Risky Insurance: Life-cycle Insurance Portfolio Choice with Incomplete Markets*

Joseph Briggs Ciaran Rogers Christopher Tonetti
Goldman Sachs HEC Paris Stanford GSB and NBER

September 2025 (Version 0.95)
Preliminary Results

[Link to most updated version]

We provide survey evidence that individuals believe there is substantial nonpayment risk in annuity, life insurance, and long-term care insurance (LTCI) products. Using simple statistical analysis we show that nonpayment beliefs predict insurance ownership and that the insurance ownership rate would be much larger if people believed there was zero nonpayment risk. To better quantify how nonpayment risk affects insurance ownership and how different features of insurance products affect consumer welfare, we develop an incomplete-markets life-cycle model of the demand for life insurance, annuities, LTCI, and a risk free bond. We incorporate features of real-world insurance products such as perceived nonpayment risk, high loads above actuarial fair prices, and quantity restrictions (e.g., age restrictions on purchases, short-selling constraints). Both high prices and nonpayment risk substantially decrease insurance ownership. Compared to our rational expectations baseline, the welfare loss from suboptimally owning zero insurance is 0.7 percent in consumption equivalent units. If the products had no risk and were sold at actuarially fair prices, the welfare cost of zero insurance ownership is much larger at 4.0 percent.

*This research received much appreciated financial support from the Alfred P. Sloan Foundation through the NBER Household Finance Small Grant program.

1 Introduction

A large number of financial products are available that directly address many of the life-cycle risks Americans face. Life insurance insures an individual against early death and provides a vehicle to transfer money to an individual's estate. Annuities insure an individual against outliving their wealth by providing an income stream that is guaranteed for as long as the insured individual is alive. Long-term-care insurance (LTCI) provides coverage of expensive health and care costs should an individual need assistance with the activities of daily living or require a stay in a nursing home. Furthermore, the products that are available in the market have a variety of maturities and payoff structures, making it possible for consumers to choose a customized portfolio of annuity, life insurance, and LTC insurance policies. Despite a rich set of available insurance options, most individuals choose not to participate in these insurance markets: only 6% of individuals aged 60-65 own a private annuity, 30% own a private life insurance policy, and 10% own a private LTC insurance policy.

Given this puzzling lack of insurance, a number of studies have attempted to understand consumer demand for annuities (Brown (2001), Inkmann et al. (2011)), life insurance (Bernheim (1989), Inkmann and Michaelides (2012), Hong and Rìos-Rull (2012)), and long-term care insurance (Brown and Finkelstein (2008), Lockwood (2012), Ameriks et al. (2018), Mommaerts (2015)). Most closely related to our paper, Koijen et al. (2016) highlighted the importance of considering the demand for these three products jointly to account for overlap in payouts across states. In a model with complete markets, Koijen et al. (2016) develop two risk measures that define necessary conditions for an optimal insurance portfolio, characterize the optimal life-cycle consumption path, and estimate substantial welfare losses due to households holding sub-optimal insurance portfolios.

In this paper we develop both a new dataset and a new model to better understand consumer demand for annuities, life insurance, and LTC insurance. We first design a new survey to measure the expected distribution of insurance policy payments. We then incorporate these measures into an incomplete markets life-cycle model of demand for life insurance, annuities, LTC insurance, and risk free bonds. The model features uninsurable health and mortality risk and the insurance products have uninsurable non-payment risk, cost more than the acturially fair price, have wedges between purchase and selling prices, are subject to short-selling constraints, and there is a maximum age for new policy issuance.

Our study makes several contributions that advance understanding of consumer insurance demand. First, we show that perceived nonpayment risk is substantial in all three of the insurance markets we study. Our elicited distributions of expected payouts imply that on average consumers expect to receive 87.2% of the promised payout from a typical

life insurance policy, 81.5% for annuities, and 76.2% for LTC insurance, with standard deviations of payouts averaging 16.7%, 21.5% and 23.9%, respectively. In addition, the full default probability—defined as the probability that a policy becomes worthless and makes no subsequent payment prior to the policy's contractual expiration—is 0.23 for life insurance, 0.29 for annuities, and 0.32 for LTC insurance. Our three risk measures all indicate that life insurance is perceived as the most safe and that LTC insurance is the most risky. Consumers perceive much higher risk in insurance markets than what would be inferred from historical observations of insurer default, possibly suggesting high paperwork or contractual complications as the main barrier to full payment.

