NBER WORKING PAPER SERIES

LATE-IN-LIFE RISKS AND THE UNDER-INSURANCE PUZZLE

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

Working Paper 22726
http://www.nber.org/papers/w22726

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2016

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 Mariacristina De Nardi, Rebecca Diamond, Sebastian Di Tella, Eric French, Vìctor Rìos-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/. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Late-in-Life Risks and the Under-Insurance Puzzle
John Ameriks, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti
NBER Working Paper No. 22726
October 2016
JEL No. D14,D91,E21,G22,H31,I13,J14

ABSTRACT

Individuals face significant late-in-life risks, including needing long-term care (LTC). Yet, they hold little long-term care insurance (LTCI). Using both "strategic survey questions," which identify preferences, and stated demand questions, this paper investigates the degree to which a fundamental lack of interest and poor product features determine low LTCI holdings. It estimates a rich set of individual-level preferences and uses a life-cycle model to predict insurance demand, finding that better insurance would be far more widely held than are products in the market. Comparing stated and model-predicted demand shows that flaws in existing products provide a significant, but partial, explanation for this under-insurance puzzle.

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

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

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

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

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

The Vanguard Research Initiative (VRI) website is available at
http://ebp-projects.isr.umich.edu/VRI
Vanguard Research Initiative Technical Report: Long-term Care Model is available at
http://ebp-projects.isr.umich.edu/VRI/papers/VRI-TechReport-LTC-Model.pdf
Vanguard Research Initiative Technical Report: Long-term Care Strategic Survey Questions is available at
http://ebp-projects.isr.umich.edu/VRI/papers/VRI-TechReport-LTC-SSQs.pdf

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 well-located and 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). The resulting $300,000 (three years times $100,000) to self-insure against this risk is 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]

There is substantial uncertainty about the determinants of low observed LTCI holdings. It could be that there is a fundamental lack of interest in insuring against this health realization. There could be, however, substantial unmet demand, with low LTCI ownership reflecting poor quality insurance products. There are many features of the currently available products that might be limiting demand, such as default risk, premium risk, inflation risk, or the uncertain and potentially adversarial claims process based on expense reimbursement. 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.

We investigate the motives generating late-in-life insurance demand in the Vanguard Research Initiative (VRI), a newly-created panel study that includes batteries of questions we designed for this purpose. A key element of our analysis is the estimation of a rich set of person-specific preferences related to risk, bequests, and long-term care. For each respondent—who differ in preference, financial, and demographic variables—we use a state-of-the-art life-cycle model to predict demand for insurance products. To assess peoples' interest in insuring LTC need, abstracting from the products sold in the market, we study a simple insurance product that does not have these negative properties. Specifically, we model demand for "Activities of Daily Living" insurance (ADLI), a state-contingent asset that delivers wealth precisely when individuals have difficulties with ADLs and may therefore need long-term care.

Our main finding is that about 60 percent of panel members would purchase ADLI, about three times the number that actually hold LTCI, which is 22 percent in the VRI. The estimated intensive-margin demands are also high for those who purchase, averaging a $67,000 payout in any year in which help is needed. The gap between modeled demand and actual holdings is present across the wealth and income distribution, survives in more representative subsamples, and is robust to reasonable loads, alternative model specifications, and different estimation strategies. Furthermore, modeled demand is relatively price-inelastic and consumer surplus is often high. We conclude that there is a significant "LTCI puzzle" in the form of a gap between actual holdings and estimated holdings in reasonably calibrated models that match many dimensions of heterogeneity. We also find the annuity puzzle in our sample (see Yaari (1965), Modigliani (1986), Brown (2007)), with modeled demand far in excess of holdings. Hence, under-insurance of late-in-life risks appears to be pervasive.

To what extent does the LTCI puzzle result from the yawning chasm between the better insurance product that we model and those available in the market? To assess the extent to which poor features of available LTCI products explain the puzzle, we posed stated demand questions in the VRI in which people directly report how much ADLI they would purchase given the description of the asset and a price. Hence, the VRI has multiple observations on demand for insurance against LTC risk: modeled ADLI demand, stated ADLI demand, and actual LTCI purchases. The modeled ADLI demand is based on a precisely-specified product in a precisely-specified circumstance. The stated

[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.

ADLI demand is based on the identical product, but for the individual's actual circumstance. Differences between them thus can be attributed to differences between modeled and actual circumstance (e.g., availability of care from family members), but not taste or product characteristics. The difference between model and stated ADLI demand represents the remainder of the LTCI puzzle after controlling for insurance product features. The actual LTCI purchase reflects an individual's circumstance plus the characteristics of LTCI products available to them. Hence, the difference between actual LTCI purchase and stated ADLI demand identifies the role of insurance products' availability and quality.

We find that 31 percent of respondents have positive stated demand for ADLI. Moreover, 44 percent of respondents express a desire to insure the need for LTC by either owning LTCI or stating positive ADLI demand. This finding closes about half of the gap between model predicted ownership of ADLI and actual LTCI ownership. The agreement between the model-predicted and stated demand measures indicates substantially more interest in ADLI than LTCI. Therefore, poor features of available LTCI substantially contribute to low insurance against LTC risk. Furthermore, the workhorse model and estimated preferences capture well the substantial latent demand for insuring late-in-life health risks.

This LTCI puzzle stands in stark contrast to the annuity puzzle, for which the model grossly over-predicts the fraction of the population that would demand annuities and how much they would demand. This finding holds when comparing either to actual annuity ownership or stated demand, although more people state demand for the high quality annuity product than own an annuity. There seems to be something fundamental to the lack of annuity demand beyond distaste for particulars of the annuity products available in the market. Hence, our measurements suggest that improvements in LTCI products have more potential to generate demand for purchasing insurance than improvements in the annuity market.

Even in the case of LTCI, this better ADLI product does not fully resolve the puzzle as people state lower interest than our model implies. We close the paper by analyzing the gap between modeled and stated ADLI demand. First, we show that the gap is concentrated among high wealth individuals. We then identify variables that predict this gap related to the structure of the model, survey comprehension, and private information about health.

Our analysis rests critically on purpose-designed survey instruments. Particularly central are "strategic survey questions" (SSQs) that are crafted to identify key preference parameters.[2] Given their centrality, we dedicate significant effort to verifying the quality of these SSQs. As part of a broader process of quality control, we posed a series of comprehension tests before allowing respondents to answer these questions. These were generally very well answered. As a result of this and other tests of reason, we have confidence that they were answered with deliberation and honest purpose. Moreover, stepping back from the particular structure of the model, the SSQ responses give a coherent account of desire to insure against LTC risk. Most VRI respondents indicated a desire to spend on themselves when they need long-term care even at the expense of leaving a smaller bequest. Reliance on government provision of long-term care via Medicaid is not seen as an attractive alternative to private provision (as discussed by Pauly (1990) and similar to Hubbard, Skinner, and Zeldes (1994)) and bequest motives do not overwhelm precautionary motives (Lockwood (2016) and Koijen, Van Nieuwerburgh, and Yogo (2016)).

The introduction concludes with a literature review. Sections 2–5 present background material, the framework for analysis, and a description of the data. Sections 6–9 contain all results. Specifically, Section 2 provides an overview of the long-term care insurance market. Section 3 presents the model. Section 4 introduces the VRI and the key data items

[2] An essential feature that makes SSQs of value in model estimation is that they are quantitative rather than qualitative. In this sense, there are analogies with the quantitative questions about beliefs pioneered by Juster (1966) and Manski (1990).

on which our analysis rests. Section 5 discusses strategic survey questions, including parameter identification, survey instrument design, and response credibility and coherence. Section 6 details the estimation strategy, the individual-specific parameter estimates, and economic interpretation of the estimated preferences. Section 7 contains the main results of the paper by providing model-based estimates of demand for ADLI. Section 8 presents stated ADLI demand. Section 9 contains the findings for annuities. Section 10 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. 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 out from government provided care (Pauly (1990), Brown and Finkelstein (2008)) with means tested programs having effects on both low wealth and affluent households (Braun, Kopecky, and Koreshkova (2016a)). The perceived value of this publicly provided care is certainly a determinant of the demand for private LTCI. Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011) estimated preferences reflecting a sizable degree of public care aversion. Furthermore, Hackmann (2015) 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. We complement this body of work by developing measures of the counterfactual demand for high quality private insurance in a model with heterogeneous preferences and government provided care—isolating the supply from the demand-side contributions to the low observed LTCI holding.

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), 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 money hampers identification. Thus, many papers match moments on insurance demand to estimate models. Inference achieved by matching measured to model-implied ownership moments is of particular use when insurance products sold in the market closely resemble the corresponding state-contingent model objects. For LTCI, as documented in Section 2, there is ample reason to believe the product does not closely resemble a simple state-contingent claim that pays without risk to be used as desired in the perfectly-verifiable ADL state. Essentially, 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. We identify preferences using SSQs, in contrast to the more common approach of estimating preferences that explicitly match insurance holdings (e.g., Lockwood (2016) and Koijen, Van Nieuwerburgh, and Yogo (2016)). Thus, by design, our estimation strategy does not necessarily explain low insurance ownership by weak preferences for such products. We view these two approaches as complementary for estimating preferences in general, but find our approach more suitable for predicting the latent demand for unavailable insurance products when it is difficult to map imperfections in existing market products to modeled insurance products.

Health-state Utility. Since demand for LTCI depends crucially on the desire to have wealth in the state of the world when help is needed with ADLs, we are intimately connected to the literature that estimates health-state dependent utility functions. 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. Our approach uses stated preference survey methods, which complements previous research using other techniques, such as the health and consumption dynamics approach of Lillard and Weiss (1997), health and utility proxy dynamics approach of Finkelstein, Luttmer, and Notowidigdo (2013), and the compensating differentials approach of Viscusi and Evans (1990). 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 using a general poor-health state, since estimates may be highly contextual and LTC is a distinct health state associated with different care, maladies, and behavior. 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. Furthermore, the demand for LTCI is driven not just by preferences in the ADL state of the world, but also by preferences in all possible health states. Thus, our work also contributes to the literature estimating risk aversion and bequest utility parameters.

Strategic Survey Questions. Survey measurement of model preference parameters was initiated by Barsky, Juster, Kimball, and Shapiro (1997) who estimated risk tolerance using stated preferences over lotteries. The methodology was refined in Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011) who called this approach strategic survey questions (SSQs). SSQs are central to the surveys of the Vanguard Research Initiative (VRI). They specify choice scenarios engineered to aid in estimation of a life-cycle model in which the ability to identify parameters using available behavioral data is very limited. For these questions, as for the expectations questions before them, one must establish credibility using additional measurements. A particular innovation in this paper is the extensive process for designing SSQs, including detailed yet simple descriptions of the scenarios and explicit tests of subject comprehension (see Section 5).

This work forms part of a growing body of research that incorporates answers to theoretically-inspired survey questions in estimating structural models of important life-cycle choices. The pioneering work on expectations measurement due to Juster (1966) and Manski (1990) is designed with just such estimation possibilities in mind and, following their lead, such data has been gathered by the Health and Retirement Study in many domains (as discussed in Attanasio (2015)). van der Klaauw and Wolpin (2008) make explicit use of HRS data on retirement and longevity expectations in estimating a structural model of the link between social security and the retirement and saving behavior of low-income households. Wiswall and Zafar (2015) implement a specific information intervention and combine it with belief measurement to separate informational from preference-based determinants of college major choice. As the importance of measuring expectations has been increasingly realized, so the fundamental reasonableness of numerical answers has been confirmed (Manski, 2004) and further improvements have been made in the design process (Delavande and Rohwedder, 2008). Specifically in the context of medical insurance demand, Kesternich, Heiss, McFadden, and Winter (2013) show that hypothetical choice questions have predictive power on willingness to pay and market shares that were later revealed in real Medicare Part D plan choices.

