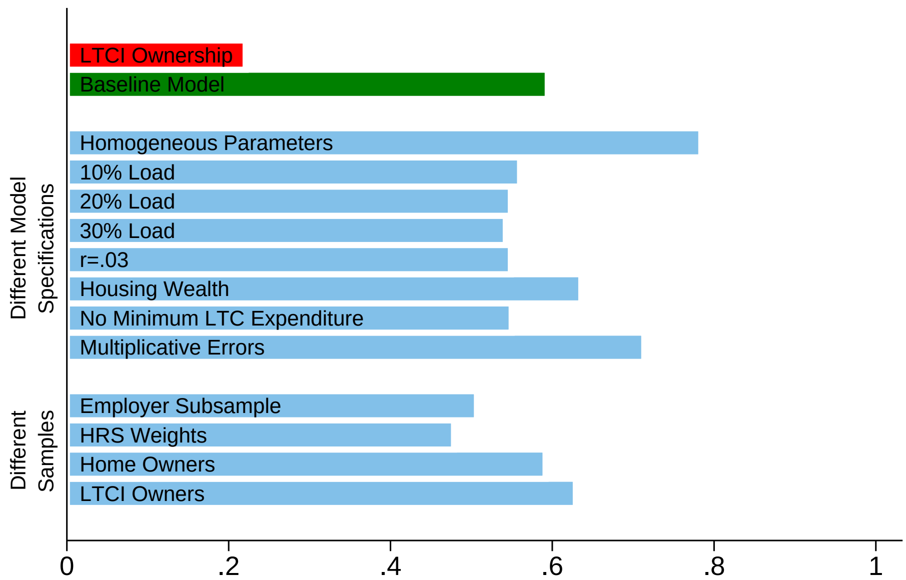

Figure 7: LTCI/ADLI Ownership Rates: Sensitivity to Alternative Parameters and Samples. This figure presents the fraction of the population that is predicted to have positive ADLI demand according to various changes to the model and sample. The top row shows the 22 percent of people who own LTCI in the VRI. The second row shows the prediction from the baseline specification of 59 percent ownership. The third row shows that the homogeneous preference model predicts 78 percent ownership. Subsequent rows present the sensitivity results.

the fraction of the population with positive demand for ADLI only drops from 59 percent at baseline to 56 percent under a 10 percent load and 54 percent under a 30 percent load. This foreshadows that those who buy ADLI have large consumer surplus, such that their demand is inframarginal of such price changes.

We then show that the existence of the LTCI puzzle is not sensitive to reasonable changes in the risk-free rate. We compute ADLI demand for the case in which consumers receive a risk free return of $r = 0.03$ on savings, while insurance products are still priced using $r = 0.01$. This exercise addresses two concerns. First, respondents might expect a higher return on wealth than the risk free rate, and so the baseline model might understate the saving motive. Second, this introduces a sizable load above actuarial fair pricing (equivalent to 18-35 percent on ADLI for males aged 55–85). Again, there is a small drop in the fraction of people with positive demand from 59 percent to 55 percent, suggesting a low return on savings in the model is not driving the puzzle.

In our baseline the wealth state variable is set equal to measured net financial wealth, but the results are robust to treating wealth as the sum of financial and housing wealth. Houses are complicated assets, since they have financial value that is difficult to calculate given indivisibility, search frictions, and fixed costs of sale, but also because they provide individual specific flow utility. As an upper bound, we add the full equity value of the primary home to financial wealth and predict ADLI demand. This increase in wealth only further exacerbates the puzzle, increasing the fraction of the population with positive demand to 64 percent.

Capturing the fact that LTC provision is essential for those in need and private long-term care is an expensive and lumpy expenditure, in the baseline model we set $\chi_{ADL} = \$40K$ as a minimum level of expenditure needed to obtain private LTC, i.e., $c \geq \$40K$ if $s = 2$ (needs help with ADLs) and no government care is provided. Results are again robust to removing this minimum expenditure constraint, with 60 percent predicted to demand ADLI. We also show that results are robust to assuming an error term that is log additive, as opposed to additive used in the baseline, leading to 71 percent of people with positive demand.

Finally, to address concerns about robustness outside of the VRI sample we repeat the analysis on different samples. First, we use a subsample of individuals restricted to respondents with employer sponsored Vanguard plans. The employer subsample is less wealthy than the general population, as displayed in Appendix Table E.2, and did not elect by themselves to become Vanguard clients. Thus, concerns of sample selection might be less severe amongst these individuals. We find that all qualitative results hold for this sample, with 51 percent of this population estimated to have positive demand for ADLI. Second, we reweight the population using weights that match the HRS on wealth and demographic variables (see Ameriks, Caplin, Lee, Shapiro, and Tonetti (2014)). Similar to the employer subsample, when reweighting to the HRS, even though the model predicts a lower 48 percent extensive margin of demand for ADLI there is still a clear prediction of high interest in these products relative to observed holdings. Third, we split the population into those who own LTCI and those who don't, with slightly more LTCI owners predicted to demand ADLI, at 63 percent of the population relative to 58 percent for LTCI non-owners. Lastly, demand is positive for 59 percent of homeowners, the same as for non-homeowners.

Thus, the clear model prediction of high interest in ADLI—and the puzzle that emerges when comparing this prediction to observed LTCI holdings—is significant and robust to alternative pricing, alternative measures of wealth, alternative preference estimation strategies, and in a number of subsamples. Incomplete markets and state-dependent preferences also robustly deliver ADLI demand well below 100 percent.

6.4 Those Who Would Buy ADLI Would Buy A Sizable Amount

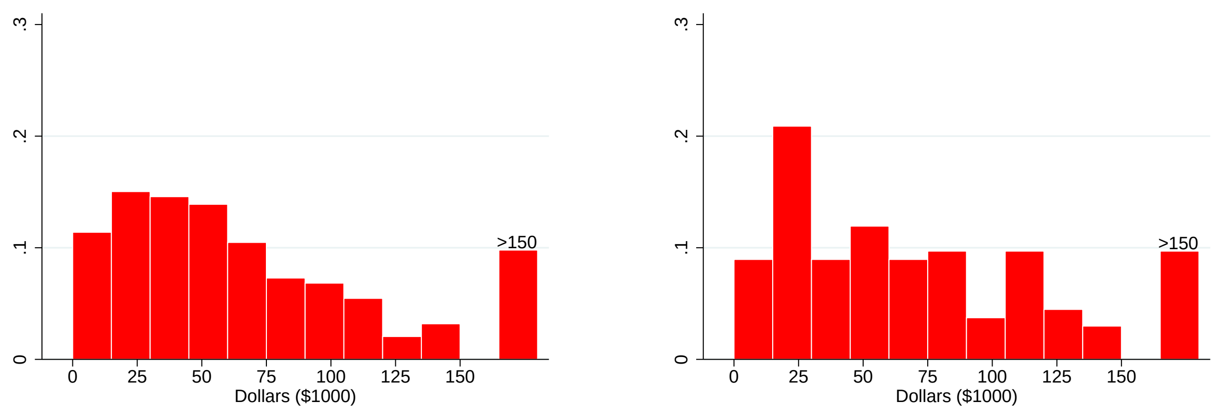

(a) LTCI Non-owners

(b) LTCI Owners

Figure 8: ADLI Quantity Demanded: This figure presents the histogram of the ADLI annual payout purchased in the model. The left panel plots ADLI demand for the 58 percent of the population of LTCI non-owners with positive modeled demand. The right panel plots ADLI demand for the 63 percent of the population of LTCI owners with positive modeled demand.

To establish the importance of the LTCI puzzle, we show that in addition to the large difference on the extensive margin between LTCI ownership and modeled demand for ADLI, the modeled quantity demanded is sizable. Because we do not have a measure of the quantity of insurance owned by those who hold private LTCI, we often restrict our analysis to the 78 percent of the population who do not own any private LTCI. This is the only population for whom we know the amount of private LTCI owned (zero) so that we can compare modeled demand to the known holdings. Nonetheless, we present in Figure 8 ADLI demand measures for both LTCI owners and LTCI non-owners, showing very similar model predictions. This similarity suggests ownership of LTCI is not driven by differences in demographic and financial variables or preferences, but features not captured in the model, e.g., opportunities to purchase LTCI linked to employer benefits.

For the baseline specification, 58 percent of people who do not own any private LTCI have positive modeled ADLI demand. Furthermore, as presented in Figure 8a, for those who have positive demand the average quantity demanded is about $67K in annual payout, the 10th percentile of demand is about $9K and the 90th percentile is around $150K. Compared to not owning any LTCI, these individuals are predicted to demand relatively large amounts of insurance. To put the quantity in context, the purchasers' median demand of a $55K payout is larger than the median income of an 80 year old, more than doubling income in the ADL-state during ages when help is most likely to be needed. Demand is also substantial for those who are likely able and eligible to purchase LTCI. Healthy females (males) aged 55–64 have median annual income of $58,000 ($62,000) and financial wealth of $455,000 ($405,000). Conditioning on this income, health, age, sex, and buying ADLI, median demand (across wealth and preferences) is $33,400 ($39,400) paid each year LTC is needed purchased at a one-time cost of $72,200 ($49,800). The size of the payouts seems reasonable, keeping in mind that an average one year stay in a nursing home costs $92K per year and costs of $150K per year are common in upscale nursing homes. Just as with the extensive margin, Appendix Table E.1 documents that the intensive margin of demand remains robust to many alternative assumptions and samples. To explore why the LTCI puzzle is so robust, we move beyond just analyzing quantity demanded to examining the estimated demand function for ADLI.

6.4.1 Consumer Surplus From ADLI Would Be Large

While the amount demanded at given prices is informative, there is further information in the properties of the demand function. It could be that people demand a large amount of ADLI, but they are near indifferent between the optimal ADLI purchased and no ADLI at all. To show that there is strong desire for better LTCI, we document that the elasticity of demand to price increases is small and the consumer surplus is large for most people.[18]

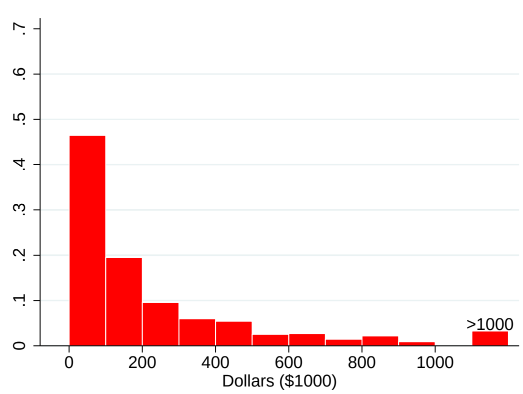

(a) Distribution of Consumer Surplus

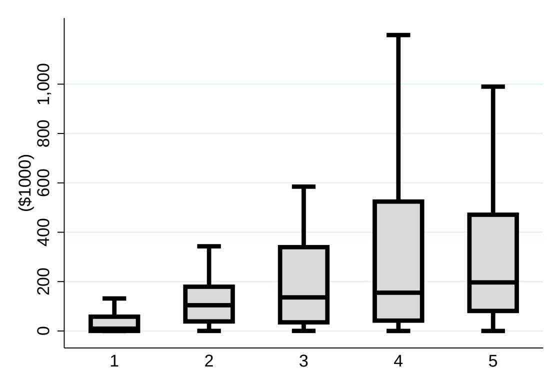

(b) Consumer Surplus by Wealth Quintile

Figure 9: Consumer Surplus: The left panel presents the histogram of consumer surplus for those who do not own LTCI. Consumer surplus is the maximum amount an individual would be willing to pay to purchase their desired amount of insurance above the price they actually paid. The right panel presents a box plot of the consumer surplus by wealth quintile.