Our second key contribution is a statistical analysis of subjective beliefs and insurance ownership. Individual's perceived risks are highly correlated with actual insurance product holdings. For example, respondents with higher expected payouts are significantly more likely to own life insurance, respondents reporting higher standard deviations of annual payments and higher full default probabilities are significantly less likely to own annuities, and respondents with higher full default probabilities are significantly less likely to own LTC insurance. While the direction of causality can not be fully determined from such cross-sectional analyses, the negative relationship between perceived risk and policy ownership are robust to including demographic controls and behavioral measures (e.g., proxies for trust, cognitive ability, risk aversion, propensity to plan, and financial literacy) and overall suggest that perceived risk reduces policy ownership. In fact, counterfactual predictions of probit regressions indicate that annuity and LTC insurance ownership rates would be twice as high if households were to perceive insurance products as risk-free, while life insurance ownership rates would be over nine percentage points higher.

Our third key contribution is the structural model analysis. To better quantify how non-payment risk affects insurance ownership and how different features of insurance products affect consumer welfare, we develop an incomplete-markets life-cycle model of the demand for life insurance, annuities, LTCI, and a risk free bond. Health and mortality are the key risks in the model. We incorporate features of real-world insurance products such as perceived nonpayment risk, high loads above actuarial fair prices, wedges between the purchase and resale prices for insurance products, and quantity restrictions (e.g., age restrictions on purchases, short-selling constraints). We employ a mix of calibration and estimation so that insurance in the model reflects products actually available in the market. This paper is the first of which we are aware to study the joint demand for annuities, life, and LTC insurance in an incomplete market model that reflects the decision problem households face as they age into retirement.

To quantify the model we need to select values for parameters that control properties of

insurance products, health and mortality risk, income, preferences, and rates of return. We use the estimated nonhomothetic health-state-dependent utility function from Ameriks et al. (2020) with one extension: we allow the strength of the bequest motive to vary over the life cycle. We also follow that paper in estimating income age profiles and health and mortality risk using Health and Retirement Study (HRS) data, but in this paper we work with a population that is representative of older Americans. We calibrate loads above actuarially fair prices on insurance products to those reported in Hong and Rìos-Rull (2012) for life insurance and Brown and Finkelstein (2008) for annuities and LTC insurance. Finally, we choose rates of return on risk free bonds, the discount rate used to determine actuarial fair prices, the household discount factor and intertemporal elasticity of substitution, and the a parameter that controls the life-cycle profile of bequest motives to match the 25th, 50th, and 75th percentiles of the wealth distribution by age and the average ownership rate of the 3 insurance products.

The model predicts insurance holdings in line with those observed empirically: consumer ownership of annuities and long-term care insurance is very low, and ownership of life insurance is far from universal. Furthermore, we match the life-cycle pattern of insurance ownership, which is not a targeted moment. Both high prices and nonpayment risk substantially decrease insurance ownership. Compared to our baseline, the welfare loss from sub-optimally owning zero insurance is 0.4 percent in consumption equivalent units. If the products had no risk and were sold at actuarially fair prices, the welfare cost of zero insurance ownership is much larger at 7.9 percent. This large welfare cost of suboptimal insurance ownership is reminiscent of the complete markets analysis in Koijen et al. (2016). Thus, we conclude that the welfare costs of suboptimally low insurance ownership is small given the products available in the market and, simultaneously, risky and incomplete insurance markets result in significant welfare losses to older consumers.

Because many of our main findings are derived from survey responses, we take a number of steps to establish the credibility of our key measures and findings. In addition to insurance product ownership, we find that our subjective expectations of nonpayment risk correlate with other behavioral measures in intuitive ways. For example, respondents that have experienced financial fraud in the past report higher probabilities of full default. In addition, several patterns in the data suggest that respondents were thoughtful when answering the survey. For example, respondents expect lower payouts from annuities and life insurance during economic downturns, while expectations of long-term care insurance payouts are invariant to aggregate fluctuations. Finally, we included a number of randomization and internal consistency checks in our survey design. We find that randomized framings and orderings do not predict survey responses, while approximately 85% of respondents ranked

the riskiness of insurance products in a manner consistent with their perceived risks.[1]

1.1 Relation to the Literature

In terms of measurement, our paper is most similar to Luttmer and Samwick (2018), which measures the perceived nonpayment risk surrounding Social Security. Our key survey instrument is derived from that used in this study, and applies the bins-and-balls method (Delavande and Rohwedder (2008)) to elicit the subjective payment distribution for each respondent. We also follow Luttmer and Samwick (2018) and measure indifference points between insurance policies currently available and a policy that offers certain, but lower, payouts. These certainty equivalent measures imply risk premia of 5.7% for life insurance, 7.7% for annuities, and 3.3% for LTC insurance, which are comparable to the 8% risk premia Luttmer and Samwick (2018) estimate for Social Security. Although similar in measurement methodology, our study focuses on more and different types of assets, thus providing additional insights into the perception of financial risks in retirement.