While different in structure from measurement of expectations and preferences, our approach is spiritually akin to the work of Paweenawat and Townsend (2012) on measurement of household and village financial accounts and of Attanasio, Cunha, and Jervis (2015) and Attanasio and Cattan (2015) in the estimation of human capital production functions. In both cases, as in ours, a complete theoretical framework was developed to guide the design of survey questions.

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 (2016) 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 (2015) 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 while this one estimates how preferences differ individual-by-individual and focuses on the demand for insurance.

Bequests have also long been accepted as an important saving motive, with Kotlikoff and Summers (1981) and Hurd (1989) modeling and estimating their contributions to wealth accumulation. The bequest motive itself reflects a variety of intergenerational linkages, from joy of giving (Abel and Warshawsky (1988)) to strategic bequest motives (Bernheim, Shleifer, and Summers (1985)). A workhorse in modern quantitative models, Nardi (2004) introduced a flexible end of life bequest functional form, and estimated a luxury bequest motive for individuals with large financial resources. Since both precautionary and bequest saving motives could drive observed saving behavior, in our paper, we allow for flexible functional forms for both bequest and health-related utility functions, and use SSQs to estimate preferences and see how they contribute to insurance demand.

2 The Long-term Care Insurance Market

The low levels of private LTCI ownership reveal that most consumers assign little value to LTCI products available on the market. In well-functioning competitive insurance markets, low demand indicates little desire to insure the states covered by the product and, by extension, a relatively low desire for wealth in those states. Although this lack of desire to insure could be caused by a perception that the risk is insignificant, public insurance crowding out private insurance demand, or a low marginal utility of resources in the insured state, in the case of LTCI it could additionally reflect substantial imperfections of private LTCI products. In this section we summarize features of the private LTCI market that complicate separating a lack of desire to insure the ADL-health state from a lack of desirability of available products.

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 existing policies do not differentiate prices by gender, 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.

In addition to being under-developed in many ways, the private LTCI market actually appears to be regressing. 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

[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.

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 might 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 (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 (2016b)).

The state of the LTCI market raises the question of how to interpret low LTCI holdings, limiting the ability to separate a lack of desire to insure LTC expenses from a lack of desire to purchase seemingly poor quality products. Thus, the map between existing products and modeled products should be accounted for when inferring preference parameters from observed LTCI holdings. Analysis is further complicated by multiple product flaws because, 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." For the remainder of this paper, 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 holdings of 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 (2015), 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 (2016). Since the 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 state-of-the-art model otherwise similar to those used in the recent literature.

The model considers individuals who are heterogeneous over wealth, income age-profile, age, sex, health status, and preferences. An individual's health status can either 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 commonly regarded as provoking need for long-term care. Health status evolves according to a Markov process conditional on age, gender, and prior health status. Individuals start at age 55 and live to be at most 108 years old. 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.[4] Each individual has a perfectly foreseen deterministic income sequence and receives a risk free rate of return of $(1+r)$ on savings. The risk free return is calibrated to a baseline 1 percent, although Section 7.3.1 shows that results are robust to allowing for a 3 percent rate. The only uncertainty an individual has is over health/death.

When in good or poor health, consumers value consumption according to standard CRRA preferences with parameter $\gamma > 0$:

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

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

$$(\theta_{ADL})^{-\gamma}\,\frac{(c+\kappa_{ADL})^{1-\gamma}}{1-\gamma}.$$

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. Finally 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

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

Both ADL state and bequest utility are each governed by two key parameters: θ and κ. θ scales the marginal utility of an additional dollar spent and κ controls the degree to which the expenditure is a luxury or a necessity by deviating from homotheticity. An increase in θ decreases the marginal utility of a unit of expenditure; an increase in κ indicates that expenditure is more of a luxury. Negative κ 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.[5] 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 $c = \psi_G$. The value ψ_G parameterizes the consumer's value of public care, since that parameter essentially determines the utility of an individual who needs help with ADLs and chooses to use government care. There is no borrowing, and the retiree cannot leave a negative bequest.

Let wealth be $a \in [0,\infty)$, income age-profile be $y \in \{y_1, y_2, \ldots, y_5\}$, age be $t \in \{55, 56, ..., T = 108\}$, gender 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), health cost be h, and $G \in \{0, 1\}$ be the government care indicator. Then, written recursively, the consumer

[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.

[5] This is a parsimonious way to model that public assistance is only available to people with sufficiently low financial resources.

problem is:

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

See "Vanguard Research Initiative Technical Report: Long-term Care Model" for further description and computation of optimal decision rules.[6]

Together $\Theta_i := \{\gamma, \theta_{ADL}, \kappa_{ADL}, \theta_{beq}, \kappa_{beq}, \psi_G\}$ define an individual's preferences over risk, expenditure in the ADL-state, and bequests. Depending on the estimated values at the individual level, 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 prefer spending while healthy. Together these preferences, demographic and financial variables, and the estimated health and longevity risks, determine demand for ADLI. The key challenge is estimation of Θ_i, especially because the typical wealth data used to identify such preferences are weakly informative. A main contribution of this paper is estimation of Θ_i using new data and methods as described in the next few sections.

4 Sample and Data Overview

4.1 The Vanguard Research Initiative

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. Because singles are better suited for research that does not directly model family interaction, singles are over-sampled in the VRI. 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.

VRI respondents participated in three surveys used in this paper that were administered between June 2013 and

[6] The non-concavity of the value function and the discontinuity in the optimal saving policy introduce computational complications. We use a modified endogenous grid method, building on insights from Fella (2014).

August 2014.[7] VRI Survey 1 measures all of the state variables of the model for each respondent (wealth, income, age, gender, 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**			**Gender**	
	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 Gender: 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 this paper we also make use of additional data measured in the VRI. In addition to the SSQs and stated demand questions, respectively detailed in Sections 5 and 8, 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.

4.2 Health and Mortality Estimation

In addition to using the financial and demographic data from the VRI, we estimate age and gender specific Markov transition matrices across health states using longitudinal data from the HRS subsample meeting the VRI age and wealth criteria. Individuals are defined as in good health if they report health being good, very good, or excellent, and are defined to be in poor health if they report health being 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.

To highlight the magnitude of LTC risk, Figure 1 presents the estimated distribution of the number of years spent

[7] Participants in this study receive a small incentive for participation in each survey in the form of sweepstakes for prizes such as an iPad, 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.

needing help with ADLs 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 not need any help with ADLs while alive. 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. Given that the average cost of one year in a nursing home is $84K, this substantial probability of needing care for many years highlights the large magnitude of LTC risk.

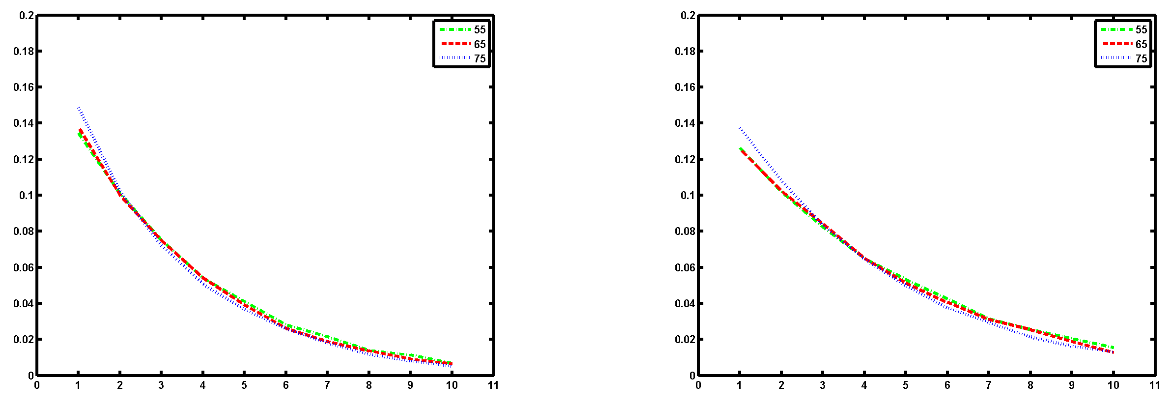

(a) (b)

Figure 1: Years Needing Help With ADLS: Panel (a) presents the distribution 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 Strategic Survey Questions

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. Since these SSQs are at the heart of this paper, we dedicate this section to detailing the SSQ identification strategy, survey instrument design, and the validity and coherence of responses.

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 choice environment is subject to restrictions 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.[8] 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.

Below we describe in detail how SSQs are designed to identify preference parameters by construction, how the survey questions were designed to help respondents understand the situation and choice while trying to make the verbal problem adhere as closely as possible to the math problem, how the survey instrument was designed to help the respondent record their choice, and some checks that the responses are logically intelligible.

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. For exposition, we sketch the identification argument for θ_{beq} and κ_{beq} using SSQ 3 assuming that γ, θ_{ADL}, and κ_{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:

$$\begin{aligned} \max_{z_1,z_2} \quad & (\theta_{ADL})^{-\gamma} \ \frac{(z_1+\kappa_{ADL})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma} \ \frac{(z_2+\kappa_{beq})^{1-\gamma}}{1-\gamma} \\ \text{s.t.} \quad & z_1 + z_2 \le W \\ & z_1 \ge 0; \ z_2 \ge 0. \end{aligned} \tag{1}$$

The optimal allocation rule is given by

$$z_1 = \begin{cases} 0 & \text{if} \ \ (\theta_{beq} \ (W+\kappa_{beq}))^{-\gamma} - (\theta_{ADL} \ \kappa_{ADL})^{-\gamma} > 0 \\ W & \text{if} \ \ (\theta_{ADL} \ (W+\kappa_{ADL}))^{-\gamma} - (\theta_{beq} \ \kappa_{beq})^{-\gamma} > 0 \\ \frac{\theta_{beq}(W+\kappa_{beq})-\theta_{ADL}\kappa_{ADL}}{\theta_{ADL}+\theta_{beq}} & \text{otherwise.} \end{cases} \tag{2}$$

Conditional on γ , θ_{ADL}, and κ_{ADL}, the interior response is linear in wealth, and thus θ_{beq} and κ_{beq} are identified by two interior responses to the question posed at different wealth levels, W. 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. These responses identify all relevant structural model parameters.[9] The optimization problem and optimal allocations (as a function of preference parameters) corresponding to each SSQ is presented in "Vanguard Research Initiative Technical Report: Long-term Care Strategic Survey Questions".

[8]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.

[9]The parameters ω_G and β are not identified by any of the SSQs, and thus are calibrated to standard values from the literature.

5.2 Design of SSQ Survey Instruments

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. We illustrate the resulting implementation by detailing a particular SSQ (SSQ 3). This SSQ focuses on the tradeoff between expenditure when in need of help with ADLs and leaving a bequest.[10] We begin with a broad introduction to the subject of interest and then present the scenario.

> We are now going to ask about a different situation where you are older and definitely need long-term care. In this situation, you are asked to make tradeoffs between spending on your long-term care and leaving a bequest. This scenario is hypothetical and does not reflect a choice you are likely ever to face.
>
> Suppose you are 85 years old, live alone, rent your home, and pay all your own bills. You know with certainty that you will live for only 12 more months and that you will need help with *ADLs for the entire 12 months.
>
> You have **$100,000** that you need to split into Plan E and Plan F.
>
> - Plan E is reserved for your spending. From Plan E, you will need to pay all of your expenses, including long-term care and any other wants, needs, and discretionary purchases.
> - Plan F is an irrevocable bequest.

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., 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.

> - You have no money other than the $100,000.
> - Other than Plan E, you have no other resources available to help with your long-term care. **You** have to pay for any long-term care you may need from Plan E.
> - No one—including friends or family—can take care of you for free. Long-term care must be purchased at market rates.
> - Any money in Plan E that you do not spend cannot be given away or left as a bequest.
> - Bequests from Plan F are not subject to any taxation.
> - Once you make your choice of plans, you cannot change how you split your money.
> - You have full insurance that covers all of your hospital, doctor, and medications, but you have no long-term care insurance.
> - There is **no public-care option or Medicaid** if you do not have enough money to pay for a nursing home or other long-term care.