Consumer Surplus. Consumer surplus is defined as the maximum amount an individual would be willing to pay in excess of the amount they actually paid for the quantity they demanded at the given price. This varies across people because they faced different prices (prices conditioned on sex, age, and health status), because the quantity demanded

[18] All analysis is presented for LTCI non-owners. As documented in Appendix E, results are very similar for LTCI owners.

differs as a function of demographics, financial variables, and preferences, and because the dollar value of the same quantity of ADLI at the same price depends on those individual-specific states. As documented in the left panel of Figure 9, many people have a consumer surplus above \$100K, with a non-trivial fraction of the population having a consumer surplus larger than \$200K. Our model thus predicts sizable demand for ADLI, suggesting a substantial missing market for higher quality LTCI. The right panel of Figure 9 shows that the median consumer surplus is small in dollar terms for the lowest wealth quintile, but is around \$100K for wealth quintile two rising slowly but steadily to \$200K for quintile five. In contrast, the consumer surplus for those who strongly value ADLI, as measured by the surplus at the 75th percentile, grows substantially from \$58K in wealth quintile one through \$524K in quintile four, but drops to \$454K in quintile five. Those in the highest wealth quintile have enough savings to self-insure and smooth consumption in all states of the world so the value of insurance is not as large to them. Those in the lowest wealth quintile are most likely to value the means tested Medicaid option, which implicitly lowers their value of private ADLI.

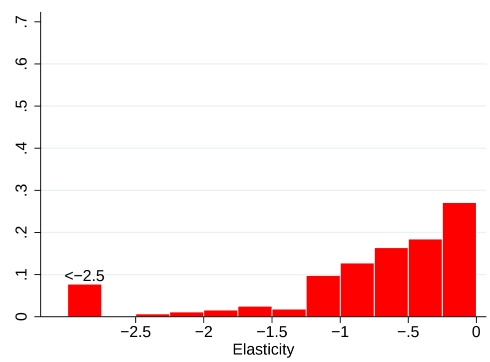

Figure 10: Distribution of the Price Elasticity of Demand: This figure presents the histogram of the elasticity of demand with respect to price for those who do not own LTCI. It plots the distribution of the percent change in demand to a one percent increase in price, local to the optimal demand level and given price.

Price Elasticity of Demand. Figure 10 plots the distribution of the price elasticity of demand, defined as the percentage change in quantity demanded for a one percent increase in price, local to the optimal quantity demanded. Overall, demand is not very price elastic, with around 80 percent of people having less than unit elasticity and about 50 percent having an elasticity less than 0.5 in absolute value. That the price elasticity is small and, as documented in Appendix Table E.1, that the intensive margin of demand does not change much between actuarially fair pricing and a 30 percent load suggests that the price of LTCI may not be the main unattractive feature of products currently in the market. It also suggests that consumer surplus is likely to remain large even at reasonably higher prices. As discussed in Section 2, there are many features of LTCI products that may contribute to low demand other than price.

6.5 Summary: Model-predicted ADLI Demand and Consumer Surplus are Large

Taken together, the large fraction of the total population predicted to demand ADLI, the robustness of this extensive margin of demand to alternative assumptions and samples, the large quantity of modeled demand and its robustness, the small price elasticity, and the large consumer surplus all document the LTCI puzzle: there is substantial demand for insuring the state of the world in which help is needed with ADLs which is at odds with the low holding of LTCI in the data. While heterogeneous preferences help to shrink the LTCI puzzle, a sizable puzzle remains. In the next section, we show that differences between LTCI available in the market and the idealized insurance in the model explains much of the remaining puzzle and suggests that there is a large unmet demand for insuring this health risk.

7 Results: Stated Demand for Activities of Daily Living Insurance

The LTCI puzzle is that modeled demand for ADLI is significantly higher than actual holdings of LTCI. To what extent does the LTCI puzzle derive from a quality gap between LTCI and modeled ADLI, given that ADLI is very different from the LTCI available in the market place? As discussed in Section 2, LTCI products have many unattractive features: consumers face default risk, possible unilateral increases in future premia, high loads, and a potentially adversarial claims process that has strict and uncertain conditions on when holders can claim. In this section we use additional information from VRI Survey 2: stated choice questions on the demand for improved insurance products. Stated choices provide a model-independent measure of demand for exactly the same ADLI product as modeled in Section 6. We have two completely different measures of the same demand. Both point to high demand for ADLI, even among those who, via lack of ownership, reveal low demand for available LTCI. This higher demand for a better insurance product suggests that the low quality of the available LTCI does indeed contribute to low LTCI holdings.

7.1 The Survey Instrument

Our survey elicits stated demand for ADLI.[19] For ADLI, a challenge in gathering this demand is that it concerns a form of insurance that is not available in the market place. For that reason the demand questions were preceded by the definition of the ADL state, defined as "needing significant help with activities such as eating, dressing, bathing, walking across a room, and getting in or out of bed." Moreover, when gathering demand information, we explicitly ask respondents to "make choices in hypothetical financial scenarios." The ADLI product is presented in the following frame.

> Please suppose that you are offered a hypothetical new form of insurance called **ADL insurance** with the following features:
>
> - You pay a one-time, nonrefundable lump sum to purchase this insurance.
> - If you need help with activities of daily living (ADLs), you will immediately receive a monthly cash benefit indexed for inflation.
> - For each **$10,000** you pay for this insurance, you will receive $Y per month indexed for inflation in any month in which you need help with ADLs.
> - The monthly cash benefit is set at the time of purchase and is not dependent on your actual expenses.
> - There is **no restriction** on the use of the insurance benefits. You are free to use benefits in any way you wish: to pay for a nursing home; a nurse to help at home; for some other form of help; or in literally any other way you would like.
> - An impartial third party who you trust will verify whether or not you need help with ADLs immediately, impartially, and with complete accuracy.
> - The insurance is priced fairly based on your gender, age, and current health.
> - There is no risk that the insurance company will default or change the terms of the policy.

When gathering stated demand information, we price the product for each individual at the expected value of payouts conditional on age, sex, and current health based on the estimated health transition probabilities, determining "$Y" in the frame above.[20] This is reinforced by the qualitative statement that the pricing is fair. After all information is provided, demand is collected in two steps. We first ask respondents whether or not they would have any interest

[19] Beshears, Choi, Laibson, Madrian, and Zeldes (2014) use stated choice questions to study determinants of annuity demand, specifically to examine what improved features of annuity products could increase demand. Brown, Kapteyn, Luttmer, and Mitchell (2017) elicit stated purchase and sale values for annuities and link these spreads to potential explanations for heterogeneity in financial decision-making abilities.

[20] To price the insurance products in the stated demand survey instrument, we used a health transition matrix estimated on an HRS sample that is representative of the U.S. population. Modeled demand when using the U.S. representative health transition matrix is little changed: for the wealthier VRI sample the lower per-year probability of needing help with ADLs is offset by the longer life expectancy.

in purchasing ADLI were it available. If the answer is affirmative, we ask how large of a monthly benefit they would purchase, updating and reporting the up-front cost of the purchase in real-time as they enter their decision. In the top right corner of the answer screen we present a link to a hover screen that presents the full specification of the product in case the respondent would like to review any features prior to reporting their demand.

Credibility of Stated Demand. While there are valid concerns whether stated preferences match normative preferences, Beshears, Choi, Laibson, and Madrian (2008) note that the likelihood of significant disparities decreases when decisions require active choice, are simple, are familiar, are not influenced by third-party marketing, and limit intertemporal considerations. By forcing individuals to make an active choice we attempt to limit fall-back to the default option. Comprehension checks on the definition of ADLs, careful design of product presentation, use of hover screens to make forgotten information available, and an answer screen that dynamically highlights the trade-off to purchasing this product as the choice is made serve to reduce the complexity. In addition, the question makes it clear that the product is a one-time offer to reduce concerns surrounding intertemporal decisions, and because ADLI does not exist in practice concerns around third party marketing are minimal. Thus, our stated demand questions are designed to address factors that facilitate reporting of normative preference.

To analyze the coherence of the stated demands, we conduct a probit regression of the decision to buy and an OLS regression on the amount purchased in the subsample of respondents that reported positive demand. Full results are included in Appendix Table D.8. Respondents who report higher probabilities of experiencing extended time in the ADL state are more likely to state demand for ADLI. This suggests that the prices quoted to these individuals may be cheaper than actuarially fair and that adverse selection affects ADLI purchases. There is also evidence that individuals who indicate a more favorable opinion of publicly provided LTC have less of a desire to purchase. Conditional on having positive demand, we observe that respondents that own private LTC insurance and that have higher than average LTC cost beliefs purchase more, while those that report a more favorable opinion of publicly provided LTC purchase less. Few demographic variables are significant, likely reflecting the survey practice of calculating actuarially fair pricing conditional on sex, age, and health status.

7.2 Differences between ADLI and LTCI Explain much of the Remaining LTCI Puzzle

Thirty one percent of respondents reported that they would purchase a strictly positive amount of ADLI. Preexisting LTCI holdings may have crowded out ADLI demand, causing individuals that would otherwise desire ADLI not to demand any more. A measure combining individuals who either own LTCI or state a demand to purchase ADLI yields 46 percent of the population expressing a desire to insure ADL risks. Thus, a combined extensive margin measure of stated demand suggests ownership three-quarters that of modeled demand. These different measures of ownership are summarized in Figure 11. Comparing the model with homogeneous preferences to measured LTCI holdings would suggest a sizable puzzle of 56 (i.e., 78-22) percentage points. Accounting for heterogeneous preferences and the difference between LTCI products in the market and idealized insurance against the ADL health state explains much of the apparent difference between model and data, resulting in a 13 (i.e., 59-46) percentage point puzzle remaining. That the union of stated ADLI demand and actual LTCI holdings is significantly larger than holdings of LTCI, shows that there is latent demand for higher quality insurance products for ADL risks. That this measure is lower than modeled ADLI demand suggests that not all of the difference between modeled and actual holdings is attributable to specific features of the LTCI products currently available in the market.

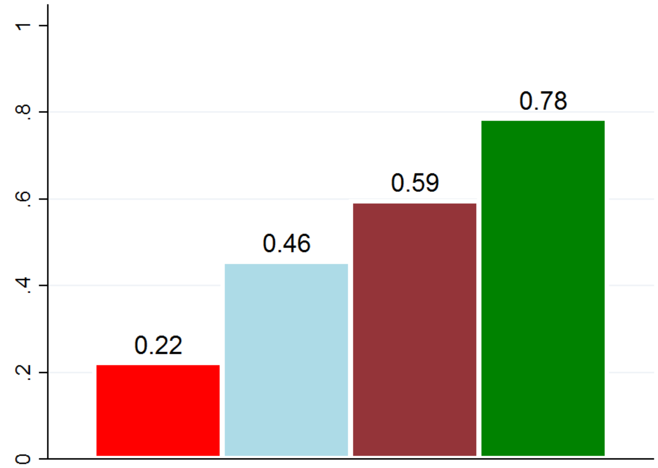

Figure 11: Fraction of Population Owning LTCI or Demanding ADLI: This figure presents various measures of the fraction of the population with positive LTCI ownership. Column 1 is actual holdings of a private LTCI in the sample. Column 2 is the union of private ownership and stated demand, capturing those with an expressed desire to insure the ADL health state. Column 3 is modeled ADLI demand in our baseline model. Column 4 is modeled ADLI demand in the homogeneous preference benchmark model.

The Difference between Modeled and Stated ADLI Demand is Concentrated in High Wealth Individuals. As seen in Figure 12, for lower wealth individuals the combined owned-or-stated measure of the fraction of the population with positive demand is much closer to the corresponding model prediction. At the lowest wealth quintile, the amount of people with positive stated-or-owned demand is 89 percent of the model predicted. At the second wealth quintile the owned-or-stated measure is a remarkable 98 percent of that predicted by the model. At higher wealth quintiles owned-or-stated goes from 78 percent of model predicted in quintile three to 58 percent in quintile five. The close fit of stated and modeled demand suggests that, for the lower wealth quintiles, a large part of the puzzle can be explained by the low quality of LTCI products available in the market. For the higher wealth quintiles, however, there seems to be an additional source significantly contributing to the LTCI puzzle.

The Modeled Quantity of ADLI Demand is Larger than Stated Demand. Figure 13 presents the histogram of stated ADLI demand for those with positive demand. 30 percent of people who do not own LTCI and 33 percent of LTCI owners state positive ADLI demand. Although median stated demand is zero, there is sizable stated demand: one third of LTCI non-owners that report positive demand indicate a desire to purchase more than a \$20K yearly payout, while the 90^{th} percentile of this conditional demand distribution is \$48K. For those who do own LTCI, there is more interest in ADLI, with less demand in the \$0–10K payout range, and more concentration in the \$20–40K payout range.[21] These stated-demands are similar to the average annual LTCI benefits received, which was \$25K in 2012 (The American Association for Long-Term Care Insurance (2018)).