In terms of the model and research question, our paper is most similar to Koijen et al. (2016), which provides the key insight that demand for annuities, life, and LTC insurance should be studied jointly given the overlap in payouts from these products. We compliment that paper's complete-market framework by characterizing optimal insurance holdings in an incomplete market setting, and show that this extension has significant implications for the welfare costs of observed household insurance holdings. Whereas Koijen et al. (2016) find substantial welfare costs from sub-optimal portfolio choice we find that foregone gains are much smaller after modeling the products that consumers perceive as being available in the market.

In addition to these two studies, this paper contributes to a number of broader literatures. First, we relate to papers that study the demand for annuities (Yaari (1965), Brown (2001), Davidoff et al. (2005), Inkmann et al. (2011)), life insurance (Bernheim (1989), Inkmann and Michaelides (2012), Hong and Rìos-Rull (2012)), and long-term care insurance Brown and Finkelstein (2008), Lockwood (2012), Ameriks et al. (2018), Mommaerts (2015). In addition, a few studies have previously considered the joint demand for annuities and life insurance (Inkmann et al. (2011), Hubener et al. (2013)), while the aforementioned Koijen et al. (2016) study considers demand for all three insurance products. We contribute to this literature by characterizing joint demand for insurance products in a model with incomplete markets and by incorporating nonpayment risk.

Our study also contributes to the literature that studies the determinants of late-in-

[1] All results are robust to restricting our sample to households that pass this internal consistency check.

life saving behavior. As noted in Poterba et al. (2013), households spend down their wealth slowly in retirement. A number of studies have tried to explain this slow draw-down by allowing for uninsurable medical expense risk (Hubbard et al. (1994), Palumbo (1999), Anderson et al. (2004), De Nardi et al. (2010), uninsurable LTC spending risk (Kopecky and Koreshkova (2014), Lockwood (2016), Ameriks et al. (2020)) and a strong bequest motive (Hurd (1989), Nardi (2004), Kopczuk and Lupton (2007), Love et al. (2009), Lockwood (2012)) in incomplete market models. The model that we use is similar to the models used in these studies and incorporates many of the key features that generate slow wealth decumulation in retirement.

We also relate to papers that measure and model how uninsurable risk and incomplete markets affect consumer behavior. Studies of uninsurable labor income have been shown to have significant impact on asset prices (Krusell and Smith (1998), Gomes and Michaelides (2008), Storesletten et al. (2007), Guvenen (2009)) and consumption inequality (Blundell et al. (2008)). At an individual level, higher exposure to uninsurable background risk has been shown to make consumers less willing to invest in equities (Curcuru et al. (2004), Fagereng et al. (2016b)). We contribute to this literature by demonstrating that incomplete insurance markets have significant affects on consumer choices and welfare, thus highlighting the importance of measuring and modeling these features in a new domain.

We additionally relate to a number of studies that have examined how trust and counterparty risk affect demand for assets. In the case of annuities, several studies have noted that small probabilities of nonpayment can make individuals unwilling to participate in annuity markets (Lopes and Michaelides (2007), Pashchenko (2013), Jang et al. (2013)), but there is no consensus as to whether this is a significant or insignificant channel for limiting annuity demand. We address this question by (1) showing that perceived nonpayment risk in annuity markets is large and (2) it significantly limits annuity ownership. In equity markets, a small perceived probability of nonpayment or loss of capital is widely acknowledged as a cause of nonparticipation. Guiso et al. (2008) show evidence that more trusting individuals are willing to participate in equity markets and Fagereng et al. (2016a) demonstrate that this channel is important for matching empirical portfolio holdings with a life-cycle model. Our study demonstrates that similar perceived risks are important in explaining consumer behavior in insurance markets as well.