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:

[10] To reinforce the definition of needing help with ADLs, respondents were given a comprehension test on the definition prior to this SSQ. Furthermore, we make the definition available in a hover button whenever *ADL appears.

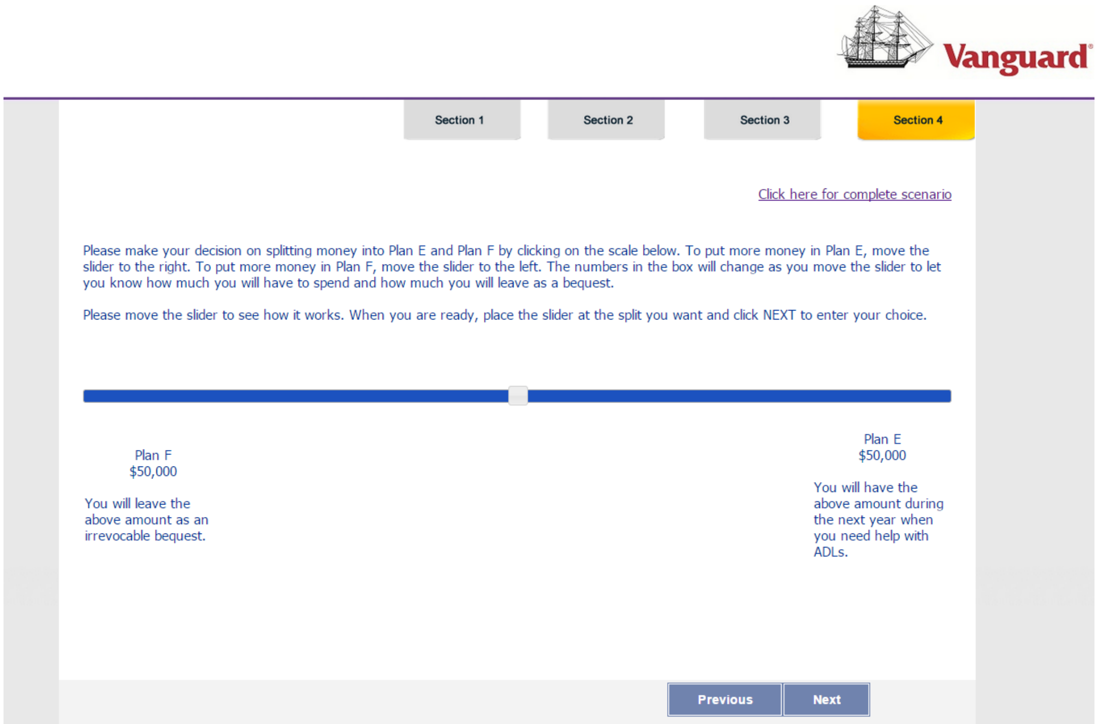

Figure 2: SSQ Response Slider

> 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 for this and the other SSQs 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 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 can be seen in Figure 2.

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

	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	$\max_{z_1,z_2} \ \pi \frac{z_1^{1-\gamma}}{1-\gamma} + (1-\pi)\frac{(\theta_{LTC})^{-\gamma}(z_2+\kappa_{LTC})^{1-\gamma}}{1-\gamma}$	(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	$\max_{z_1,z_2} \ (\theta_{LTC})^{-\gamma} \ \frac{(z_1+\kappa_{LTC})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma} \ \frac{\left(z_2+\kappa_{beq}\right)^{1-\gamma}}{1-\gamma}$	(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	$W^* : (\theta_{LTC})^{-\gamma} \ \frac{(\psi_G+\kappa_{LTC})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma} \ \frac{\left(W^*+\kappa_{beq}\right)^{1-\gamma}}{1-\gamma} =$ $(\theta_{LTC})^{-\gamma} \ \frac{(z_1+\kappa_{LTC})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma} \ \frac{\left(W^*-z_1+\kappa_{beq}\right)^{1-\gamma}}{1-\gamma}$	(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.

respondents to move the slider at least once, which helps also to clarify the connection to the chosen allocation.[11] 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 stayed within the scenario and asked respondents to make new choices with different scenario parameters. In the above question, answers were gathered not only concerning division of $100,000, but also of $150,000 and $200,000.

In addition to this SSQ, we posed three other SSQs.

- SSQ 1 asks about willingness to take a risky bet over annual expenditure, using an analogous survey question and identification strategy to those developed in Barsky, Juster, Kimball, and Shapiro (1997) and Kimball, Sahm, and Shapiro (2008).
- SSQ 2 asks individuals facing uncertain future health to allocate wealth to states either when healthy or when in need of help with ADLs.
- SSQ 4 asks individuals how much wealth they would need to have in order to purchase private LTC instead of using government provided care.

A brief summary of these SSQs and their variants is presented in Table 2. For all of the SSQs, the survey instruments have the same structure as that described for SSQ 3: statement of the scenario and rules, comprehension-verification questions, restatement of the scenario, and slider visualization for recording responses. In "Vanguard Research Initiative Technical Report: Long-term Care Strategic Survey Questions" we present the text for each SSQ, including all rules and a full list of comprehension questions. The results of these comprehension tests are summarized in the next section.

[11] 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.

5.3 Credibility of SSQ Responses

A number of features of the VRI data give us high confidence in the credibility and quality of the survey measures. In this section we summarize a few main findings on the validity of the data. First we summarize objective and subjective measures of comprehension and then we discuss measures of the internal and external coherence of SSQ responses. We refer to Appendix B for details of this analysis and additional evidence of SSQ credibility.

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. In the case of SSQ 1 (where there were 6 questions) 46 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. Similar patterns hold across other SSQs (see Table B.1), indicating a high level of comprehension.

In addition to objective measures of respondent comprehension, the production survey includes a set of subjective comprehension questions motivated by the real-time pilot chat responses. Survey respondents report it was easy to comprehend the tradeoffs asked about in the SSQs: 90 percent of respondents found the tradeoffs either very clear or somewhat clear, 84 percent indicated that they were able to place themselves in the hypothetical scenario either moderately or very well, and 82 percent had given the underlying issues at least a little thought before taking the survey. 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.

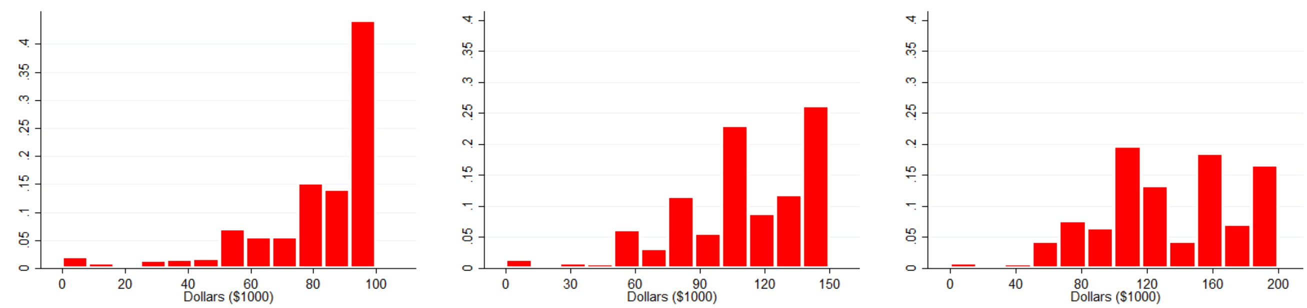

(a) \$100K **(b)** \$150K **(c)** \$200K

Figure 3: SSQ 3 Response Distributions: We ask SSQ 3, the SSQ presented in the section above, for wealth values of \$100,000, \$150,000, and \$200,000. For these three wealth values, the figure shows the distributions of allocation to spending on self when needing help with ADLs (with the remainder left as a bequest).

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. We see clear indications of internal coherence both at the aggregate and individual level. First, we observe meaningful systematic patterns in the response distributions when changing scenario parameters of a given SSQ. For example, in SSQ 3, all respondents were asked to allocate not only \$100,000, but also \$150,000 and \$200,000. Most respondents place almost all of their wealth in the ADL state when wealth is \$100,000, about two-thirds in the ADL state when wealth is \$150,000, but only about half when wealth is \$200,000, as illustrated in Figure 3. Second, we observe a strong positive correlation in within-individual responses to SSQs of the same type. In the case of SSQ 3, this indicates that individuals that allocated more money to the LTC state when wealth was \$100,000 also allocated more wealth to the LTC state when wealth was \$150,000 and \$200,000. Similar patterns hold across all SSQs, alleviating concerns that respondents replied with a large degree of randomness.

Finally, 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 B.4 through B.7 report regressions of SSQ responses on demographic and economics covariates. Indeed, there is little significance in coefficients on age, income, health, and wealth, suggesting that this design is successful.

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.

6 Estimating 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 documented 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 $\Theta_i = \{\gamma, \theta_{ADL}, \kappa_{ADL}, \theta_{beq}, \kappa_{beq}, \psi_G\}$. 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.

To derive the likelihood function, we denote the true response to the k^{th} SSQ as $z_k(\Theta_i)$ and assume each individual's response is reported with normally distributed errors. That is, let observed responses be

$$\hat{z}_{k,i} = z_k(\Theta_i) + \hat{\epsilon}_{k,i}, \tag{3}$$

where $\epsilon_{k,i} \sim \mathbb{N}(0, \sigma^2_{k,i})$ and $\hat{\epsilon}_{k,i}$ denotes the realization of individual i's response error to SSQ variant k. For robustness, in Section 7.3.1 we show results for a multiplicative error structure that assumes $\log \hat{z}_{k,i} = \log z_k(\Theta_i) + \hat{\epsilon}_{k,i}$.

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 $\sigma^2_{k,i}$ 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 $\bar{\sigma}_i$, so that $\sigma_{k,i} = \sigma_i \times W_k$. The idiosyncratic component accounts for differences in the precision with which individuals report answers. The question specific component takes into account the different scales of the nine SSQ variations and thus normalizes the error standard deviation according to the feasible response size. Note that W_k is naturally defined in each question by the budget constraint, except in SSQ 4. In SSQ 4, W_k is set to the 95^{th} percentile of the survey responses, resulting in \$500,000 as the maximum response in the cleaned data.[12]

This specification yields the following closed form expression for the likelihood of observing a response to each question as a function of (Θ_i, σ_i):

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

$$\mathcal{L}_k(\Theta_i, \sigma_i | \hat{z}_{k,i}) = \begin{cases} F_{\sigma^2_{k,i}}(-z_k(\Theta_i)) & \text{if } \hat{z}_{k,i} = 0 \\ f_{\sigma^2_{k,i}}(\hat{z}_{k,i} - z_k(\Theta_i)) & \text{if } 0 < \hat{z}_{k,i} < W_k \\ 1 - F_{\sigma^2_{k,i}}(W_k - z_k(\Theta_i)) & \text{if } \hat{z}_{k,i} = W_k. \end{cases} \tag{4}$$

The boundary cases take into account error truncation due to the budget constraint, and $F_{\sigma^2_{k,i}}$ and $f_{\sigma^2_{k,i}}$ denote the mean-zero normal cdf and pdf with variances $\sigma^2_{k,i}$. We assume independence of survey response errors, yielding a multiplicatively separable likelihood function for the full response set $\hat{Z}_i$:

$$\mathcal{L}(\Theta_i, \sigma_i | \hat{Z}_i) = \prod_{k=1}^{9} \mathcal{L}_k(\Theta_i, \sigma_i | \hat{z}_{k,i}). \tag{5}$$

We use MLE to estimate individual parameter sets, such that

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

This provides a consistent estimate of each individual's parameter set for those whose parameters are identified. Parameters 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.