Since the only people for whom we know the quantity of insurance owned against the ADL health realization is those who own zero LTCI, we next compare stated and modeled demand for those who do not own LTCI. While on the extensive margin stated and modeled demand are quite close, modeled demand is systematically larger on the intensive margin, as seen by comparison of Figures 8 and 13. Table 5 documents the distributions of stated and modeled demands for LTCI non-owners. Comparing the distributions of demand in rows 1 and 2 of Table 5, we

[21]The somewhat higher intensive and extensive margin demand by LTCI-holders is not informative about the existence of crowding-out, since we do not observe their counterfactual demand when they do not own LTCI.

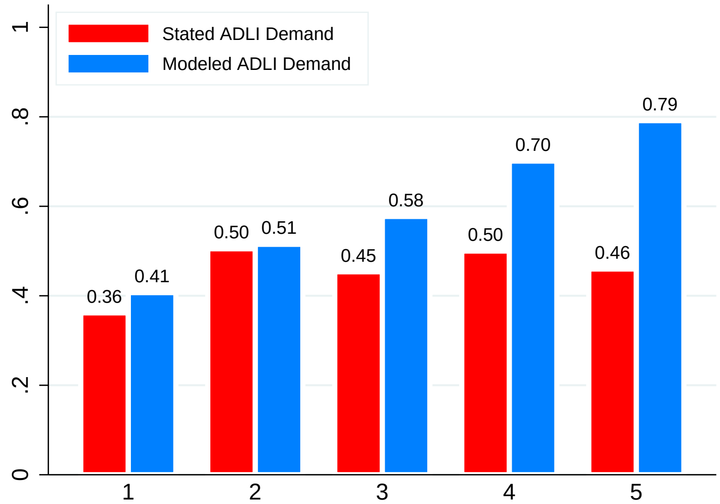

Figure 12: ADLI Demand by Wealth Quintile: Stated and Modeled Demand This figure presents the fraction of the population with positive demand for ADLI by wealth quintile according to stated demand and modeled demand. The red bars on the left show the fraction of the population in a given quintile who either own LTCI or state positive demand for ADLI in the survey, while the green bars on the right are the corresponding modeled demand.

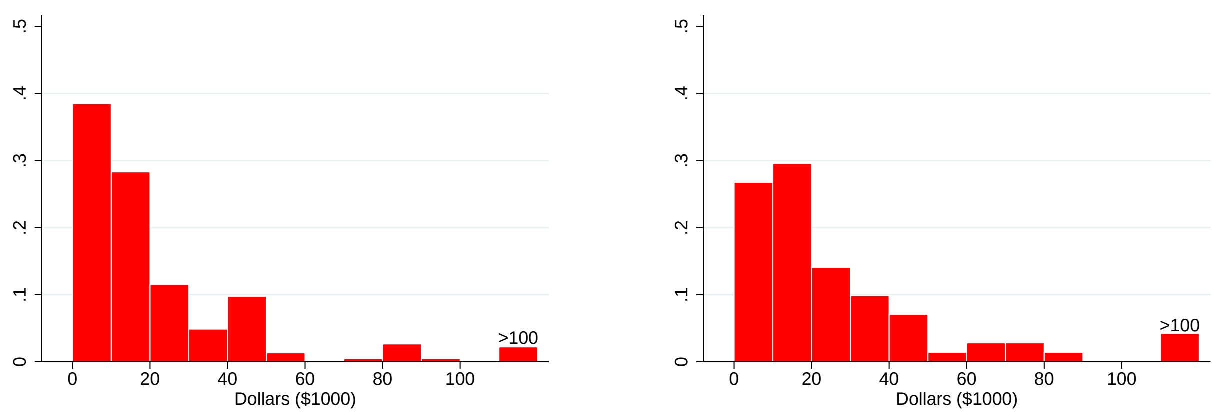

(a) LTCI Non-owners

(b) LTCI Owners

Figure 13: Stated ADLI Quantity Demanded: This figure presents the histogram of the ADLI annual payout purchased as stated by survey respondents. The left panel plots stated ADLI demand for the 30.1 percent of the population of LTCI non-owners with positive stated demand. The right panel plots ADLI demand for the 33.3 percent of the population of LTCI owners with positive stated demand.

observe that the mean, median, and all percentiles of modeled ADLI demand distribution are at least as large as those in the stated ADLI demand distribution. This is seen more directly in the distribution of differences in the third row of Table 5. The median demand difference is $11K and mean difference is $32K, suggesting for many individuals that the model predicts higher demand.

	%>0	mean	p5	p10	p25	p50	p75	p90	p95
Modeled	58	39,282	0	0	0	17,347	62,204	118,820	155,650
Stated	30	6,793	0	0	0	0	6,000	20,400	40,800
Modeled−Stated		32,489	-18,720	-8,859	0	10,585	57,859	105,877	151,377

Table 5: Distribution of Differences in ADLI Demand: This table presents the distribution of each of the ADLI demand measures for those individuals that do not own LTCI. The top line presents the distribution of modeled demand and the middle line presents the distribution of stated demand from the survey. The bottom line presents the distribution of the differences between modeled and stated demand (not the difference of the distributions).

In summary, both stated and modeled ADLI demand are significantly larger than existing holdings of LTCI. First, these two independent measures indicate a robust finding of substantial desire to insure against possible LTC need. Second, the fraction of the population with positive ADLI demand is similar across measures, suggesting a sizable part of the low purchase of LTCI and the LTCI puzzle is driven by the low perceived quality of LTCI products. Given that the extensive margin of stated and modeled ADLI demand differ for higher wealth people, there are likely other motives generating an LTCI puzzle. Last, there still exists an intensive margin LTCI puzzle in which the model predicts more ADLI demand than people state. In the next section we provide an quantitative exploration into possible features driving the intensive margin LTCI puzzle.

7.3 Predictors of the Estimated vs. Stated ADLI Demand Gap

In this section we analyze our model-based and model-free demand measures to provide insight into possible reasons for their difference. Generally, there are two reason why the model and stated demand measures might not align. First, factors included in our demand measures might not be properly specified. Second, we might exclude considerations from our demand measures that should be taken into account. To identify whether such omitted considerations contribute to the difference between modeled and stated demand, we develop a general econometric method that identifies sources of model misspecification both related to included state variables or preferences and omitted variables. We define an omitted variable as any variable that respondents may consider when forming demand that is not included in the model. Such omitted variables, denoted q, bias model estimates of demand from an individual's true demand.

Defining the difference between modeled and stated demand as

$$\eta_i \equiv \text{Modeled}_i - \text{Stated}_i, \tag{10}$$

we decompose the difference into factors related to state variables, preferences, and omitted variables q. We do so by estimating the following equation, with details on the derivation of the estimation equation included in Appendix F.[22]

$$\eta_i = \beta^x C_i^x + \beta^\Theta C_i^\Theta + \Gamma q_i + \epsilon_i \tag{11}$$
$$H_0: \ \beta^\Theta = 0; \ \beta^x = 0; \ \Gamma = 0.$$

We allow the difference to be a nonlinear function of financial and demographic states and preferences, modeled non-parametrically by partitioning individuals into regions of the state and parameter space. Variables C are indicators of the partition element to which each individual belongs. That is, individuals of a similar age, sex, income, health, and wealth will be grouped into the same element of the partition C_i^x. Analogously, those with similar preferences will be grouped into the same element of C_i^Θ. Estimation of $\Gamma > 0$ indicates model misspecification related to variable q that generates higher demand for insurance relative to stated, while $\Gamma < 0$ indicates model misspecification that generates lower demand for insurance.

	ADLI difference (Γ)							
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
$\mathbb{I}_{Transfers}$	9,786							6,821
	(7,057)							(7,312)
$\mathbb{I}_{child}$		-5,986						-2,725
		(7,409)						(7,498)
$\mathbb{I}_{Real\ Estate}$			-6,235					-5,643
			(6,524)					(6,530)
$\mathbb{I}_{College}$				-1,742				1,284
				(6,607)				(6,804)
$\mathbb{I}_{Comp.\ Test}$					-11,541			-10,868
					(6,144)			(6,271)
$\mathbb{I}_{Family\ Care}$						-3,922		-514
						(6,007)		(6,227)
$\mathbb{I}_{ADL\ help}$							-10,379	-9,867
							(6,199)	(6,206)

Table 6: Explaining the Difference Between Modeled and Stated Demand for ADLI. This table presents the Γ coefficient from estimation of equation 11 on the sample of respondents who do not own LTCI. The coefficients on β^x and β^Θ are omitted, but in all estimations these coefficients are jointly significant at the 1% level. See text and Appendix F for discussion of β^x and β^Θ. Standard Errors are included in parentheses.

Table 6 presents results from estimating Equation 11 on the sample of people who do not own LTCI with q

[22]Note that the above specification ignores misspecification caused by the interaction of state variables and preferences. Attempts to control for these interaction effects through partial correlations of individual preference parameters and state variables do not significantly change any of the results presented in this paper, although estimates become less precise. Furthermore, we do not find significant evidence that omitted factors predict demand measures separately.

defined as variables related to omitted model elements, omitted motives that would be difficult to model, and potential behavioral biases. For all variables considered (except college education and having a child), we define an indicator that is equal to one if the respondent's characteristic is above the median value of that characteristic for the sample. For example, $\mathbb{I}_{ADL\ help}$ is equal to one if the respondents subjective probability of needing help with the activities of daily living for at least one year is above the median respondent's. To address concerns of error around the estimated parameters and demands included in this regression we follow Rubin (1987) and estimate this equation for multiple replicates generated by resampling from the estimated parameter error distribution. Reported coefficients and standard errors reflect this multiple imputation/wild bootstrapping approach.

We find that the gap is smaller for those who have in the past made large intervivos transfers to a descendant. This gap is consistent with the idea that the warm-glow bequest specification that is the current workhorse in the quantitative literature since De Nardi (2004) is not a complete account of the bonds between generations. Model enrichment to capture other family-related motives may be warranted, for example the family-provided care studied by Mommaerts (2016).[23] With regard to survey comprehension, the gap is smaller for those who performed better at the SSQ comprehension tests. This suggests that individuals whose responses better reflect their preferences have stated demands that better align with economic models (as proposed by Beshears, Choi, Laibson, and Madrian (2008)). It is therefore plausible that demand in a working ADLI market would be somewhat higher than stated preferences indicate. Finally, the gap is smaller for those with adverse private information on the likely length of needing care. This suggests that adverse selection may be significant problem, and that market provision of actuarially fair LTCI may be infeasible (Hendren (2013), Braun, Kopecky, and Koreshkova (2017a)). Variables such as real estate holdings, education, and the probability of receiving care from family, among others, do not significantly predict the difference.

8 Conclusion

Older Americans face many risks as they age. Foremost among these risks is needing assistance with activities of daily living as health declines. This assistance can be provided either in home or in a long-term care facility. The cost of this long-term care is high and need for care can be prolonged.

Why, then, do so few have private long-term care insurance? This paper uses the Vanguard Research Initiative to investigate the factors that low measured LTCI holdings reflect. The VRI includes batteries of questions that we designed to elicit the demand for insurance against long-term-care risks. Using answers to these questions together with a structural model of decision-making in the face of late-in-life risks, the paper sheds light on whether the lack of demand for LTCI reflects individual preferences, individual circumstances, or defects in the LTCI products available in the market.

Our ability to distinguish among preferences, circumstances, and market defects as explanation for low purchase of LTCI derives from having multiple measures of demand. We define an idealized insurance product, "Activities of Daily Living Insurance," that provides income when individuals need long-term care. ADLI has none of the defects of the LTCI available in the marketplace. Using the VRI measures, we present both modeled and stated demand for ADLI. Modeled and stated demand for ADLI are both substantial. Fifty nine percent of respondents have positive modeled demand. Conditional on positive modeled demand, the amount demanded is substantial. For those with positive modeled ADLI demand, median demand for typical females (males) aged 55–64 is \$33K (\$39K) paid each year LTC is needed at a one-time cost of \$72K (\$50K).