Our study also contributes to an active literature that studies how subjective beliefs regarding investments and financial decisions are correlated with actual consumer behavior (see Manski (2004) for a survey on measuring beliefs and their importance in explaining behavioral observations). Most directly comparable, Beshears et al. (2014) find that perceived risks and investments limit annuity purchases. Additionally, a number of recent studies have

established that subjective beliefs of house price returns significantly predict home purchase decision (Armona et al. (2018)), while Adelino et al. (2018) show that the second moment of the expected return distribution is highly predictive as well. Additionally, expectations of returns have been shown to affect investment in the stock market (Hurd et al. (2011)), human capital (Wiswall and Zafar (2017)), and labor market search effort (Conlon et al. (2018)). Our finding that expectations of insurance product returns predict ownership contributes to the mounting evidence that subjective expectations are an important determinant of consumer behavior.

Finally, we tangentially relate to an active literature that examines the investments and financial stability of insurance providers (Merrill et al. (2012), Becker and Opp (2013), Becker and Ivashina (2015), Ellul et al. (2015)). Recent work on life insurance providers (Koijen and Yogo (2015), Koijen and Yogo (2016)) suggests that the growth of shadow insurance is a growing risk in the stability of these market, suggesting that the perceived counterparty risk consumers report might not be without merit. Although we do not model the supply side of insurance markets and treat nonpayment risk as exogenous, this study provides new measures of consumer beliefs of nonpayment risks from insurance providers that potentially provide new insights into sources of financial instability.

The remainder of the paper is structured as follows: Section 2 introduces the data set, including documenting patterns in observed insurance product holdings. Section 3 outlines our survey, while Section 4 analyzes survey responses and examines their relationship with product ownership and other individual characteristics. Section 5 describes the baseline structural model and details our quantification approach. Section 6 analyzes predictions from our structural model and calculates the welfare gains associated with insurance products, under scenarios that vary consumer beliefs and properties of the insurance products. Finally, Section 7 concludes.

2 Data

We use two data sets that play distinct roles in our analysis. The first is the Understanding America Study (UAS), a nationally representative internet panel of approximately 6,000 adults hosted by the USC Dornsife Center. Our working sample comprises 1,040 respondents aged 45 and over who completed our purpose-designed survey on insurance products in May-June 2018 (see Section 3). The second is the Health and Retirement Study (HRS), a nationally representative panel of more than 37,000 Americans over age 50 with rich measures of health, income, assets, and insurance from 1991 to the present.

We use the two sources in complementary ways. The UAS provides *initial state variables*

and beliefs - age, sex, income, liquid and housing wealth, and observed insurance ownership - to set initial state variables and to measure perceptions of nonpayment risk and their cross-sectional correlates. The HRS provides *inputs and targets*: we calibrate the stochastic processes for health and longevity, health cost shocks, and income age profiles from the HRS, and we target HRS moments for the wealth distribution and life-cycle insurance participation in our model simulations. Our primary empirical analysis is on our UAS sample, but HRS allows us to document key facts and estimate processes using a widely understand data set.

After appropriate reweighting, both the UAS and HRS samples are representative of older Americans, so we combine measures across sources without significant external-validity concerns. We first present key facts on insurance ownership from the HRS, then turn to the UAS and describe our survey and constructed variables in detail.

2.1 HRS Sample and Dynamics of Insurance Ownership

We use the HRS to (i) estimate health, income and medical expenses that inputs for our model, and (ii) estimate wealth and insurance ownership dynamics that are moments targeted in calibrating our structural model described in Section 5. Our estimation of these profiles and processes follow literature standards, so methodological details are relegated to Appendix C.

Figure 1 plots insurance participation rates by age in 2015 from the HRS. Two patterns stand out. First, late-life ownership of annuities and LTCI is low: at age 65, roughly 12% of households hold LTCI and fewer than 5% hold a private annuity, consistent with the classic annuity puzzle (Yaari (1965), Modigliani (1986)) and LTCI puzzle (Ameriks et al. (2018); Braun et al. (2016)). Second, LI participation is much higher - about 60% at age 65 - but declines steadily with age, consistent with earlier term policies purchased to insure children's welfare rolling over time (Hong and Rìos-Rull (2012)). Overall, Figure 1 indicates limited demand by older Americans for the annuity or LTC insurance products available in the market.