In this paper, identification is achieved via multiple responses to SSQ variants at different scenario parameterizations. 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 assume 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.

6.1 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 the 10th/25th/50th/75th/90th percentiles of the marginal distributions for the estimated population parameter distribution. The median marginal estimates suggest a relative risk aversion parameter $\gamma = 4.45$, ADL expenditure as a necessity ($\kappa_{ADL} < 0$) with high marginal valuations ($\theta_{ADL} < 1$), bequests as a significant luxury ($\kappa_{beq} > 0$) with a 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 estimate of the dollar equivalent of public long-term care corresponds to an equivalent utility level of an expenditure

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
Ameriks, et.al (2016)	5.85	1.57	-45.65	0.59	7.88	85.11

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.

of $41,063 in a model without state dependent preferences.[13] For a point of comparison, the bottom row of Table 3 presents the estimates from Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2015), which assumes homogeneous preferences and uses an estimator matching both SSQs and moments of the wealth distribution. Given the difference in estimation procedures and that the presented marginals do not account for correlation between parameters, there is no clean mapping of parameters across studies. Comparison does, however, show consistency in qualitative patterns. Furthermore, previous estimates with homogeneous preferences are contained between the $25^{th} - 75^{th}$ percentiles of the estimated parameter distribution.

The parameter sets are quite well identified. 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, given 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 4.

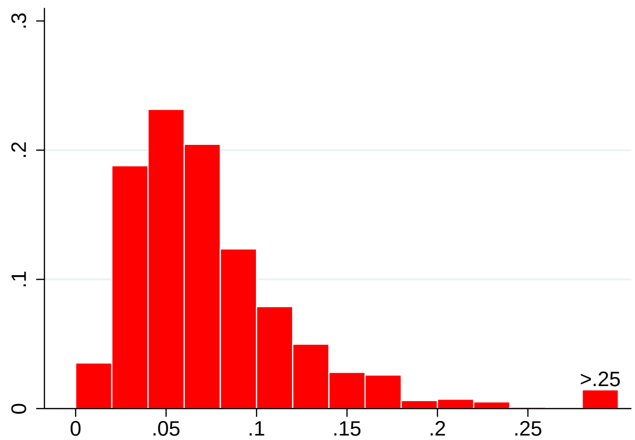

Figure 4: Distribution of Individual Response Error Standard Deviation σ_i

Given that it is hard to interpret preference parameters in isolation, partly because parameter values are inherently difficult to interpret and partly because the interpretation of any one parameter depends on values of other parameters, we interpret 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 to answer the real questions of interest.

Interpreting Preference Parameters Using Simple Synthetic Choice Problems. In Figure 5, we present the 10th/25th/50th/75th/90th percentiles of allocations to the ADL state in response to SSQ 3 that are implied by the

[13]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: $\frac{\bar{\psi}^{1-\gamma}}{1-\gamma} = (\theta_{ADL})^{-\gamma} \frac{(\psi_G+\kappa_{ADL})^{1-\gamma}}{1-\gamma}$.

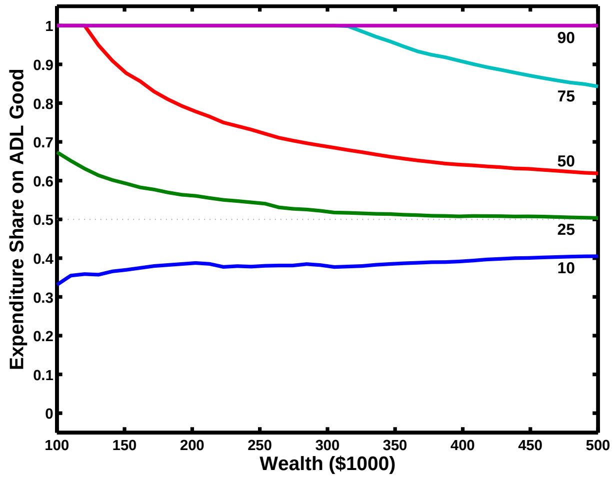

Figure 5: Distribution of SSQ 3 Responses Implied by Estimated Parameters: This figure plots quantiles of the distribution of allocations to the ADL state in response to SSQ 3, for different levels of wealth W, that are implied by the estimated distribution of preference parameters. The problem is not 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.

estimated distribution of preference parameters.[14] Due to the κ parameters, the allocations are not invariant to the wealth level. At $100K, even those in the 25th percentile of allocations to the ADL state spend much more than 50 percent of their wealth on own expenditure when needing help with ADLs and leaving around 35 percent as a bequest. The 50th percentile person leaves no bequest. At $200K, some still value spending on self when in need of help with ADLs relative to leaving a bequest so highly that they leave no bequest, but the median person spends about $160K on self and leaves a $40K bequest; the 25th percentile person almost splits the money evenly. In summary, at lower levels of wealth, spending on self when needing help with ADLs dominates leaving a bequest for almost everyone in the sample; even up through a relatively high level of annualized wealth the large majority of individuals spend more on self than leave a bequest. Some people do value bequests highly, however. At the 10th percentile, spending on a bequest is valued greater than spending when in need of LTC at all wealth levels. Even for these people with the strongest bequest motive relative to self spending when needing help with ADLs, they still leave only 60 to 70 percent as a bequest.[15]

Figure 6 presents statistics on expenditure in a 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

[14] 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 ADL utility is too strong, the figure provides a rough lower bound on the strength of the ADL saving motive in the population.

[15] 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. This representation makes the ADL related saving and insurance demand motive even stronger, since even in a 1 year static problem the allocation skews towards the flow utility state.

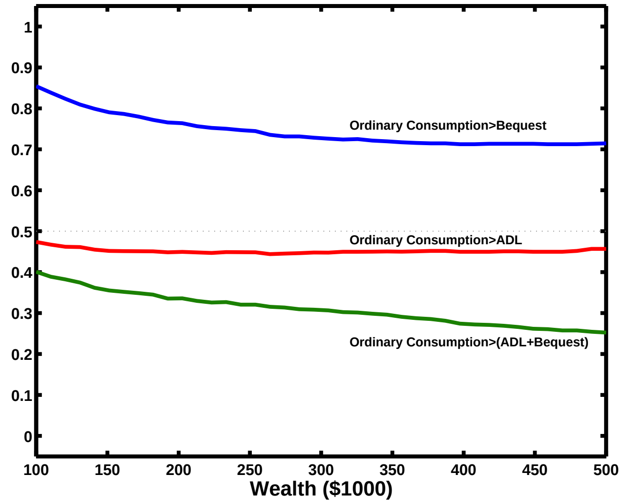

Figure 6: Expenditure in the three good synthetic choice problem: The above figures present statistics on expenditure in the three good synthetic choice problem in which we treat the utility function associated with ordinary health, needing help with ADLs, and bequests as the utility function associated 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."

"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:

$$\begin{aligned} \max_{x_1,x_2,x_3} \quad & \frac{(x_1)^{1-\gamma}}{1-\gamma} + (\theta_{ADL})^{-\gamma}\frac{(x_2+\kappa_{ADL})^{1-\gamma}}{1-\gamma} + (\theta_{beq})^{-\gamma}\frac{(x_3+\kappa_{beq})^{1-\gamma}}{1-\gamma} \\ & \text{s.t. } x_1+x_2+x_3 \leq W \\ & \quad x_1,\ x_2,\ x_3 \geq 0;\ x_2 \geq -\kappa_{ADL};\ x_3 \geq -\kappa_{beq}. \end{aligned} \tag{6}$$

About half the population 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 being valued on the margin. As we show in Section 7, it is exactly these patterns in preferences that largely determine the substantial model-predicted demand for activities of daily living insurance.

7 The Long-term Care Insurance Puzzle

Having developed the model and solved for the associated optimal policies and value function, collected the financial and demographic data, estimated health and mortality risk, and estimated preference parameters at the individual level, we are now able to predict demand for Activities of Daily Living insurance for each person in the sample.

This section documents that the long-term care insurance puzzle—that model predicted insurance demand is larger than observed holdings—is sizable and robust. We detail how to calculate ADLI demand and show that predicted demand is significantly determined by health-state dependent preferences. We also show that many more people are predicted to demand ADLI than actually hold LTCI, that predicted demand is robust to many different model specifications and in different subsamples, that demand is large on the intensive margin for many, as is the associated consumer surplus.

7.1 Calculating Activities of Daily Living Insurance Demand

Using each individuals' financial and demographic states and estimated preference parameters, we calculate the model-implied demand for insurance products. ADLI is modeled as a state contingent security that pays out whenever an individual is in the ADL health state $(s = 2)$. Purchasing this product entails paying a lump sum of $\$\tilde{y}_i \times p(t_i, s_i, g_i)$ at current age t in return for payout $\tilde{y}_i$ in each year that assistance with ADLs is needed for the remainder of life. The demand is thus determined by preference over expected future consumption streams as a function of preference parameter set, Θ_i, the set of state variables, X_i, and the price, $p(t_i, s_i, g_i)$, that individuals must pay to purchase an additional unit of state contingent income.

The pricing function is such that the product is actuarially fair conditional on an individual's gender, age, health state, and access to a risk free outside asset promising 1 percent annual return.[16] Actuarially fair is defined such that the insurer selling this product makes zero expected profit (using the same health transition matrix as in the decision problem).[17] 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.

[16] In the baseline specification, ADLI has no load. We show results for different loads in Section 7.3.1.

[17] The realized period payouts for annuities and ADL insurance depend on health state s. An annuity pays out while $s = 0, 1$ or 2, while ADLI pays out while only when $s = 2$. Thus, the vector of period payouts across health states $s \in \{0, 1, 2, 3\}$ for annuities is $\vec{\tilde{y}} = [\tilde{y},\ \tilde{y},\ \tilde{y},\ 0]'$, while for ADLI it is, $\vec{\tilde{y}} = [0,\ 0,\ \tilde{y},\ 0]'$. Let $\vec{s}$ be an indicator vector that has elements s_i for $i \in \{0, 1, 2, 3\}$ equal to zero for $s \neq i$ and equal to 1 if $s = i$. The insurance product is priced to equal the expected discounted stream of payments. Thus, an insurance product that pays out $\vec{\tilde{y}}$ per period for a person of age t, gender g, with current health status s has price

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

Given prices, demand for insurance is calculated as

$$D(a, y, t, s, h, g) = \arg\max_{\tilde{y}} V(a - p(t, s, g)\tilde{y}, \hat{y}, t, s, h, g) \tag{7}$$
$$\hat{y} = \{y_\tau + \tilde{y}(s)\}_{\tau=t}^{T},$$

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

To account for uncertainty around estimated parameter values when calculating model implied demand, we resample five parameter sets for each person from the 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 predicted demand caused by parameter uncertainty. For the remainder of the paper, all reported baseline results reflect these bootstrapped estimates.[18]

7.2 Estimated Activities of Daily Living Insurance Demand

We estimate that 59 percent of respondents have positive demand for ADLI. This indicates that the majority of individuals assign a high valuation to wealth in the ADL state and, if offered suitable insurance products, would like to insure wealth in this state. While many are estimated to want ADLI, there is a substantial minority for whom purchasing is not predicted to be attractive. Majority interest is not built into the specification, but rather a result of desires as inferred from the responses to SSQs.[19]

Although estimates are a function of observable demographics and financial wealth, preferences significantly affect ADLI demand, such that different survey responses would have produced completely different estimates. Table 4 compares the mean parameter values for individuals predicted to purchase versus not purchase ADLI. Most differences between the groups are as expected. ADLI purchasers are significantly 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 more of a necessity. 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. Similarly, purchasers of ADLI have a lower θ_{beq} and κ_{beq}. The comparison of bequest motives is less theoretically 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 large expenditures when in the ADL state. Table 4 suggests that the second motive is dominant, since purchasers have stronger bequest motives than non-purchasers, viewing it as substantially less of a luxury good (lower κ_{beq}). This suggests that those predicted to demand zero ADLI are motivated by the desire to spend on self when healthy.