[23] See Barro (1974), Becker (1974), Bernheim, Shleifer, and Summers (1985), Barro and Becker (1988), Altonji, Hayashi, and Kotlikoff (1997), McGarry (1999), Light and McGarry (2004) for different treatments of intergenerational motives. Abel and Warshawsky (1988) provides discussion of different modeling approaches for rationalizing bequests.

Modeled and stated demand are correlated within individual, though stated demand is lower. The difference between modeled demand, which is derived from circumstances that we determine in the construction of the model, and stated demand, which depends on individual circumstances, arises from differences between modeled and actual circumstances. For example, stated and modeled demand could differ due to unmodeled differences in circumstances like an expectation of care from a child. The similarity in popularity of ADLI across measures suggests that these unmodeled circumstances do not loom too large.

While providing a partial explanation for this under-insurance puzzle, our results suggest that flaws in existing products do not fully explain it, especially for the highest wealth respondents. The differences in stated ADLI demand and actual LTCI purchase should largely be due to difference in product characteristics. This gap is large, suggesting substantial unmet insurance demand in the market place. Accounting for differences in individuals' financial holdings, demographics, and health-state dependent preferences, model predictions indicate that better quality LTCI would be far more widely held than are products in the market, be held in large quantities, and generate substantial consumer surplus.

This paper is able to make progress on quantifying explanations for the demand (or the lack of demand) for insurance against late-in-life risks. It combines the strategic survey questions (SSQ) approach, which allows us to estimate relevant preference parameters at the individual level, with modeling of choices in the face of the large-scale risks that older households experience. The SSQs elicit choices in hypothetical circumstances, but they are based on scenarios that are highly relevant as individuals prepare for retirement and then make choices about spending and health care during retirement. These purpose-designed measures of preference parameters, together with rich information on individual economic and health circumstances from the VRI, allow the choices of individual respondents to be studied through well-defined economic models. The paper discusses how to design and implement SSQs that provide credible estimates of individual preference parameters, and then shows that the SSQ responses have substantial internal and external validity. This paper, by posing and then partially answering the long-term care insurance puzzle, demonstrates the usefulness of this approach.

There are substantial challenges in providing market solutions to the need for long-term care insurance. Our findings imply, however, that there is substantial unmet demand for improved insurance against the need for long-term care and suggest that improvements in insurance offerings would be a boon to older Americans.

References

ABEL, A. B., AND M. WARSHAWSKY (1988): "Specification of the Joy of Giving: Insights from Altruism," *The Review of Economic Statistics*, 70(1), 145–149.

ALTONJI, J., F. HAYASHI, AND L. KOTLIKOFF (1997): "Parental Altruism and Inter Vivos Transfers: Theory and Evidence," *Journal of Political Economy*, 105(6), 1121–1166.

AMERIKS, J., J. BRIGGS, A. CAPLIN, M. D. SHAPIRO, AND C. TONETTI (2016a): "Long-term Care Model," *Vanguard Research Initiative Technical Report*, `http://ebp-projects.isr.umich.edu/VRI/papers/VRI-TechReport-LTC-Model.pdf`.

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

——— (2017): "Long-Term-Care Utility and Late-in-Life Saving," *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," *Scandanavian Actuarial Journal*, 3, 1–42.

BARCZYK, D., AND M. KREDLER (2015): "Altruism, Exchange, Attachment to the House, or by Accident—Why do People Leave Bequests?," *working paper*.

BARRO, R., AND G. BECKER (1988): "A Reformulation of the Economic Theory of Fertility," *Quarterly Journal of Economics*, 103(1), 1–25.

BARRO, R. J. (1974): "Are Government Bonds Net Wealth?," *The Journal of Political Economy*, 82(6), 1095–1117.

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.

BECKER, G. S. (1974): "A Theory of Social Interactions," *Journal of Political Economy*, 82(6), 1063–1093.

BERNHEIM, B. D., A. SHLEIFER, AND L. H. SUMMERS (1985): "The Strategic Bequest Motive," *The Journal of Political Economy*, 93(6), 1045–1076.

BESHEARS, J., J. J. CHOI, D. LAIBSON, AND B. C. MADRIAN (2008): "How are Preferences Revealed?," *Journal of Public Economics*, 92(8), 1787–1794.

BESHEARS, J., J. J. CHOI, D. LAIBSON, B. C. MADRIAN, AND S. P. ZELDES (2014): "What Makes Annuitization More Appealing?," *Journal of Public Economics*, 116, 2–16.

BRAUN, R. A., K. A. KOPECKY, AND T. KORESHKOVA (2017a): "Old, Frail, and Uninsured: Accounting for Puzzles in the U.S. Long-Term Care Insurance Market," *working paper*.

——— (2017b): "Old, Sick, Alone and Poor: A Welfare Analysis of Old-Age Social Insurance Programs," *Review of Economic Studies*, 84, 580–612.

BROWN, J. R., AND A. FINKELSTEIN (2007): "Why is the Market for Long-term Care Insurance so Small?," *Journal of Public Economics*, 91, 1967–1991.

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

——— (2011): "Insuring Long-Term Care in the United States," *The Journal of Economic Perspectives*, pp. 119–141.

BROWN, J. R., G. S. GODA, AND K. MCGARRY (2012): "Long-term Care Insurance Demand Limited by Beliefs about Needs, Concerns about Insurers, and Care Available from Family," *Health Affairs*, 31(6), 1294–1302.

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

BROWN, J. R., A. KAPTEYN, E. F. P. LUTTMER, AND O. S. MITCHELL (2017): "Cognitive Constraints on Valuing Annuities," *Journal of the European Economic Association*, 15(2), 429–462.

CARROLL, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112(1), 1–55.

COHEN, M., R. KAUR, AND B. DARNELL (2013): "Exiting the Market: Understanding the Factors Behind Carriers' Decision to Leave the Long-Term Care Insurance Market," in *U.S. Department of Health and Human Services Assistant Secretary for Planning and Evaluation Office of Disability, Aging and Long-Term Care Policy*. U.S. Department of Health and Human Services Assistant Secretary for Planning and Evaluation Office of Disability, Aging and Long-Term Care Policy.

CUTLER, D. (1996): "Why Don't Markets Insure Long-Term Risk?," *working paper*.

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

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

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

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.

FINKELSTEIN, A., AND K. MCGARRY (2006): "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review*, 96(4), 938–958.

FRIEDBERG, L., W. HOU, W. SUN, A. WEBB, AND Z. LI (2014): "New Evidence on the Risk of Requiring Long-Term Care," *Boston College Center for Retirement Research WP 2014-12*.

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

GODA, G. S. (2011): "The impact of state tax subsidies for private long-term care insurance on coverage and Medicaid expenditures," *Journal of Public Economics*, 95(7), 744–757.

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

HACKMANN, M. B. (2017): "Incentivizing Better Quality of Care: The Role of Medicaid and Competition in The Nursing Home Industry," *working paper*.

HENDREN, N. (2013): "Private Information and Insurance Rejections," *Econometrica*, 81(5), 1713–1762.

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

HONG, J. H., AND J.-V. RÌOS-RULL (2007): “Social Security, Life Insurance and Annuities for Families,” *Journal of Monetary Economics*, 54(1), 118–140.

——— (2012): “Life Insurance and Household Consumption,” *The American Economic Review*, 102(7), 3701–3730.

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.

HURD, M. D., P.-C. MICHAUD, AND S. ROHWEDDER (2017): “Distribution of Lifetime Nursing Home Use and of Out-of-Pocket Spending,” *Proceedings of the National Academy of Sciences*, 114(37), 9838–9842.

INKMANN, J., P. LOPES, AND A. MICHAELIDES (2011): “How Deep is the Annuity Market Participation Puzzle?,” *Review of Financial Studies*, 24(1), 279–319.

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

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

KOIJEN, R. S. J., AND M. YOGO (2015): “The Cost of Financial Frictions for Life Insurers,” *American Economic Review*, 105(1), 445–475.

——— (2016): “Shadow Insurance,” *Econometrica*, 84(3), 1265–1287.

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

LAITNER, J., D. SILVERMAN, AND D. STOLYAROV (2015): “The Role of Annuitized Wealth in Post-Retirement Behavior,” *working paper*.

LIGHT, A., AND K. MCGARRY (2004): “Why Parents Play Favorites: Explanations for Unequal Bequests,” *The American Economic Review*, 94(5), 1669–1681.

LOCKWOOD, L. (2012): “Bequest Motives and the Annuity Puzzle,” *Review of Economic Dynamics*, 15(2), 226–243.

——— (2018): “Incidental Bequests and the Choice to Self-Insure Late-Life Risks,” *American Economic Review*, forthcoming.

MANSKI, C. F. (2004): “Measuring Expectations,” *Econometrica*, 72(5), 1329–1376.

MCGARRY, K. (1999): “Inter Vivos Transfers and Intended Bequests,” *Journal of Public Economics*, 73(3), 321–351.

MOMMAERTS, C. (2016): “Long-Term Care Insurance and the Family,” *working paper*.

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

PAULY, M. V. (1990): “The Rational Nonpurchase of Long-Term-Care Insurance,” *Journal of Political Economy*, 98(1), 153–168.

RUBIN, D. B. (1987): *Multiple Imputation for Nonresponse in Surveys*. John Wiley & Sons.

Rubin, L., K. Crowe, A. Fisher, O. Ghaznaw, R. McCoach, R. Narva, D. Schaulewicz, T. Sullivan, and T. White (2014): "An Overview of the U.S. Long-Term Care Insurance Market (Past and Present): The Economic Need for LTC Insurance, the History of LTC Regulation & Taxation and the Development of LTC Product Design Features," in *Managing the Impact of Long-Term Care Needs and Expense on Retirement Security Monograph*. The Society of Actuaries.

Stallard, E. (2011): "Estimates of the incidence, prevalence, duration, intensity, and cost of chronic disability among the US elderly," *North American Actuarial Journal*, 15(1), 32–58.

The American Association for Long-Term Care Insurance (2018): "Long-Term Care Insurance Facts—Statistics," Discussion paper, http://www.aaltci.org/long-term-care-insurance/learning-center/fast-facts.php.

Zeldes, S. P. (1989): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, 104(2), 275–298.

Appendices

A Estimated Inputs for the Model

A.1 Health and Mortality

Mapping Health States to Data. 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 $s = 1$).This 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.

The second set of questions is used to determine whether an individual is in the LTC/ADL state ($s = 2$). There are three 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 five 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 an intermediate measure: we categorize and individual as needing help with ADLs if they have difficulty with at least one ADL and they also receive outside help completing the ADL task. We choose this state definition since it is most consistent with the ADL definition presented in the VRI survey.[24] Since we are not using stays in a nursing home to represent our health state $s = 2$, not modeling spending when in need of care as an out-of-pocket expenditure shock, and not using a U.S. representative sample, our model-generated health and spending patterns will be different from those in Hurd, Michaud, and Rohwedder (2017) and Friedberg, Hou, Sun, Webb, and Li (2014).

Estimating the Health-State Transition Matrix. Using the health state definitions above, we estimate a sequence of health transition matrices conditional on a vector $x_{i,t}$ which includes individual i's age, t, and sex, 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.

[24] 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.

Figures A.1 and A.2 display the estimated health state transition probabilities ($\pi_g(s'|t, s)$).