2.2 UAS Sample and Beliefs

The Understanding America Study (UAS) is an internet panel run by the USC Dornsife Center. Researchers field survey modules to the panel, and respondents are randomly invited to join (so there's no voluntary self-selection). Before fielding, the UAS team programs and tests all modules, and respondents are compensated. The UAS also collects general-purpose measures: every two years, all participants complete a web-adapted HRS survey, providing harmonized data on health, work, finances, and behavioral traits. After a one-year embargo,

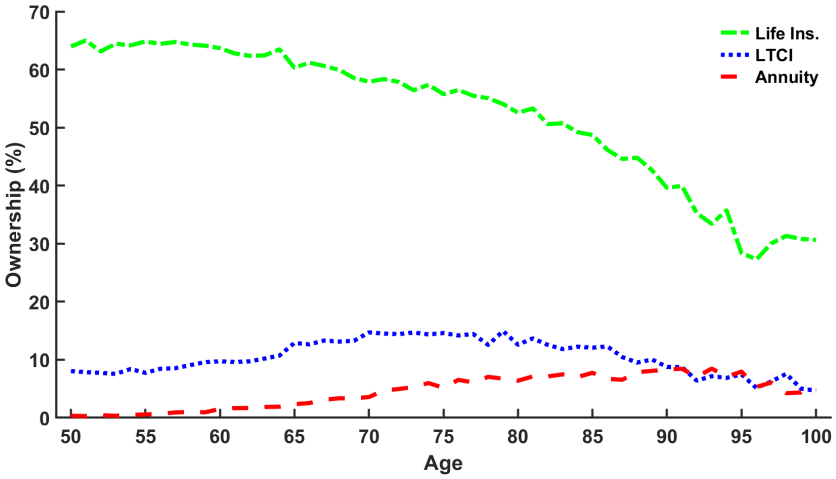

Figure 1: **Age profiles of ownership rates for annuities, life insurance, and long-term care insurance in the HRS.**

modules are publicly available and can be linked across time via unique respondent IDs. More information is at `https://uasdata.usc.edu/`.

Our survey was fielded in May 2018. Among invited panelists, the response rate was 82%[2]. In total, 1,104 started and 1,088 completed the survey; average completion time was 17 minutes. Item non-response was minimal: 1,055 answered all life insurance items, 1,046 all annuity items, and 1,040 all LTC items. To maximize sample size, we do not restrict analyses to a common respondent set across products.

Summary statistics for our UAS sample are in Column (2) of Table 1, and compared to the UAS sample in Column (1). Some differences versus the HRS are mechanical (English-only sampling makes our sample less Hispanic and more White). Others are substantive (our respondents are wealthier with higher incomes), but importantly, insurance ownership rates are similar across data sets.

We link our module to prior UAS measures using fixed respondent IDs. Standard demographics, such as age, gender, marital status and labor market status, are preloaded; we also merge behavioral covariates - trust, cognition, financial literacy, numeracy, recent fraud experience, subjective risk aversion, and propensity to plan - that are collected by linking respondent IDs to answers on prior UAS surveys. Finally, we construct cohort-specific stock-return exposure from ages 18-25 using Bob Shiller's historical returns, merged by birth year. These variables appear as controls in selected specifications alongside standard demographics.

[2]Lower bound; the link closed when the purchased quota was met.

Table 1: **HRS vs. UAS Samples: Summary Statistics.** We report mean values for each variable. Column (1) is HRS; Column (2) is UAS.

	HRS (1)	**UAS (2)**
Demographics		
Male	0.46	0.47
Age	66.9	61.1
Married	0.64	0.57
Education		
High school	0.64	0.68
Some college	0.23	0.18
College & above	0.13	0.14
Race		
White	0.83	0.81
Black	0.10	0.11
Hispanic & other	0.07	0.08
Health		
Good	0.74	0.79
Bad	0.23	0.17
LTC	0.03	0.04
Income (K)	69.7	100.6
Wealth (K)	582	461
Insurance ownership		
Annuity	0.04	0.12
Life insurance	0.61	0.58
LTCI	0.12	0.10
N	18747	1085

3 Survey Description

This section describes the UAS survey we fielded to measure beliefs about insurance non-payment risk. Our key survey instruments are motivated by those developed in Luttmer and Samwick (2018), who first measure the perceived distribution of Social Security benefits and the certainty equivalent to that distribution for each survey respondent. The difference of the expected value of payouts and the certainty equivalent measure yields the risk premia, i.e., the compensation an individual would require to hold the risk associated with payment uncertainty. While we adopt the Luttmer and Samwick (2018) baseline survey instrument, we tailor it in a number of ways to better inform risks in insurance markets. In this section, we provide details of our survey instrument.