[18]For the robustness exercises in Section 7.3.1, due to computational run time limitations, we use one sample of preference parameters that is held constant across exercises.

[19]Lockwood (2016) shows that a sufficiently strong bequest motive limits interest in either LTCI or annuities due to a preference for liquid wealth at the end of life, and Lockwood (2016) and Koijen, Van Nieuwerburgh, and Yogo (2016) both match observed insurance products holdings in their estimations. A key methodological difference with the current study is use of observed insurance holdings as a source of identification. While we agree that insurance holding patterns contain information about preference for wealth in insured states, targeting low LTC insurance holdings ensures that the model delivers a low preference for wealth in this state. Given the differences between modeled state-contingent securities and products available in the market, we instead identify preferences from SSQs and analyze their implications for predicted holdings.

	Preference Parameters					
Purchase	γ	θ_{ADL}	κ_{ADL}	θ_{beq}	κ_{beq}	ψ_G
Yes	5.57	4.02	-11.26	2.07	118.17	76.82
No	4.31	23.11	33.86	22.69	278.84	77.42

	State Variables				
Purchase	**Age**	**Income Quint**	**Wealth**	**Gender**	**Health**
Yes	69.2	3.24	820,510	0.45	1.06
No	67.1	2.84	517,129	0.40	1.07

Table 4: Parameter sets and ADLI purchase: This table presents averages of demographic, financial, and preference variables for two groups: those with zero ADLI demand and those with positive ADLI demand.

7.3 Documenting the Long-term Care Insurance Puzzle

While the model predicts that 59 percent of the sample would want to purchase ADLI, only 22 percent of the VRI sample own private LTCI.[20] This large difference between actual LTCI holdings and predicted 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 Americans. Figure 7 compares actual LTCI ownership and model predicted ADLI ownership 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. Note that both observed holdings and model predictions of ADLI ownership are increasing in wealth and income. Note also that the difference between modeled and observed holdings is large and significant at all quintiles, confirming the robustness of the puzzle. We therefore conclude that there exists a puzzle regarding the lack of LTCI ownership: Observed insurance holdings among older wealth-holders are well below the levels suggested by the model.

7.3.1 Robustness of the Long-term Care Insurance Puzzle

To document the robustness of the LTCI puzzle, we present in Figure 8 ADLI demand calculated for different model specifications and for different subsamples. First, we show the existence of the LTCI puzzle is not sensitive to reasonable increases in the price of ADLI. To document this, 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, 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.

We also predict 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 positive load (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

[20]Moreover we do not know the extent to which this private ownership is due to deliberate purchase as opposed to being a job benefit, making this an upper bound on the fraction of individuals in the sample who have actively purchased private LTCI.

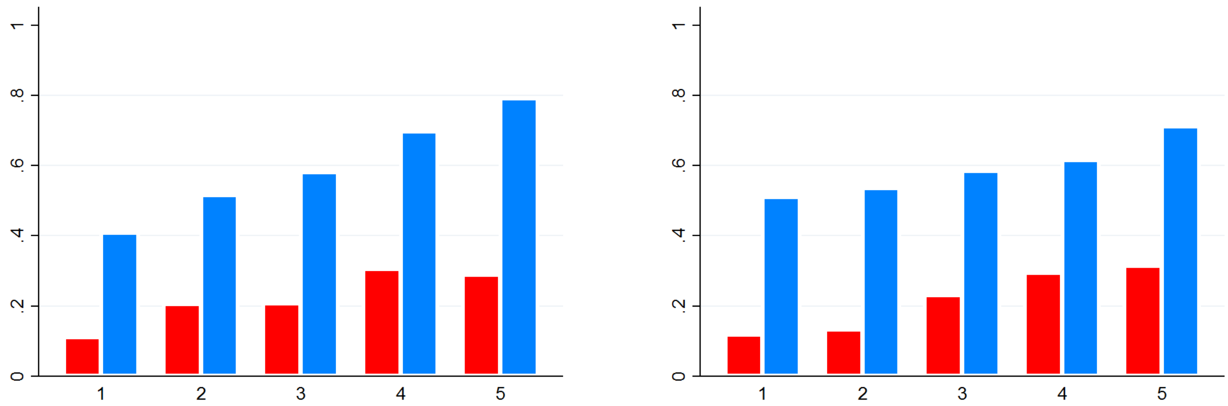

(a) LTCI/ADLI Ownership by Wealth Quintile

(b) LTCI/ADLI Ownership by Income Quintile

Figure 7: Comparing Ownership Measures: 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 green bars on the right are the corresponding model predictions.

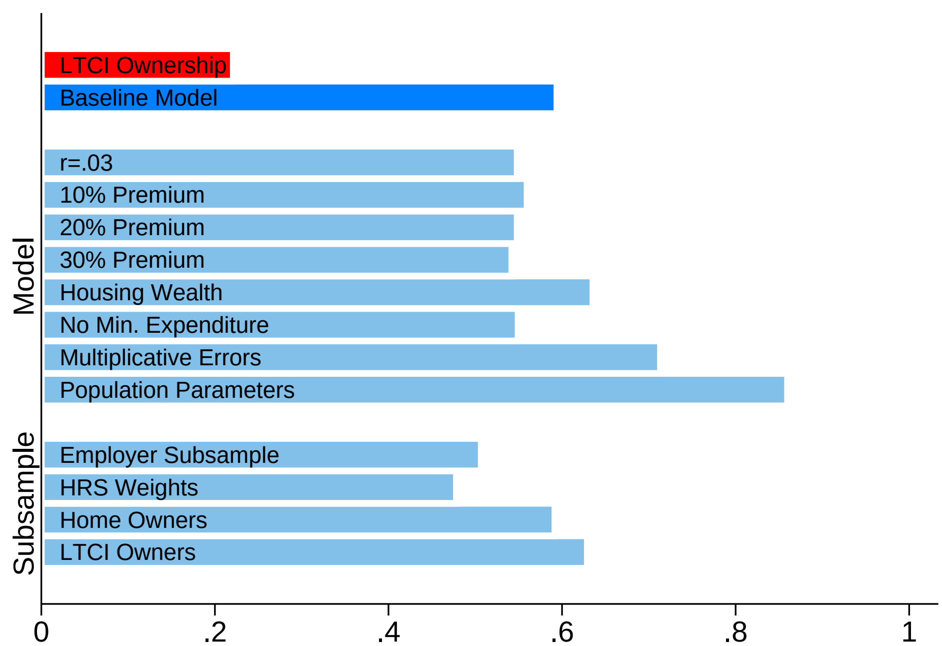

Figure 8: LTCI/ADLI Ownership Rate Robustness—Alternative Parameters and Samples: This figure presents the fraction of the population that is predicted to have positive ADLI demand according to various changes to the model and sample. The top row shows the 22 percent of people who own LTCI in the VRI. The second row shows the prediction from the baseline specification of 59 percent ownership. Subsequent rows present the robustness results.

low returns on investments are not driving the puzzle.

We map net financial wealth to the wealth state variable, 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 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 we model a minimum level of expenditure needed to obtain private LTC, i.e., $e_{LTC} \geq \$40K$ 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.

Since demand is driven in such large part by estimated preferences, we show that the puzzle remains even when estimating preference parameters using more traditional methods common in the literature that do not use SSQs. We calculate ADLI demand using a parameter set from Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2015) that was estimated using the same model but assuming homogeneous preferences and exclusively targeting cross-sectional moments of the wealth-age distribution (the 25^{th}, 50^{th}, and 75^{th} percentiles of the wealth distribution by 3 year age bins). Under this parametrization the model predicts significantly higher ADLI demand, with 86 percent of the population having positive demand. 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 C.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.

7.4 How Much Activities of Daily Living Insurance Would People Demand?

To establish the economic significance of the LTCI puzzle, we show that in addition to the large difference on the extensive margin between ownership and predicted demand for ADLI, the predicted 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 model predicted demand to the known holdings. Nonetheless, we present in Figure 9 ADLI demand measures for both LTCI owners and

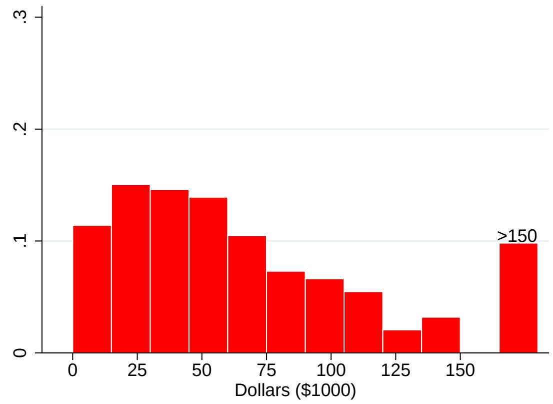

(a) LTCI Non-owners

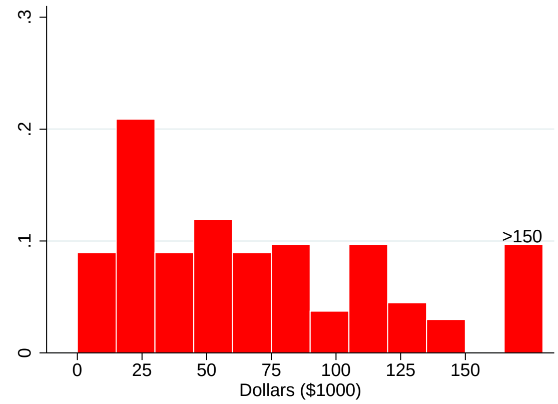

(b) LTCI Owners

Figure 9: ADLI Quantity Demanded: This figure presents the histogram of the ADLI annual payout purchased predicted by 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.

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 are predicted to have positive ADLI demand. Furthermore, as presented in Figure 9a, 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, gender, and buying ADLI, median demand (across wealth and preferences) is $33,400 ($39,400) paid each year LTC is needed 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 C.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.

7.4.1 The ADLI Demand Function

While the amount demanded at given prices is informative, but 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.[21]

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

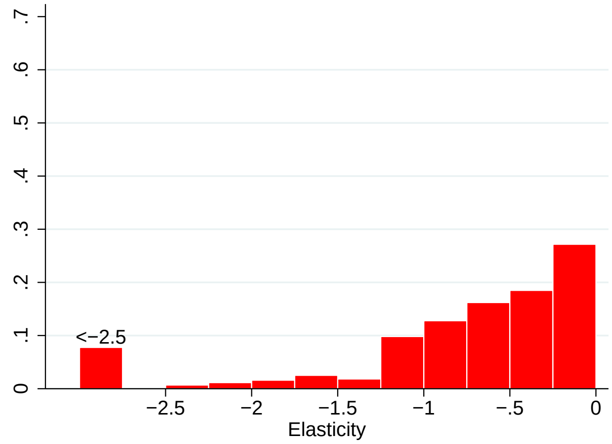

Figure 10: Distribution of the Price Elasticity of Demand: This figure presents the histogram of the elasticity of demand with respect to price for those who do not own LTCI. It plots the distribution of the percent change in demand to a one percent increase in price, local to the optimal demand level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

The Elasticity of Demand to Price. 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 C.1, that the intensive margin of demand is 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. As discussed in Section 2, there are many features of LTCI products that may contribute to low demand other than price.