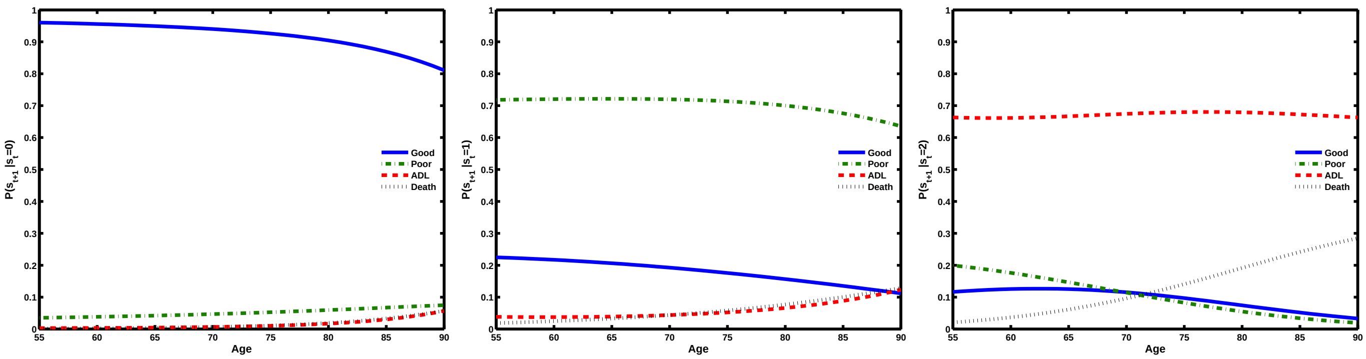

(a) Healthy **(b)** Sick **(c)** ADL

Figure A.1: Male Health State Transition Profile

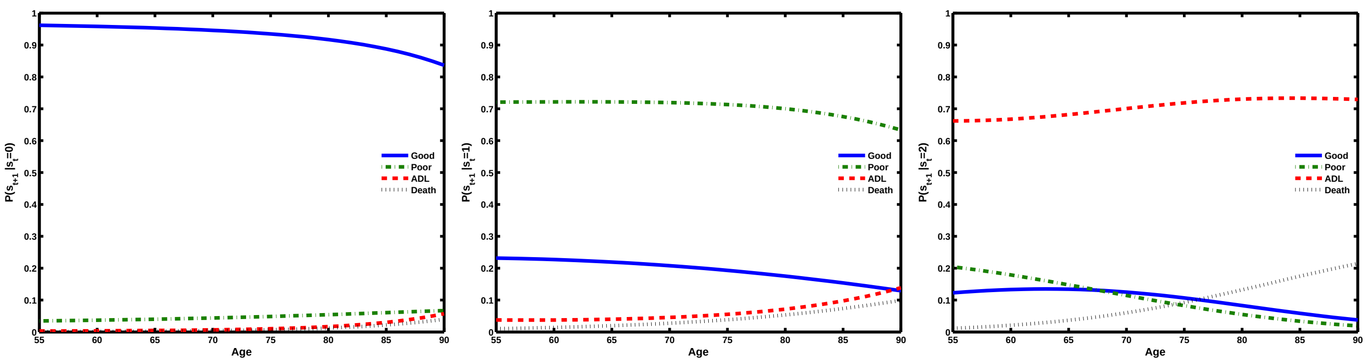

(a) Healthy **(b)** Sick **(c)** ADL

Figure A.2: Female Health State Transition Profile

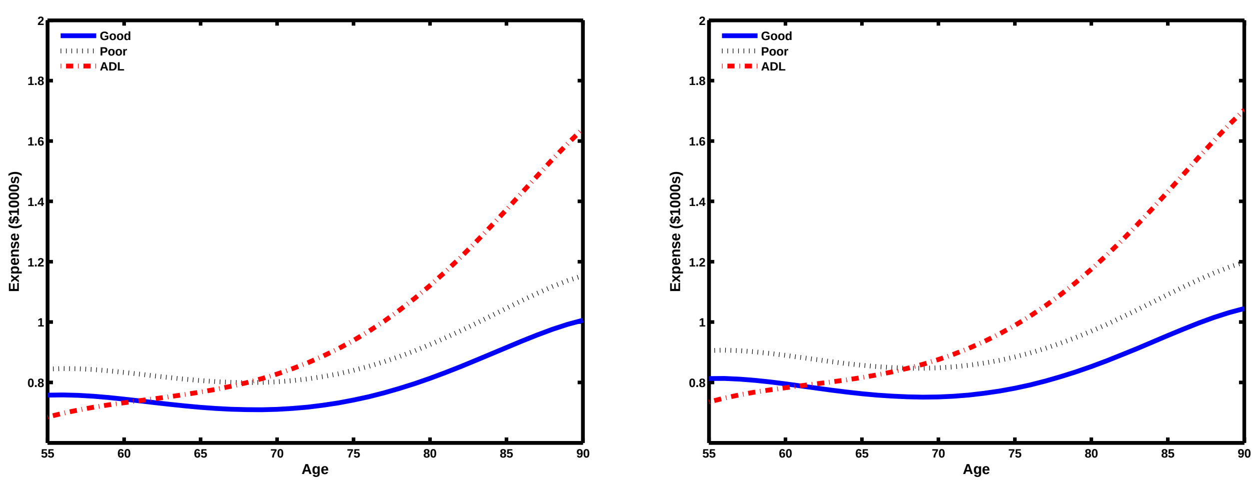

(a) Male **(b)** Female

Figure A.3: Median Health Cost Profile

Health Cost. To estimate the mean of the health cost distribution, $\mu_{med}(t, g, s)$, we regress log out-of-pocket medical expenditures on age, sex, 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, $\sigma^2_{med}(t, g, s)$. Discretizing the error term $\epsilon_t \sim N(0, 1)$ into separate health cost states determines the medical expense process.

Out-of-Pocket Health Cost Shocks. Figure A.3 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.

A.2 Income

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, sex, and the interaction of sex and age. To ensure that income is positive in all periods, we estimate a quantile regression of log income on these variables. Because we allow for 5 income profiles, the quantile regression is estimated for the 10^{th}, 30^{th}, 50^{th}, 70^{th}, and 90^{th} percentiles of the income distribution. We calibrate our income processes to the resulting estimates and group individuals into income profile quintiles.

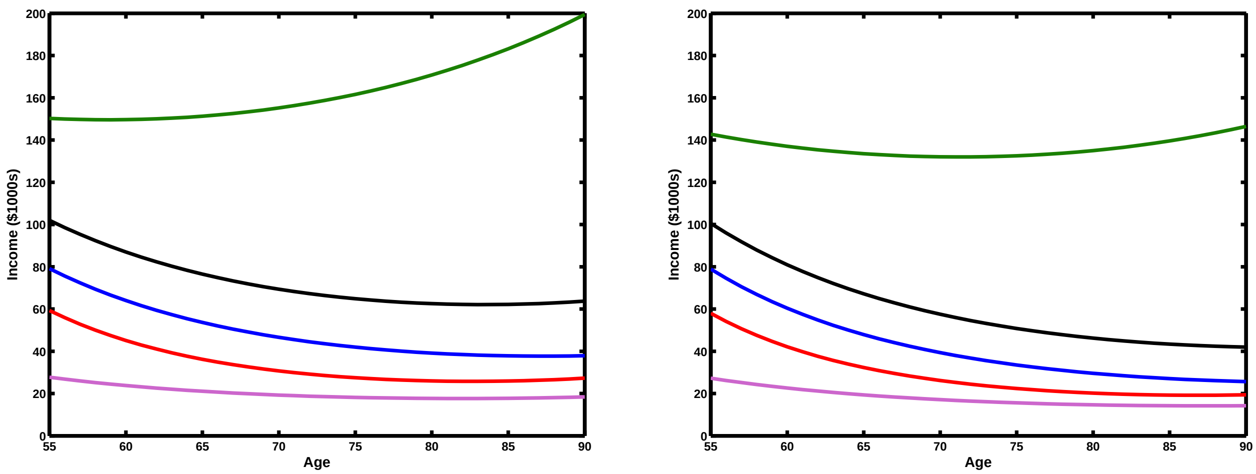

(a) Male **(b)** Female

Figure A.4: Income Profile Quintiles

B SSQ Design

The complete text for all SSQs are available in "VRI Technical Report: Strategic Survey Questions." An interactive demonstration of the SSQ survey instruments is available at `http://ebp-projects.isr.umich.edu/VRI/survey_2.html`. "VRI Technical Report: Strategic Survey Questions" also contains the math problems and first order conditions used to derive the likelihood function and estimate the parameters.

Since SSQs require respondents to comprehend and imagine complex scenarios, their design involved rich interaction with test respondents who were given cognitive interviews conducted by psychologists on the research team at the Survey Research Center at the University of Michigan. Furthermore, a sample of pilot survey respondents were interviewed via a scripted electronic real-time chat that was modeled after these cognitive interviews with input from survey experts from Vanguard and IPSOS. The resulting feedback led to numerous edits to improve the wording and flow of the surveys, and motivated us to test the comprehension of survey respondents.

Each SSQ begins with a broad introduction to the subject of interest and then presents the scenario. Immediately after the scenario is presented, respondents are provided with a summary of the rules that govern their choice. This recaps the previous screen but is presented in a bulleted, easy to read format. In addition, some features that were hinted at in the first screen, e.g., in SSQ 3 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.

To further reinforce details of the scenario and obtain a quantitative measure of understanding, we ask the respondents to answer a sequence of comprehension questions. For all SSQ questions, these comprehension questions are introduced with "Again for research purposes, it is important to verify your understanding. We will now ask you a series of questions (each question no more than 2 times). At the end we will give you the correct information for any questions which you haven't answered correctly just to make sure that everything is clear." 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 they fail to answer questions correctly a second time, they are presented with the correct answers. The questions asked verified the understanding of the ADL state, what the exact tradeoffs in that question were, which plan allocated resources to which state, what restrictions there are on the use of funds, and the nature of the claims process. Because respondents who make errors review the scenario between their first and second attempt, they get to reinforce those aspects they failed to understand the first time through before reporting their demand.

Having measured and reinforced understanding, we then ask respondents to make their final choice. For SSQs 2-4, we asked respondents to split their wealth between the two plans after again presenting them with the original scenario and including a link in the top right corner to the full scenario. The actual division of money involved a custom-designed interface that presents the trade off as clearly as possible. Specifically, we use an interactive slider that presents the payoffs in different states of the world. This payoff changes as the slider is moved, allowing respondents to observe how their choice is impacted by moving the slider. Text is included instructing the respondent how to allocate money by moving the slider, as well as what their allocation implies about resources available for different uses. The exact presentation for SSQ 3 can be seen in Figure B.1.

When the slider first appears, it does not have an allocation selected. It is only when respondents themselves click on the slider that any allocation is shown. To further dampen possible anchoring and status quo bias, we ask

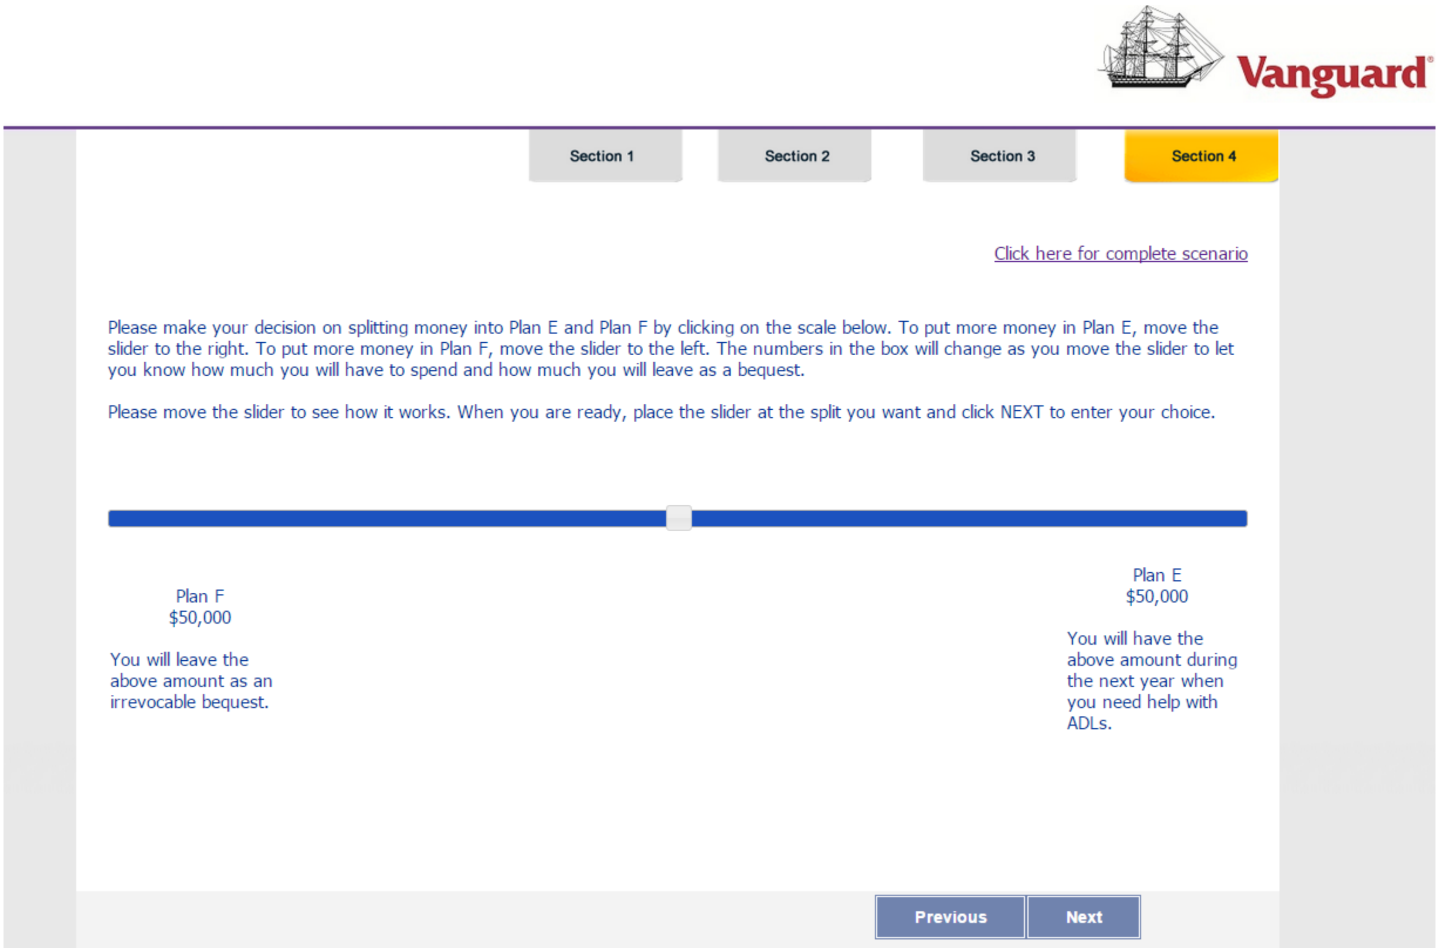

Figure B.1: SSQ Response Slider

respondents to move the slider at least once, which helps also to clarify the connection to the chosen allocation. [25] A key benefit of the slider is that it embodies the tradeoff and constraints of the choice problem, so that the respondent can experiment with them.

Having spent such a long time setting up the scenario and aiding comprehension, we often stayed within the scenario and asked respondents to make new choices with different scenario parameters. For example, in SSQ 3 answers were gathered not only concerning division of $100,000, but also of $150,000 and $200,000.

C SSQ Data

In this appendix we first present the raw data from the SSQ responses. Then, we present results from validation exercises for key survey instruments. We first consider SSQ credibility by reporting the outcomes of the comprehension tests. We then explore credibility by examining the internal coherence of the SSQ data using within-individual across-question response patterns. Then we examine how SSQ responses correlate with other variables, including demographic and financial measures. Finally, we examine similar correlations for stated ADLI and annuity demands. The SSQ response histograms are also available in “VRI Technical Report: Strategic Survey Questions.”

[25] Patterns of slider movement provide additional evidence of deliberation in the survey responses. To alleviate concern about anchoring effects for which individuals might settle immediately on their first chosen allocation, an analysis of click patterns shows that most respondents followed our suggestion and moved the slider before finalizing their choice. Regressions show that initial clicks have little predictive power for final answers, further suggestive of deliberation.

C.1 SSQ Response Histograms

In SSQs 1a and 1b, a response indicates how much income a respondent would be willing to risk. In SSQs 2a, 2b, 2c, 3a, 3b, a response indicates the amount of wealth allocated to the ADL state. In SSQ 4a a response indicates the wealth level at which a respondent is indifferent.

C.2 SSQ Data Analysis

In this section we expand upon three types of credibility analysis. First, we present results of objective and subjective comprehension tests. Second, we report responses to questions that were designed to assess how well the respondents felt they understood and internalized the SSQs. Finally, we analyze the internal coherence of responses. In short, our credibility analysis suggests that SSQ respondents understood the scenarios well and made meaningful choices, and that the SSQs are largely successful in providing measures of preferences where respondents abstract from their situations.

C.2.1 Objective Measures of SSQ Comprehension

In creating the surveys, we explicitly designed measures that permit analysis of response quality. For example, after introducing each SSQ's scenario, respondents answered questions checking comprehension of details of the scenario. In addition to reinforcing question specifics, these tests provide quantitative measures of respondent understanding. Respondents are asked to answer questions on the comprehension test at most twice. Performance on these comprehension tests is summarized in Table C.1. For SSQ 1 there were 6 comprehension questions. About 50 percent of respondents answered all questions correctly on their first attempt, with 75 percent doing so after their second attempt, and more than 90 percent making at most one error after the second attempt. For SSQ 2, there were 9 comprehension questions. Although only 20% of respondents answered all questions correctly on the first attempts, more than 55% answer all correctly and more than 80% miss at most one on the second attempt. For SSQs 3 and 4, most respondents answer all questions correctly on the first attempt and nearly all respondents miss at most one on the second attempt, albeit for fewer comprehension questions. Thus understanding of scenario details appears high, and is in practice likely even higher than tests indicate because answers to missed questions and important aspects of the scenario are reiterated before and while respondents make their final decisions.

	SSQ 1	SSQ 2	SSQ 3	SSQ 4
Number of questions	**6**	**9**	**3**	**2**
Percent all correct, 1^{st} try	46.2	18.5	55.3	77.3
Percent all correct, 2^{nd} try	75.1	55.4	81.9	94.1
Percent ≤ 1 wrong, 2^{nd} try	93.4	80.8	96.1	99.5

Table C.1: SSQ Comprehension Questions: When introducing each survey instrument, we asked a series of test questions that examined respondents knowledge of and reinforced details of each scenario. Statistics on the number of correct responses are presented in the table.

C.2.2 Subjective Measures of SSQ Comprehension

As part of the survey design process, we gathered feedback from scripted live chat pop-up interviews with a subset of the pilot sample. Additionally, a subset of the live chat questions were posed to the full production sample at the end

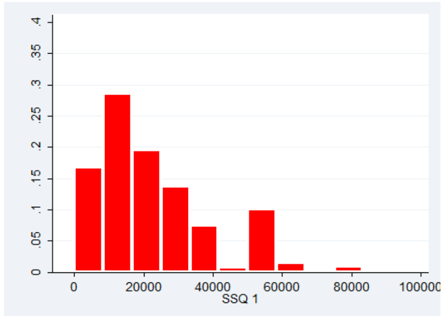

SSQ 1a: Lottery over Spending ($100K)

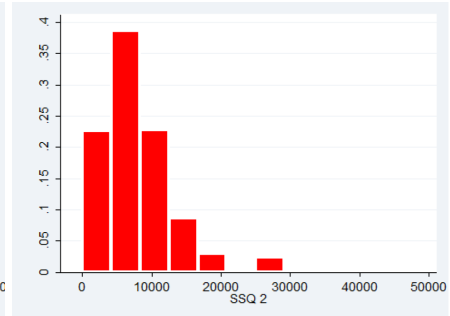

SSQ 1b: Lottery over Spending ($50K)

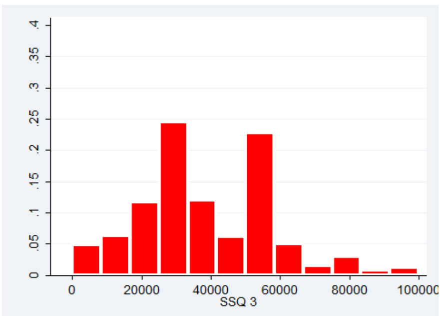

SSQ 2a: Ordinary vs. ADL
($100K $\pi = 0.75$)

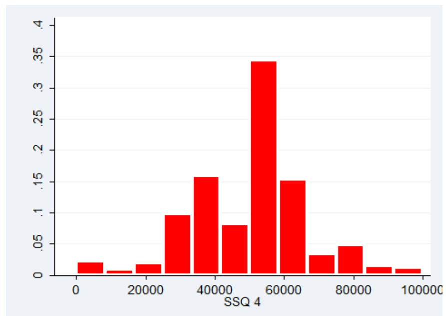

SSQ 2b: Ordinary vs. ADL
($100K $\pi = 0.50$)

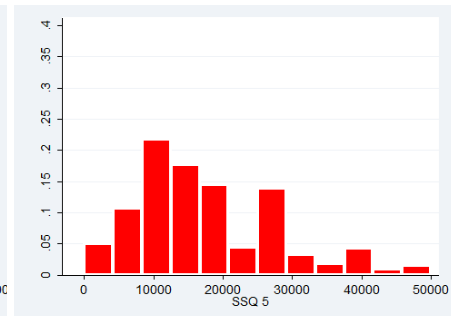

SSQ 2c: Ordinary vs. ADL
($50K $\pi = 0.75$)

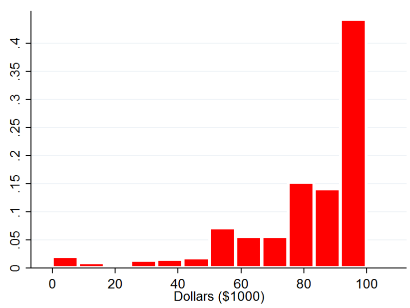

SSQ 3a: ADL vs. Bequests ($100K)

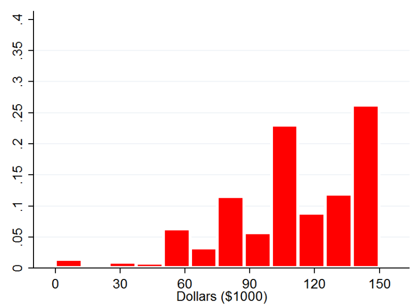

SSQ 3b: ADL vs. Bequests ($150K)

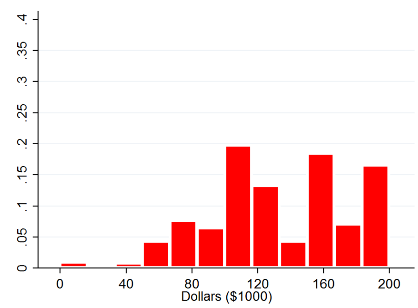

SSQ 3c: ADL vs. Bequests ($200K)

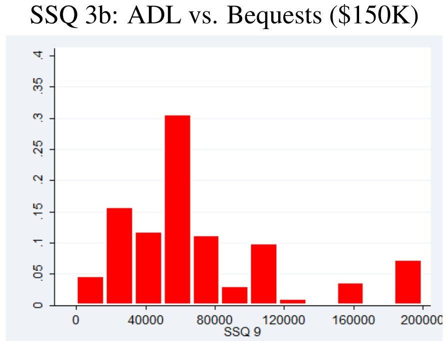

SSQ 4a: Value of Public Care

Figure C.2: Distribution of Responses to SSQs

of the survey. The full text and tabulated answers to these questions are included in Table C.2. We see that nearly 90 percent of respondents found the tradeoffs either very clear or somewhat clear, while only 1 percent reported finding the tradeoffs very unclear. Furthermore, more than 80 percent indicated that they were able to place themselves in the

Overall, how clear were the tradeoffs that the hypothetical scenarios asked you to consider?		Overall, how well were you able to place yourself in the hypothetical scenarios and answer these questions?		How much thought had you given to the issues that the hypothetical scenarios highlighted before taking the survey?	
Response	Percent	Response	Percent	Response	Percent
Very Clear	51.7	Very Well	23.0	A lot of thought	29.5
Somewhat Clear	39.7	Moderately Well	60.5	A little thought	52.0
Somewhat Unclear	7.4	Not very well	14.2	No thought	18.4
Very Unclear	1.1	Not very well at all	2.2		

Table C.2: General SSQ Comprehension Questions: Each respondent was asked each of the three questions presented in the table. This table provides the distribution of responses.

hypothetical scenario either moderately or very well, with only 2 percent reporting they were "not very well at all" able to place themselves in the scenario. There is also a significant and interesting difference, with evidence that it was harder to place oneself in the scenario and answer from that perspective than it was to comprehend the question. This difference in difficulty is consistent with our prior, and is suggestive of how seriously respondents took their charge. Finally, 82 percent had given the underlying issues at least a little thought before taking the survey, with only 18 percent having given no or very little thought to the relevant issues. These responses indicate a clear understanding of the tradeoffs, and an ability and willingness to think hypothetically. Overall, the quality checks implemented in the surveys indicate that respondents understood the scenarios and gave responses reflecting choices they would likely make if they were in the described situations.