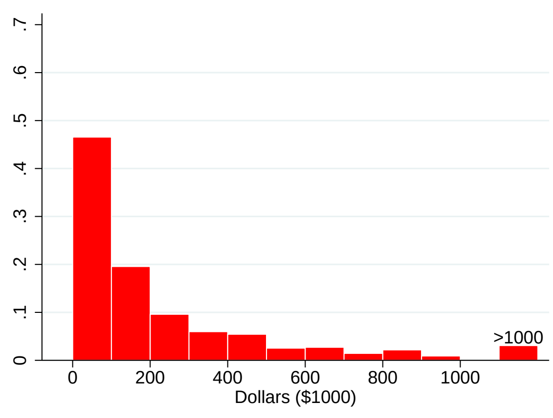

(a) Distribution of Consumer Surplus

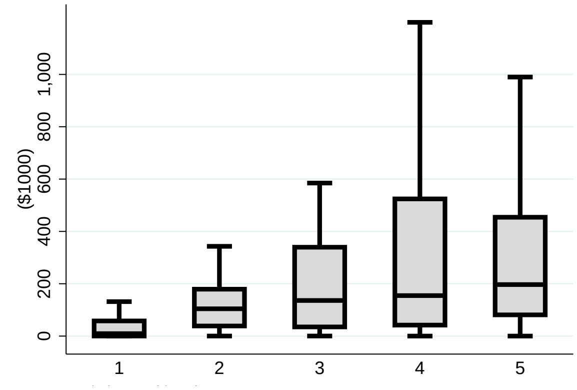

(b) Consumer Surplus Box Plot by Wealth Quintile

Figure 11: Consumer Surplus: The left panel presents the histogram of consumer surplus for those who do not own LTCI. 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.

Consumer Surplus. Consumer surplus is defined as the maximum amount people 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 gender, age, and health status), because the quantity demanded 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 and preferences. As documented in the left panel of Figure 11, many people have a consumer surplus above $100K, with a non-trivial fraction of the population having a consumer surplus larger than $200K. This shows that our model predicts sizable demand for ADLI, suggesting a substantial missing market for higher quality LTCI. The right panel of Figure 11 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 that 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.

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 predicted quantity of 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.

8 Stated Demand for Activities of Daily Living Insurance

The LTCI puzzle is that model predicted 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 Survey 2, which included stated choice questions on the demand for improved insurance products. This provides a model-independent measure of demand for the exact same ADLI product.

To the degree that model predicted and stated demand agree, we have two completely different measures of the same demand that 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. To the degree there is a puzzle that manifests as the difference between model predicted and stated ADLI demand, there is evidence that some factor other than the difference between LTCI an ADLI is driving the LTCI puzzle.

8.1 The Survey Instrument

Our survey elicits stated demand for ADLI.[22] For ADLI, a challenge in gathering this demand is that, by definition, 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." In the specific case of ADLI, the

[22] 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 (2016) elicit stated purchase and sale values for annuities and link these spreads to potential explanations for heterogeneity in financial decision-making abilities.

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, gender, and current health based on the estimated health transition probabilities, determining "$Y" in the frame above.[23] This is reinforced by the qualitative statement that the pricing is actuarially fair. After all information is provided, demand is collected in two steps. We first ask respondents whether or not they would have any interest in purchasing ADLI were it available. If the answer is affirmative, we ask how large a monthly benefit they would purchase, while simultaneously reporting to them how much their purchase of any such benefit would cost up front. 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 a OLS regression on the amount purchased in the subsample of respondents that reported positive demand. Full results are included in Table B.8. Respondents who report higher probabilities of experiencing extended time in the ADL state are more likely to purchase 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

[23] 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. Model-predicted 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.

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 predict higher average LTC costs 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 gender, age, and health status.

8.2 Stated ADLI Demand and the LTCI Puzzle

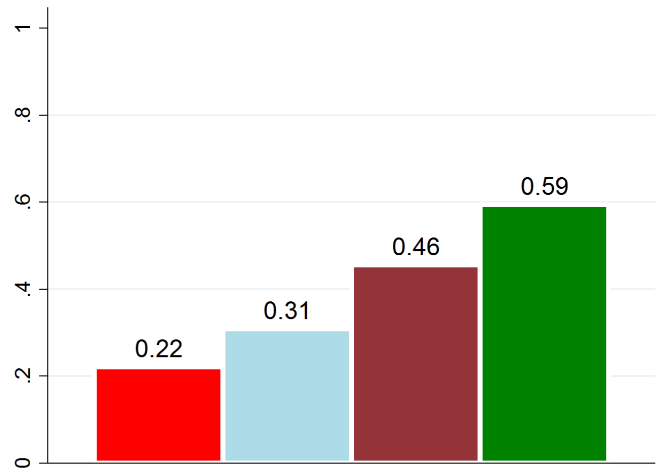

Figure 12: Fraction of Population Owning LTCI: This figure presents various measures of the fraction of the population with positive LTCI ownership. Column 1 is actual holdings of a private LTCI in the sample. Column 2 is stated ADLI demand. Column 3 is the union of private ownership and stated demand. Column 4 is model predicted ADLI demand.

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 44 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 model predicted demand. These different measures of ownership are summarized in Figure 12. 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 model predicted ADLI demand suggests that not all of the difference between predicted and actual holdings is attributable to specific features of the LTCI products currently available in the market.

As seen in Figure 13, 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 84 percent of the model predicted. At the second wealth quintile the owned-or-stated measure is a remarkable 94 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 59 percent in quintile five. The close fit of stated and model predicted 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.

Figure 14 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

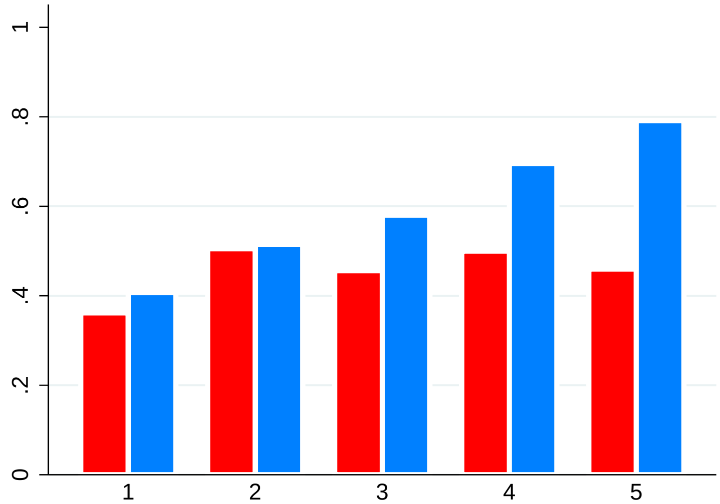

Figure 13: ADLI Ownership by Wealth Quintile: Stated and Model Predicted Demand This figure presents the fraction of the population with positive demand for ADLI by wealth quintile according to the stated demand and model predicted demand. The red bars on the left show the fraction of the population in a given quintile who either own LTCI in the VRI or state positive demand for ADLI in the survey, while the green bars on the right are the corresponding model predictions.

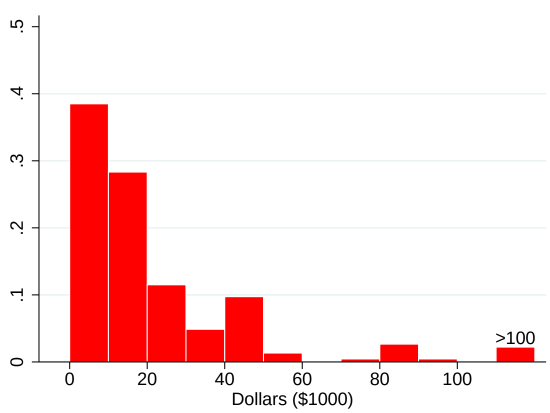

(a) LTCI Non-owners

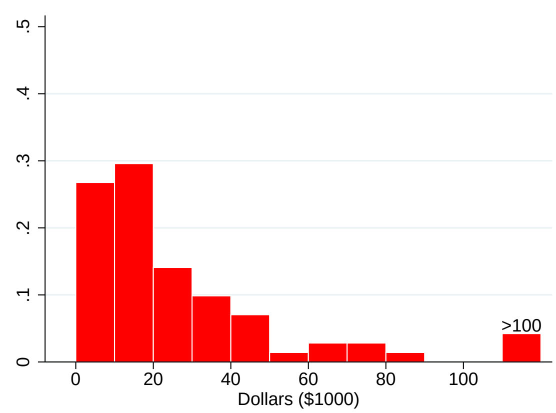

(b) LTCI Owners

Figure 14: 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.

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.[24]

Since the only people for whom we know the quantity of insurance owned against the LTC/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 model predicted demand are quite close, modeled demand is systematically larger on the intensive margin, as seen by comparison of Figures 9 and 14a. Table 5 documents the distributions of stated and model predicted demands for LTCI non-owners. Comparing the distributions of demand presented in rows 1 and 2 of Table 5, we observe that the mean, median, and all percentiles of model estimated ADLI demand distributions are at least as large as 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 model-predicted 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 model predicted 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 ownership of LTCI and the LTCI puzzle is driven by the low quality of LTCI products. Given that the extensive margin of stated and model predicted 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.

8.3 Predictors of the Estimated vs. Stated ADLI Demand Gap

In this section we analyze our model-bound 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 demands, 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

[24]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.

included in the model. Such omitted variables, denoted q, bias model estimates of demand from an individual's true demand.

Defining the difference between model-predicted and stated demand as

$$\eta_i := \text{Modeled}_i - \text{Stated}_i, \tag{8}$$

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 D.[25]

$$\begin{aligned} \eta_i &= \beta^x C_i^x + \beta^\Theta C_i^\Theta + \Gamma q_i + \epsilon_i \\ H_0 &: \ \beta^\Theta = 0; \ \beta^x = 0; \ \Gamma = 0. \end{aligned} \tag{9}$$

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, gender, 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 mis-specification 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.

Table 6 presents results from estimating Equation 9 on the sample of people who do not own LTCI with q 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 Nardi (2004) is not a fully adequate summary of the bonds between generations. Model enrichment to capture other family-related motives may be warranted.[26] 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 (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

[25] Note that the above specification ignores mis-specification caused by interaction of state variables and preferences. Attempts to control for these interaction effects through partial correlations of individual 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.

[26] 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.

	ADLI difference							
	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: Omitted Considerations, ADLI: This table presents the Γ coefficient from estimation of equation 9 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 D for discussion of β^x and β^Θ. Standard Errors are included in parentheses.

fair LTCI may be infeasible (Hendren (2013)). Variables such as real estate holdings, education, and the probability of receiving care from family, among others, do not significantly predict the difference.

9 Annuities and the Under-Insurance Puzzle

In this section we repeat the previous exercises for actuarially fair annuities. The annuity market is more developed than the market for LTCI products, and most individuals in our sample are familiar with them. Just as with ADLI, we use the model to calculate the implied annuity demands for the sample. Strikingly, all but two percent of respondents are modeled to purchase a strictly positive amount of an actuarially fair risk free annuity, much higher than the 59 percent of the population predicted to have positive ADLI demand.

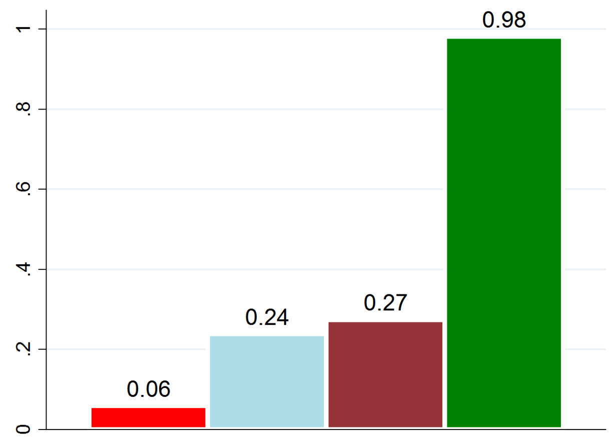

Figure 15: Fraction of Population Owning Private Annuities: This figure presents various measures of the fraction of the population with positive annuity ownership. Column 1 is actual holdings of a private annuities in the sample. Column 2 is stated demand. Column 3 is the union of private ownership and stated demand. Column 4 is model predicted demand.