C.2.3 Coherence

As Manski (2004) stresses, one necessary criterion for judging responses as meaningful is internal coherence, i.e., responses should not be self-contradictory across questions. One indication of internal coherence derives from analyzing the pattern of correlations in survey responses. SSQ 1, SSQ 2, and SSQ 3 were each asked to all correspondents with several variants, using the same scenario with different scenario parameters. Internal coherence would require a strong positive correlation in responses for each individual within each scenario across scenario parameterizations. Just such a pattern is present in the diagonal blocks of the correlation matrix presented in Table C.3.

	SSQ 1a	SSQ 1b	SSQ 2a	SSQ 2b	SSQ 2c	SSQ 3a	SSQ 3b	SSQ 3c	SSQ 4a
SSQ 1a	1.00								
SSQ 1b	**.44**	1.00							
SSQ 2a	-.01	.04	1.00						
SSQ 2b	-.04	-.01	**.61**	1.00					
SSQ 2c	-.08	.07	**.55**	**.56**	1.00				
SSQ 3a	-.01	-.08	-.11	-.04	-.11	1.00			
SSQ 3b	-.06	-.08	.04	.04	.02	**.78**	1.00		
SSQ 3c	-.08	-.08	.07	.08	.07	**.63**	**.86**	1.00	
SSQ 4	-.04	.00	.04	.05	.05	-.15	-.13	-.10	1.00

Table C.3: Correlation Matrix of SSQ responses: The correlation matrix for the SSQ responses are presented above. Correlations between SSQs of the same type are in bold.

C.3 SSQ Correlates

SSQs are designed to be invariant to the situation of the respondent, so we would not expect to see significant predictive power of demographics and economic covariates. Appendix Tables C.4 through C.7 report regressions of demographic and economics covariates. Indeed, we observe little significance in coefficients on age, income, health, and wealth, suggesting that this design is successful. There are some differences by sex, which is not inconsistent with the validity of the SSQs, since the SSQs did not ask people to respond from the perspective of a hypothetical sex.

	SSQ 1a	SSQ 1b
No equity	-2890.87	-2518.79
	(3528.23)	(1437.26)
Age	-2160.43	-1449.8
	(9093.02)	(3692.05)
Age^2	35.75	19.51
	(127.17)	(51.64)
Age^3	-0.19	-0.09
	(.59)	(.24)
Health: Poor	3396.57	-770.19
	(2645.17)	(1073.20)
Health: ADL	-7880.16	-6274.23
	(5801.23)	(2491.90)
Income Quintile: 2	651.91	-760.05
	(1733.43)	(703.60)
Income Quintile: 3	-1484.81	-239.31
	(1793.88)	(725.39)
Income Quintile: 4	-2356.25	-1866.28
	(1752.99)	(711.14)
Income Quintile: 5	-215.54	-1232.80
	(1839.01)	(745.27)
Female	4003.05	1260.96
	(1090.83)	(442.36)
College or Higher	364.78	-412.27
	(1288.13)	(521.99)
Log(Wealth)	-101.86	281.22
	(534.23)	(216.91)
N	1,086	1,086

Table C.4: Correlates of SSQs 1: This table presents the results from a Tobit regression of SSQ 1 responses on the listed covariates.

	SSQ 2a	SSQ 2b	SSQ 2c
Predicted Average Cost of ADL Care	-.01	.01	.00
	(.01)	(.01)	(.01)
Family Care Probability	-47.39	-27.11	-11.71
	(23.67)	(20.47)	(13.03)
Own Private LTCI	1413.97	1319.49	1559.72
	(1473.88)	(1292.41)	(822.26)
Public LTC Facility vs Private Ranking (1-5)	1105.05	1153.45	149.96
	(782.12)	(685.61)	(436.23)
Subj. Prob of ADL need for 1 year	-590.72	569.13	250.83
	(1180.95)	(1035.46)	(658.77)
Age	-4419.66	11786.23	4329.54
	(11916.13)	(10448.18)	(6648.32)
Age^2	62.59	-241.13	-56.46
	(169.27)	(148.42)	(94.44)
Age^3	-.27	1.09	.25
	(.79)	(.70)	(.44)
Health: Poor	23475.37	-579.42	-3628.50
	(2874.03)	(2521.13)	(1604.86)
Health: ADL	-45859.72	186.35	-4914.70
	(6235.94)	(5434.30)	(3452.86)
Income Quintile: 2	-1604.81	-557.64	-1023.98
	(1870.98)	(1640.46)	(104390.00)
Income Quintile: 3	-415.03	-1770.96	-1043.34
	(1937.79)	(1699.07)	(1080.37)
Income Quintile: 4	-4021.63	-2918.75	-1648.65
	(1903.79)	(1668.65)	(1061.68)
Income Quintile: 5	-4766.32	-2209.07	-951.83
	(1993.76)	(1747.56)	(1111.47)
Female	-359.41	-651.06	1383.57
	(119.44)	(1046.27)	(665.57)
College or Higher	61.68	1201.94	-722.12
	(1398.27)	(1225.56)	(779.56)
Log(Wealth)	-1185.08	-1069.65	-624.44
	(574.72)	(503.92)	(320.44)
N	1,086	1,086	1,086

Table C.5: Correlates of SSQs 2: This table presents the results from a Tobit regression of SSQ 2 responses on the listed covariates. Missing observations related to Family Care Probability, Predicted Average Cost of ADL Care, and the Subjective Probability of Needing help with ADLs coming from attrition between Survey 2 and 3 are addressed via dummy variables for missing observations.

	SSQ 3a	SSQ 3b	SSQ 3c
Predicted Average Cost of ADL Care	0.05	0.07	0.08
	(0.02)	(0.02)	(0.03)
Family Care Probability	-97.42	-143.80	-210.50
	(40.52)	(48.08)	(60.37)
Total Transfers to Descendants in last 3 years	-0.06	-0.09	-0.13
	(0.03)	(0.04)	(0.04)
Public LTC Facility vs Private Ranking (1-5)	-2,315	269.94	1,652
	(1,379)	(1,630)	(2,047)
Age	20,542	12,766	22,308
	(17,512)	(20,409)	(25,600)
Age^2	-279.27	-180.33	-304.29
	(244.85)	(285.24)	(357.83)
Age^3	1.23	0.80	1.32
	(1.13)	(1.32)	(1.65)
Health: Poor	2,999	1,970	5,370
	(5,009)	(5,943)	(7,467)
Health: ADL	24,130	3,624	21,009
	(11,786)	(13,135)	(16,191)
Income Quintile: 2	4,036	4,382	369.10
	(3,302)	(3,900)	(4,886)
Income Quintile: 3	170.29	-830.55	-484.33
	(3,396)	(4,018)	(5,056)
Income Quintile: 4	-91.25	3,411	-356.78
	(3,338)	(3,961)	(4,957)
Income Quintile: 5	1,316	4,537	4,542
	(3,560)	(4,214)	(5,280)
Female	134.23	-1,637	-2,644
	(2,078)	(2,455)	(3,071)
College or Higher	7,136	6,420	4,868
	(2,435)	(2,895)	(3,631)
log(Wealth)	976.54	985.06	459.47
	(1,040)	(1,222)	(1,536)
N	1,086	1,086	1,086

Table C.6: Correlates of SSQs 3: This table presents the results from a Tobit regression of SSQ 3 responses on demographic variables and the listed covariates.

	SSQ 4a
Predicted Average Cost of ADL Care	.01
	(.01)
Family Care Probability	-23.85
	(20.62)
Own Private LTCI	1306.05
	(1291.73)
Total Transfers to Descendants in last 3 years	-3.5e-3
	(.02)
Public LTC Facility vs Private Ranking (1-5)	(1170.57)
	(684.31)
Subj. Prob of ADL need for 1 year	573.83
	(1034.31)
Age	17557.61
	(10435.53)
Age^2	-237.98
	(148.24)
Age^3	1.07
	(.70)
Health: Poor	-565.05
	(2516.65)
Health: Poor	228.07
	(5425.09)
Income Quintile: 2	-599.13
	(1637.37)
Income Quintile: 3	-1765.51
	(1696.82)
Income Quintile: 4	-2897.08
	(1669.79)
Income Quintile: 5	-2146.97
	(1773.72)
Female	-658.80
	(1044.65)
College or Higher	1215.54
	(1223.56)
Log(Wealth)	-1046.09
	(506.97)
N	1086

Table C.7: Correlates of SSQ 4: This table presents the results from a Tobit regression of SSQ 4 responses on the listed covariates. Missing observations related to Family Care Probability, Predicted Average Cost of ADL Care, and the Subjective Probability of Needing help with ADLs coming from attrition between Survey 2 and 3 are addressed via dummy variables for missing observations.

D Stated Demand Correlates

To check whether stated preferences for insurance products reported in the survey are consistent with behaviors and beliefs outside of the survey we regress the extensive and intensive margins of stated ADLI and annuity demand on a host of covariates. We again find evidence that our survey measures align with behavior in meaningful ways.

Stated ADLI Demand Table D.8 presents results from regressions of demographic and other covariates on stated demand for ADLI. The first column presents a probit regression of demographic and other covariates on an indicator equal to one if the respondent reported they would purchase a positive amount of ADLI. We observe that respondents that report higher probabilities of experiencing extended time in the ADL state are more likely to purchase ADLI, and that individuals that indicate a more favorable opinion of publicly provided LTC are less likely to purchase ADLI. In the second column we present an OLS regression of the amount of ADL-contingent annual income that respondents state they would purchase in the subsample of respondents that reported they would purchase positive amounts. Here we observe that those that own LTC insurance and those that predict higher average LTC costs purchase more, while those that report a more favorable opinion of publicly provided LTC purchase less.

Both measures of stated interest correlate in generally reasonable manners with economic and demographic characteristics.

	$\mathbb{I}_{ADLI>0}$	**Annual ADLI Payout**
Owns LTCI Indicator	0.09	6,872
	(0.11)	(4,023)
Predicted Average Cost of ADL Care	1.28e-7	.07
	(9.13e-7)	(0.04)
Family Care Probability	0.002	-39.32
	(0.002)	(64.42)
Total Transfers to Descendants in last 3 years ($1000s)	-2.96e-6	.07
	(1.42e-6)	(.06)
Public LTC Facility vs Private Ranking (1-5)	-0.10	-5,565
	(0.06)	(2,331)
Subj. Prob. of Help with ADLs for 1 year	0.17	-340.29
(Above Median)	(0.09)	(3,142)
Age	-0.52	66,838
	(0.81)	(31,951)
Age^2	0.01	-955.8
	(0.01)	(452.2)
Age^3	-0.00003	4.50
	(0.00005)	(2.11)
Health: Poor	-0.12	-95.9
	(0.22)	(7,867)
Health: ADL	-0.63	17,751
	(0.58)	(27,067)
Income Quintile: 2	-0.30	-2,423
	(0.14)	(5,227)
Income Quintile: 3	-0.03	-4,213
	(0.14)	(5,000)
Income Quintile: 4	-0.17	-10,888
	(0.14)	(5,074)
Income Quintile: 5	-0.10	-650.5
	(0.15)	(5,449)
Female	0.16	16,590
	(0.09)	(3,246)
College or Higher	-0.05	28.45
	(0.10)	(3,670)
Log(Wealth)	0.05	654
	(0.05)	(1,870)
N	750	225

Table D.8: Correlates of Surveyed ADL demand measurement: This table presents how stated ADLI demand is predicted by covariates. Column 1 presents the results of a probit regression of the ADLI purchase decisions, and Column 2 presents an OLS regression on the level of ADLI annual payout demanded for those with positive demand.

E Robustness

	%>0	Mean	p5	p10	p25	p50	p75	p90	p95
Baseline	58	67,322	8,845	13,108	29,070	54,768	93,449	148,143	180,556
Different Model Specifications									
Homogeneous Parameters	75	71,732	46,224	48,807	52,462	60,012	84,084	123,398	137,228
10% load	56	66,301	7,930	12,720	30,118	53,259	88,377	149,928	175,342
20% load	54	63,000	7,255	13,227	29,734	51,417	83,496	133,913	165,648
30% load	54	59,954	6,773	12,505	28,395	50,417	81,979	126,423	155,934
r=.03	54	65,833	8,094	14,623	29,168	53,975	89,073	141,207	178,275
Housing Wealth	62	77,581	9,683	16,677	34,923	62,770	103,619	166,392	214,944
No Min. Expenditure	54	70,770	7,946	12,891	33,537	58,613	94,365	164,076	181,407
Multiplicative Errors	72	56,590	8,551	16,648	28,190	43,276	74,598	107,812	153,489
Different Subsamples									
Employer Subsample	49	50,285	8,015	8,885	20,609	41,617	70,466	102,391	141,093
HRS weights	46	47,751	6,043	6,812	10,833	34,246	65,437	107,986	150,038
Home Owners	58	69,187	10,137	15,651	31,262	57,353	97,129	146,690	177,729
LTCI Owners	63	71,520	9,612	18,826	26,057	57,610	109,535	149,846	191,177

Table E.1: Robustness of ADLI Quantity Demanded: This table presents ADLI demands for various specifications and subsamples. Demand measures are for the subsample of the population that does not own any private LTCI aside from the LTCI Owners alternative sample.

	Wealth and Income					
	Mean	**10p**	**25p**	**50p**	**75p**	**90p**
Wealth	540,510	52,473	168,150	392,926	836,400	1,161,000
Income	77,887	37,500	50,000	72,065	104,000	130,000

	Demographics							
	Age			**Health**			**Sex**	
	55-64	**65-74**	**75+**	**Good**	**Poor**	**ADL**	**Male**	**Female**
N=162	68.5%	28.4%	3.1%	95.7%	3.1%	1.2%	45.1%	54.9%

Table E.2: Summary Statistics on Wealth, Income, Health, Age, and Sex—Employer sample: This table presents the marginal distributions of wealth, income, and demographic characteristics of the subsample of respondents that have an employer sponsored Vanguard account.

ADLI Demand Function for LTCI Owners.

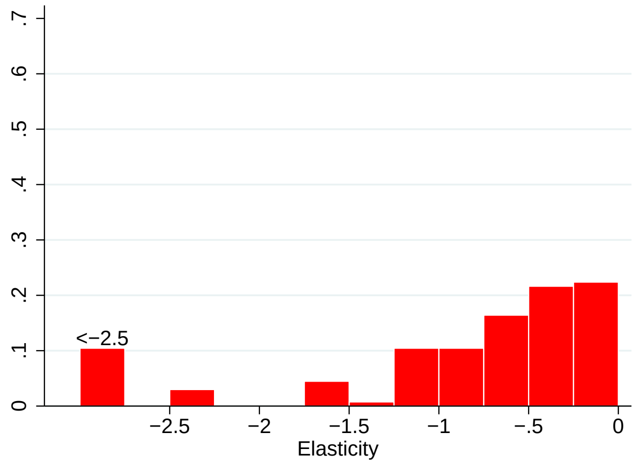

Figure E.1: Distribution of the Price Elasticity of Demand: This figure presents the histogram of the elasticity of demand with respect to price for those who own LTCI. It plots the distribution of the percent change in demand to a one percent increase in price, local to the optimal demand level and given price.

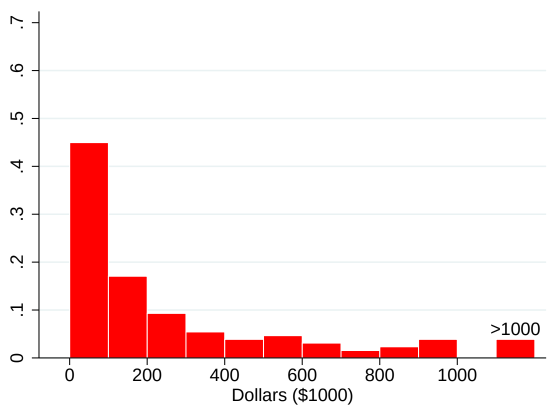

(a) Distribution of Consumer Surplus

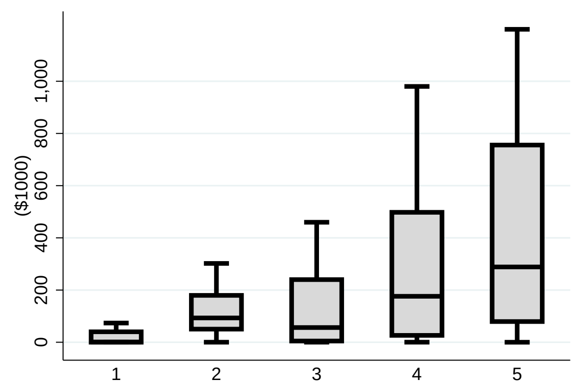

(b) Consumer Surplus Box Plot by Wealth Quintile

Figure E.2: Consumer Surplus for LTCI owners: This figure presents consumer surplus measures for the subsample who owns LTCI. The left panel presents the histogram of consumer surplus. Consumer surplus is the maximum people are willing to pay to purchase their desired amount of insurance above the price they actually paid. The right panel presents a box plot of the consumer surplus by wealth quintile.

F Exploring Modeled and Stated Demand Difference

In this section we develop an econometric method that utilizes the difference between modeled and stated demand to identify sources of model misspecification. Define η_i as the difference between model and stated demand for individual i:

$$\eta_i = \text{Modeled}_i - \text{Stated}_i.$$

Assume that this difference η_i can generally be expressed as a function of modeled state variables x_i, preference parameters Θ_i, and other, undetermined state variables q_i. Thus,

$$\eta = G(x, \Theta, q).$$

G is thus a generic function of our demand measurement error that allows for differences in demand measures from two distinct sources. First, differences in demand measurements could be caused by misspecification of included model elements as dictated by Θ and x. For example, misspecification of the functional form of preferences could cause systematic variation in η_i as a function of Θ, while use of incorrect health transition probabilities (which we model only as a function of x) could cause η_i to be dependent upon included state variables sex and age. A second cause of differences in demand measurement could be omission of relevant state variables q from our modeled demands. For example, the model in this paper does not consider the effect of children and family on the saving and insurance purchase decisions. Similarly, private information about individuals' health is omitted from the model but presumably affects stated demand.

Each of these variable sets could affect both measures of demand. Preferences Θ and x are the factors that are modeled, reflecting opinions of the model-builders that they are the relevant variables in stated insurance purchase decisions. Omitted variables q could affect decisions two ways. First, in recovering parameters Θ, SSQ responses are interpreted as being determined by a limited number of factors. Omission of these factors from the model could impact this interpretation and thus affect modeled demand. In addition, stated demand is possibly affected by factors that are not considered in the model. Given that most factors affect both demand measures simultaneously, it is difficult to determine exactly how each will affect the difference between modeled and stated demand. In general, however, one would expect omitted variables that discourage purchase of insurance products to be associated with lower model differences. Similarly, model misspecification that encourages demand for insurance products might be associated with larger differences in demand measures. Thus, omitted risks that encourage precautionary holding of liquid wealth should correspond to larger demand differences, while overstated insurable risks should correspond to smaller differences in demand measures.

Returning to the model of demand differences, we assume that G can be approximated as

$$G(x, \Theta, q) \approx g_x(x) + g_\Theta(\Theta) + g_q(q). \tag{F.1}$$

This decomposition assumes that there is no effect on demand differences due to the interaction between modeled state variables x, estimated parameter set Θ, and omitted variables q. It is thus a first order approximation to the function of interest. The separability of effects of state variable and parameter sets is primarily necessary for tractability. Further examination of this assumption does not appear to change our fundamental conclusions. The separability of omitted variables q and parameter sets Θ or state variables x likely weakens the closeness of our approximation. Given that we are primarily interested in identifying the presence of omitted factors q and not the quantitative effect however, this assumption should not be restrictive. It is only restrictive if the omitted variable q only affects the difference in demand measurements through its interaction with state variables x and Θ.

The assumptions of additive separability provide convenient interpretation. For each function g, $g \neq 0$ implies model misspecification (relative to stated demands) related to the relevant variables. Thus, $g_x(x) \neq 0$ suggests model misspecification related to modeled state variables, $g_\Theta(\Theta) \neq 0$ suggests model misspecification related to preference parameters, and $g_q(q) \neq 0$ suggests model misspecification related to omitted variables q. Furthermore, $g > 0$ suggests misspecification that causes model demand to be overstated relative to stated demand, while $g < 0$ suggest misspecification that causes model demand to be understated relative to stated demand. To estimate this function, we take a non-parametric approach that does not assume any functional form for g_Θ and g_x . Specifically, partition the space of feasible Θ and x into $\mathcal{P}^\Theta = \{P_k^\Theta\}_{k=1}^{K^\Theta}$ and $\mathcal{P}^x = \{P_k^x\}_{k=1}^{K^x}$ respectively. Using these partitions, define vectors $C_i^\Theta \ni \{C_{i,k}^\Theta = 1 \iff \Theta_i \in P_k^\Theta\}$ and $C_i^x \ni \{C_{i,k}^x = 1 \iff x_i \in P_k^x\}$. Finally, defining vectors $\beta_k^\Theta = g(\Theta)$ for any $\Theta \in P_k^\Theta$ and $\beta_k^x = g(x)$ for any $x \in P_k^x$, the functions of interest

$$g_\Theta(\Theta_i) = \beta^\Theta C_i^\Theta$$
$$g_x(x_i) = \beta^x C_i^x$$

are approximated to arbitrary precision for sufficiently fine partitions. Finally, model-omitted variables q are examined one at a time. Given primary interest in the significance and sign of $g(q)$, we approximate g_q with a linear function, such that $g_q(q) = \Gamma q$. Substituting these expressions into equation F.1 yields

$$G(x, \Theta, q) = \beta^\Theta C_i^\Theta + \beta^x C_i^x + \Gamma q_i, \tag{F.2}$$

which we use to estimate

$$\eta_i = \beta^\Theta C_i^\Theta + \beta^x C_i^x + \Gamma q_i + \epsilon_i. \tag{F.3}$$

Equation F.3 permits testing of the null hypothesis $H_0 : \ \beta^\Theta = 0; \beta^x = 0; \Gamma = 0$. Rejection of the null hypothesis for β^Θ or β^x suggests that the existing state variables included in our structural model are not incorporated in a way that fully reflects their impact on demand.[26] Similarly, a positive coefficient on Γ indicates that the variables in q cause the model to overpredict demand, while a negative coefficient on Γ indicates that the variables in q cause the model to underpredict demand. It is thus reasonable to expect any variables that reflect missing risks or savings motives that are not included in our model to be estimated to have a significant positive coefficient.

To implement this estimation, we must first construct our partitions $\mathcal{P}^\Theta$ and $\mathcal{P}^x$. $\mathcal{P}^x$ is constructed according to the discrete value of all state variables except wealth. Because wealth is continuous, we discretize it according to \$50,000 bins up to \$1,000,000, and \$200,000 bins thereafter. $\mathcal{P}^\Theta$ is a partition of continuous valued parameters. We discretize this by sorting individuals into partitions according to whether each parameter is above or below the population median.

[26]As mentioned when discussing equation F.1, the above specification does not control for effects of the interaction between preferences and modeled state variables. Attempts to control for these effects through inclusion of first order cross-partials of Θ_i and x_i weakens precision of estimates but does not impact significance of other coefficients.