Precautionary motives related to long-term care might explain lack of interest in annuities, but above some level of wealth and income, people have enough resources to be able to self insure against an expensive LTC spell using retirement income and a purchased annuity. Compared to the U.S. population, respondents of the VRI generally have high wealth as well as relatively high anticipated future income. Given their financial status and that their bequest motives are relatively weak, the model suggests that it is optimal to annuitize the bulk of their wealth.

We also collect stated annuity demand measures, the distribution of which is presented in Table 7. The direct stated demand questions concerning actuarially fair annuities specify an annuity as paying a fixed amount of income annually for remaining life. The hypothetical annuities for which demand is elicited are described as having no risk of default, being perfectly indexed for inflation, and as being fairly priced based on gender, age, and current health. In identifying respondent demand, it is specified that they pay a one-time, nonrefundable lump sum to purchase the annuity.

Despite being told explicitly that the offered annuity has no risk of default, is perfectly indexed for inflation, and is fairly priced, respondents reported limited interest in this product. Twenty four percent of respondents indicate that they would purchase some of this product, substantially more than the six percent of households who own annuities. This finding suggests some room for poor product features to explain low annuity ownership rates. Limited interest is also expressed by a modest-level of demand for the amount of annuity income. The 95th percentile of annuity demand is only $20,000. A regression of this demand on demographic correlates yields two highly significant findings.[27]

[27] The results of this estimation are presented in Appendix B.2.

	mean	p5	p10	p25	p50	p75	p90	p95
Modeled	47,082	2,987	5,540	15,146	33,708	65,926	108,981	148,414
Stated	2,968	0	0	0	0	0	10,000	20,000
Difference	44,110	1,022	4,166	12,962	31,338	64,112	106,967	140,561

Table 7: Distribution of Differences in Annuity Demand: This table presents the distribution of each of our annuity demand measures for those individuals that do not own a private annuity. The top line presents the model-predicted demand distribution, and the middle line presents the stated demand distribution. The bottom line presents the distribution of the differences between modeled and stated demand.

First, those with longer life expectancy are significantly more likely to have strictly positive demand than are those with lower life expectancy. As with ADLI, this points to possible adverse selection in the market for annuities. With respect to the extensive margin, among those who state a willingness to purchase, the quantity purchased increases strongly with wealth, as expected.

The distribution of stated annuity demand is dramatically different from model-predicted demand. Table 7 presents the demand distributions for both modeled and stated demands as well as the distribution of these differences. The table shows for actuarially fair annuities that the gap between what the model predicts individuals would demand and what individuals state they would purchase is massive. We observe that on average the model over-predicts annuity demand by more than $44K with a median over-prediction of over $30K. The model predicts most individuals should allocate most of their wealth to purchase a private annuity, while respondents state that they would generally only allocate a small share of wealth to such a purchase: almost all respondents stated demand at levels below 10 percent of their wealth. This illustrates the annuity puzzle in dramatic form, yet for a non-standard population. Figure 15 documents visually that the classic annuity puzzle is present in the VRI sample, with actual ownership and model predicted ownership drastically different.

The difference between stated-or-owned and model predicted annuity demand is radically larger than the difference in demands observed for ADLI. Given that stated demand is closer to actual annuity holdings and that both are much smaller than model predicted annuity demand, the annuity puzzle is likely not driven by differences in annuities available in the market and the modeled annuity product, but rather something missing in the model. This is in sharp contrast to ADLI and LTCI, in which a substantial part of the LTCI puzzle stems from the differences between LTCI products on sale and the preferred ADLI product.

10 Conclusion

Older Americans face many risks as they age. Foremost among these risks is needing assistance with activities of daily life 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 newly-created Vanguard Research Initiative to investigate the factors that low observed LTCI holdings reflect. The VRI includes batteries of questions that we designed to elicit the demand for insurance against late-in-life 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 Life 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 who the model predicts would buy ADLI, 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).

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 model predicted 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, we find that flaws in existing products do not fully explain it, especially for the highest wealth respondents. The VRI also includes a measure of whether respondents have LTCI. 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.

We also provide a more limited analysis of annuities based on stated demand for idealized annuities. Almost no VRI respondents have private annuities (though they have substantial retirement income through Social Security and private defined-benefit pensions). Unlike for idealized insurance against long-term care risk, the VRI respondents have little interested in idealized insurance against longevity risk. Hence, while this finding provides evidence that the lack of interest in annuities does not derive mainly from defects in the annuity market, it leaves the annuity puzzle largely unresolved.

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 in detail 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 (2015): "Long Term Care Utility and Late in Life Saving," *Vanguard Research Initiative Working Paper*.

——— (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.

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.

ATTANASIO, O. (2015): "The Determinants of Human Capital Formation During the Early Years of Life: Theory Measurement and Policies," *Journal of the European Economic Association*, 13(6), 949–997.

ATTANASIO, O., AND S. CATTAN (2015): "Inferring Beliefs about the Production Function," *working paper*.

ATTANASIO, O., F. CUNHA, AND P. JERVIS (2015): "Eliciting Beliefs about the Production Function," *working paper*.

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 (2016a): "Old, Sick, Alone and Poor: A Welfare Analysis of Old-Age Social Insurance Programs," *Review of Economic Studies*, forthcoming.

——— (2016b): "Private Long-Term Care Insurance: Why is the Market so Small and Coverage Denials so Frequent?," *working paper*.

Brown, J. (2007): "Rational and Behavioral Perspectives on the Role of Annuities in Retirement Planning," *NBER Working Papers*, 13537, 1–32.

Brown, J. R., and A. Finkelstein (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 (2016): "Cognitive Constraints on Valuing Annuities," *Journal of the European Economic Association*, forthcoming.

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., E. French, and J. Jones (2010): "Differential Mortality, Uncertain Medical Expenses, and the Saving of Elderly Singles," *Journal of Political Economy*, 118, 49–75.

Delavande, A., and S. Rohwedder (2008): "Eliciting Subjective Probabilities in Internet Surveys," *Public Opinion Quarterly*, 72(5), 866–891.

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

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

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

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.

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. (2015): “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. (1989): “Mortality Risk and Bequests,” *Econometrica*, 57(4), 779–813.

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

JUSTER, F. T. (1966): “Consumer Buying Intentions and Purchase Probability: An Experiment in Survey Design,” *Journal of the American Statistical Association*, 61(315), 658–696.

KESTERNICH, I., F. HEISS, D. MCFADDEN, AND J. WINTER (2013): “Suit the Action to the Word, the Word to the Action: Hpyothetical Choices and Real Decisions in Medicare Part D,” *Journal of Health Economics*, 32(6), 1313–1324.

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.

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

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.

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

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

——— (2016): "Incidental Bequests: Bequest Motives and the Choice to Self-Insure Late-Life Risks," *working paper*.

MANSKI, C. F. (1990): "The Use of Intentions Data to Predict Behavior: A Best Case Analysis," *Journal of the American Statistical Association*, 85(412), 934–940.

——— (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.

MODIGLIANI, F. (1986): "Life Cycle, Individual Thrift and the Wealth of Nations," *American Economic Review*, 76(3), 297–313.

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

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.

PAWEENAWAT, A., AND R. M. TOWNSEND (2012): "Village Economic Accounts: Real and Financial Intertwined," *The American Economic Review Papers and Proceedings*, 102(3), 441–446.

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.

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

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

WISWALL, M., AND B. ZAFAR (2015): "Determinants of College Major Choice: Identification using an Information Experiment," *Review of Economic Studies*, 82(2), 791–824.

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

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

Appendices

A External Estimates

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.[28]

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 gender, g. The HRS only records 2 year health state transitions which we use to identify the one-year transition probabilities in a manner similar to De Nardi, French, and Jones (2010). To do this, we write the two year transition probabilities as:

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

where,

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

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

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

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, gender, health state, and interaction terms. Using the residuals from this first regression, we

[28]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.

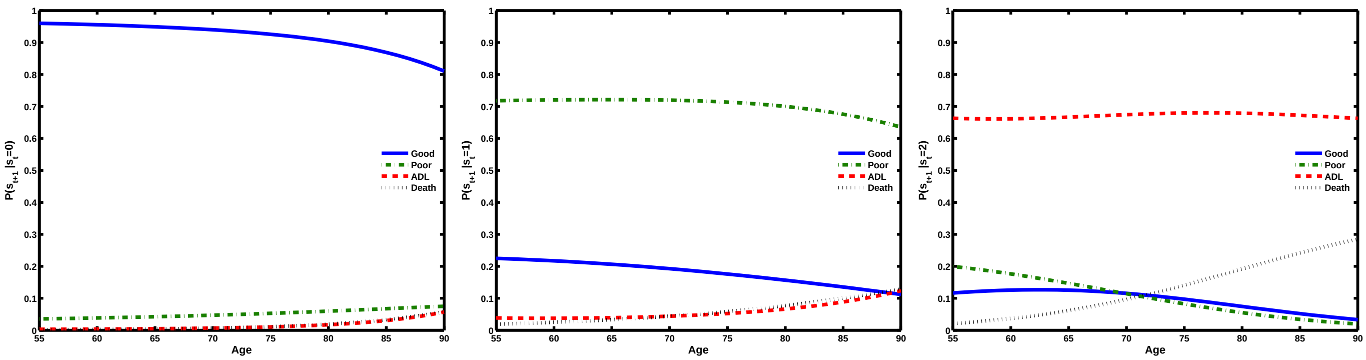

(a) Healthy **(b)** Sick **(c)** ADL

Figure A.1: Male Health State Transition Profile

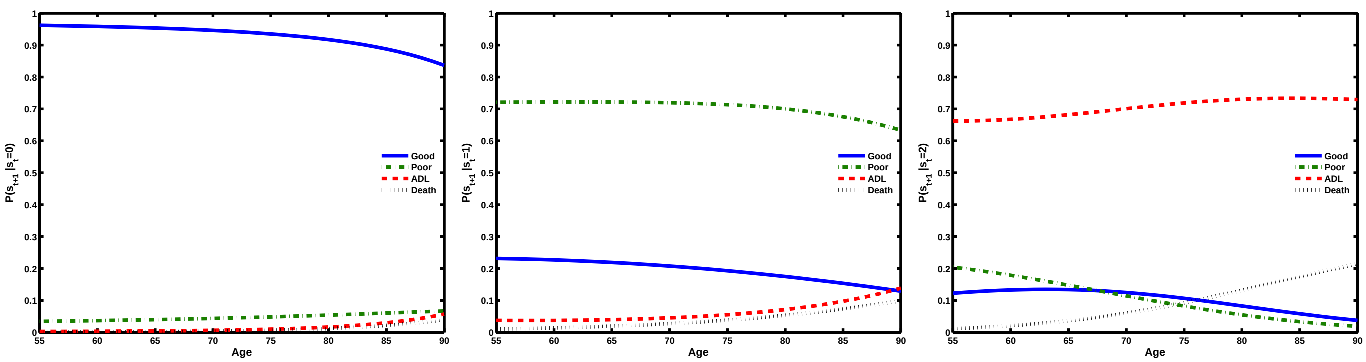

(a) Healthy **(b)** Sick **(c)** ADL

Figure A.2: Female Health State Transition Profile

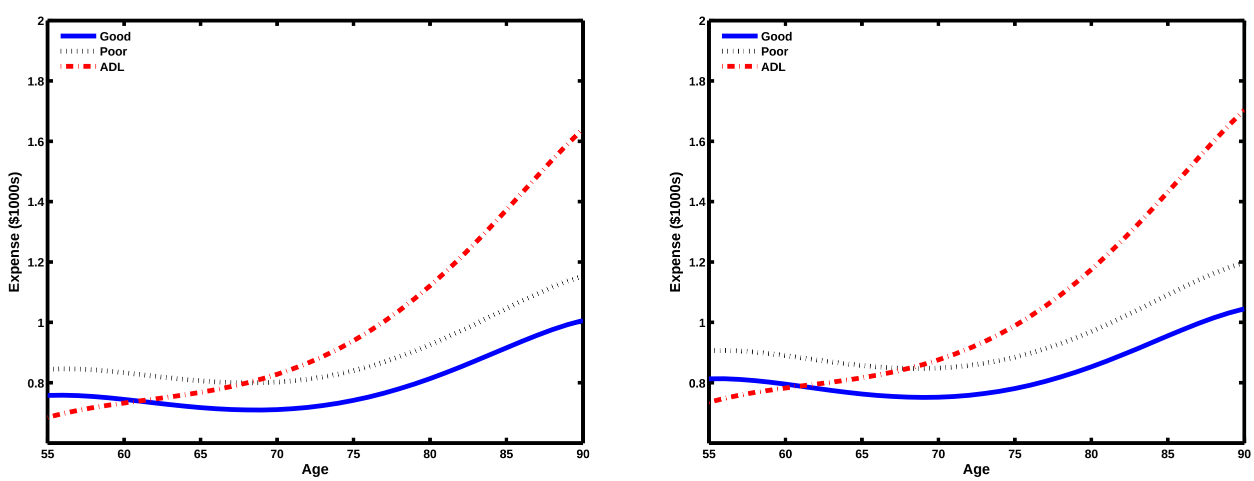

(a) Male **(b)** Female

Figure A.3: Median Health Cost Profile

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

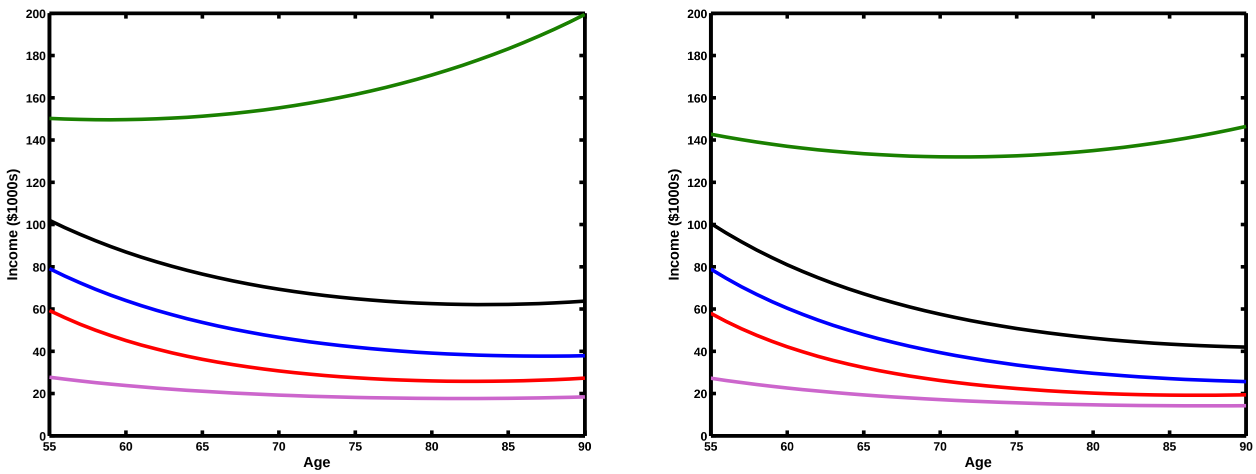

(a) Male **(b)** Female

Figure A.4: Income Profile Quintiles

B Validating Survey Instruments

This appendix presents results from validation exercises for key survey instruments. It first considers SSQ credibility beyond the summary provided in Section 5.3. It then examines how SSQ responses correlate with other variables before doing the same for stated ADLI and annuity demands.

B.1 SSQ Credibility

In this section we expand upon three types of credibility analysis summarized in Section 5.3. First, we present results of 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.

B.1.1 Comprehension Tests

The survey includes direct comprehension tests that respondents are asked to answer at most twice. Performance on these comprehension tests is summarized in Table B.1. For SSQ 1, which was discussed in the main text, there were 6 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 B.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.

B.1.2 Respondent Feedback and SSQ Design

The SSQ design process incorporates several forms of feedback that provided us with opportunities to improve the survey prior to fielding to the production sample. In addition to survey design feedback obtained as a result of cognitive interviews, we also gathered feedback from scripted live chat pop-up interviews with a subset of the pilot sample. The live chats provide feedback in free response form on issues that may trouble respondents. In addition to asking respondents for their overall reactions to the survey, we posed specific questions about each SSQ, with broadly encouraging and informative results.

Additionally, as summarized in Section5.3, a subset of the live chat questions were posed to the full production sample at the end of the survey. The full text and tabulated answers to these questions are included in Table B.2. We

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 B.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.

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 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.

B.1.3 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 B.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 B.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.

B.2 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 B.4 through B.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 gender, 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 gender.

	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 B.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 B.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 B.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 B.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.

B.3 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 B.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 B.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.

Stated Annuity Demand Table B.9 presents results from regressions of demographic and other covariates on stated demand for annuities. The first column presents a probit regression on an indicator equal to one if the respondent reported they would purchase a positive amount of the offered annuity. We observe that respondents who report higher probabilities of living 10-20 years are more likely to purchase the offered annuity.

In the second column we present an OLS regression of the amount of annual annuity income that respondents state they would purchase in the subsample of respondents that reported they would purchase positive amounts. Both measures of stated interest correlate in generally reasonable manners with economic and demographic characteristics.

	$\mathbb{I}_{Ann>0}$	**Annual Income**
Subj. Prob of Survival (Above Median)	.22	-293.77
	(.09)	(1578.73)
Age	.82	14249.78
	(1.01)	(16812.93)
Age^2	-.01	-209.82
	(.01)	(241.22)
Age^3	5.50E-05	1.02
	6.80E-05	(1.14)
Health: Poor	.23	720.12
	(.22)	(3532.51)
Health: Poor	-.08	3807.43
	(.57)	(11508.28)
Income Quintile: 2	.11	-3466.09
	(.15)	(2524.96)
Income Quintile: 3	.06	708.25
	(.15)	(2569.15)
Income Quintile: 4	.13	-265.64
	(.15)	(2547.29)
Income Quintile: 5	.07	2357.49
	(.16)	(2608.58)
Female	.18	2884.34
	(.09)	(1492.53)
College or Higher	-.04	-1801.32
	(.11)	(1816.60)
Log(Wealth)	.11	2170.86
	(.05)	(934.27)
N	1016	226

Table B.9: Correlates of Surveyed Annuity Demand Measure: This table presents how stated annuity demand is predicted by covariates. Column 1 presents the results of a probit regression of the annuity purchase decisions, and Column 2 presents an OLS regression on the level of annual payout demanded for those with positive demand. Our measure of longevity is whether an individual's expectation on the probability of living for 10-20 years is above or below median, conditional on current age.

C Robustness

	%>0	Mean	p5	p10	p25	p50	p75	p90	p95
Baseline	58	67,264	8,845	13,108	29,070	54,740	93,449	148,143	180,556
Alt. Estimates									
r=.03	54	65,764	8,094	14,623	29,057	53,962	88,995	141,207	178,275
10% load	55	66,254	7,930	12,720	30,104	53,257	88,654	149,928	175,342
20% load	54	62,962	7,255	13,227	29,512	51,212	83,720	133,913	165,648
30% load	54	59,923	6,773	12,505	28,257	50,377	82,270	126,423	155,934
Housing Wealth	62	77,545	9,683	16,677	34,923	62,749	103,619	166,392	214,944
No Min. Expenditure	54	70,712	7,946	12,891	33,537	58,492	94,365	164,076	181,407
Multiplicative Errors	71	56,563	8,551	16,648	28,157	43,247	74,715	107,812	153,489
Population Parameters	83	96,374	62,289	68,286	75,053	85,511	109,235	149,240	177,582
Alt. Subsamples									
Employer Subsample	49	50,285	8,015	8,885	20,609	41,617	70,466	102,391	141,093
HRS weights	46	47,696	6,043	6,812	10,833	34,246	65,223	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 C.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.

	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**			**Gender**	
	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 C.2: Summary Statistics on Wealth, Income, Health, Age, and Gender—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.

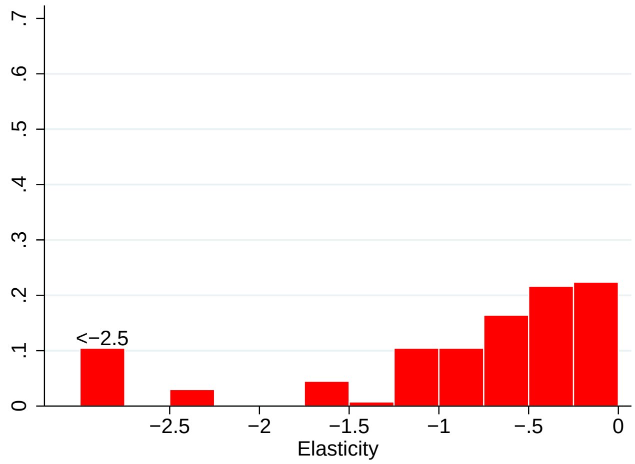

Figure C.1: Distribution of the Price Elasticity of Demand: This figure presents the histogram of the elasticity of demand with respect to price for those who own LTCI. It plots the distribution of the percent change in demand to a one percent increase in price, local to the optimal demand level and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

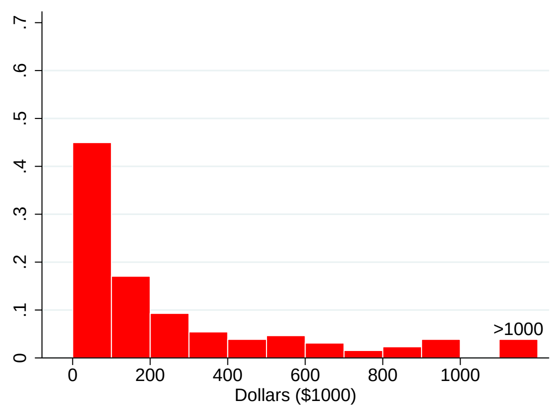

(a) Distribution of Consumer Surplus

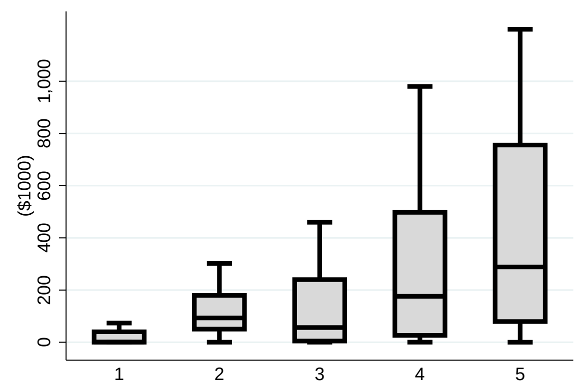

(b) Consumer Surplus Box Plot by Wealth Quintile

Figure C.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.

D Exploring the Model Predicted 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 mis-specification. 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 mis-specification of included model elements as dictated by Θ and x. For example, mis-specification 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 gender 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 mis-specification 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

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

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 mis-specification (relative to stated demands) related to the relevant variables. Thus, $g_x(x) \neq 0$ suggests model mis-specification related to modeled state variables, $g_\Theta(\Theta) \neq 0$ suggests model mis-specification related to preference parameters, and $g_q(q) \neq 0$ suggests model mis-specification related to omitted variables q. Furthermore, $g > 0$ suggests mis-specification that causes model demand to be overstated relative to stated demand, while $g < 0$ suggest mis-specification 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 D.1 yields

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

which we use to estimate

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

Equation D.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.[29] 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.

[29] As mentioned when discussing equation D.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.