We organise Section 3 as follows. Section 3.1 focuses on our three core measures: (i) the annual nonpayment probability, (ii) the full default probability, and (iii) a certainty-equivalent (CE) valuation measure. Section 3.2 shows how to evaluate conditional beliefs, while Section Section 3.3 provides additional implementation details.

3.1 Measuring Nonpayment Risk and Certainty Equivalence

Survey Design Our survey consists of three modules that ask each respondent about perceived risk for each insurance product (annuity, life, and LTC). To avoid and permit testing for anchoring of responses, we randomize the order in which respondents receive modules. Within each module, respondents are first asked whether they are familiar with or own that specific type of insurance policy. For those that own an insurance policy, we collect details of the specific policy that they own (e.g., policy size) and ask them to consider this policy when reporting nonpayment risk. For those that do not own an insurance policy of the type measured in each module, we provide them details of how the indicated insurance works, prompt them with a policy of an indicated benefit size (with benefit sizes randomly determined), and ask them to imagine the best policy of this type that they think they could buy in the market today. Having fixed a policy in mind we then ask them about their perceptions of risk associated with this policy. To provide a concrete example of our survey instrument, for the remainder of this section we will focus on annuities, with survey modules for life and LTC insurance implemented analogously.

Full Default We first ask respondents about the probability of full default. Specifically, we ask the following:

> Suppose that you own an annuity that promises to pay $[5000] each year for the rest of your life. Suppose further that you never trade this annuity for cash and hold the contract until the end of your life.
>
> We are now interested in the percent chance that the annuity becomes worthless due to no fault of your own at any point before the end of your life. This means that the annuity permanently stops making payments. This might occur if the insurance company goes out of business, they claim you violated a clause in the contract, or they ruled the policy void for some other reason.
>
> What is the percent chance this occurs?

Respondents than indicate their responses on a slider which updates to remind them what their choice indicates as they move between zero and one hundred.

Annual Payments We then proceed to collect the perceived distribution of payments in a year where the respondent qualifies for a claim. We use the bins-and-balls technique

developed in Delavande and Rohwedder (2008) that prompts respondents to place balls in categories based on the likelihood indicated events will occur. In the first screen, respondents are asked to indicate whether they will receive no payment, a fraction of their promised payment, or the full payment promised in their contract:

> Suppose that you own an annuity that promises to pay $[20,000] each year for the rest of your life. We would now like to focus on what might happen just during the next calendar year.
>
> You have been given 20 balls to put in the following bins. Each bin describes a scenario that involves the annuity payment that you are supposed to receive next year. The more likely you think a bin is, the more balls you should put in that bin.
>
> What do you think will happen to the annuity payment next year?
>
> I will receive no payment/I will receive a payment less than I am supposed to receive/I will receive a payment at least as large as I am supposed to receive

Respondents are then asked to reassign balls placed in the middle bin into 5 bins ([1-20, ... , 81-99]) to indicate the fraction of the payment they expect to receive conditional on receiving less than contractually promised.

> You put [13] ball(s) in the bin marked "I will receive a payment less than I am supposed to receive." Please distribute those balls in the following bins. The more likely you think a bin is, the more balls you should put in that bin.
>
> If you do receive a payment that is less than you are supposed to receive, how much do you think you would get?

This provides a complete characterization of the distribution of payments respondents expect to receive, which we can then use to calculate the expected value and standard deviation of expected annual payouts as a fraction of the amount promised in the contract.[3]

[3]In calculating expected values and standard deviations we assume that balls correspond to the midpoint of their final bins.

Certainty Equivalent After collecting the distribution of payments, we ask respondents whether they would be willing to accept a risk-free contract that offers certain payments of a value below those contractually promised. Each respondent (that reported at least some nonpayment risk) makes this choice once for a risk-free contract that either promises 25%, 50%, 75%, 80%, or 90% of the promised payout of the policy under consideration, with the fraction randomly determined. See Appendix A for an example of how the survey question is framed.

Because each respondent only responds to this question once we do not collect precise certainty equivalent measure for each respondent, and therefore do not (as in Luttmer and Samwick (2018)) recover the distribution of risk-premia among respondents. However, the fraction of people that would accept the risk-free contract for each level of payout gives the population CDF of certainty equivalent measures that can then be used to calculate summary statistics like mean and median certainty equivalent measures, and by extension, the mean and median of the population risk-premia.

3.2 Additional Measures of Nonpayment Risk

In addition to the core measures described above, our survey also collects a number of additional nonpayment risk measures in a supplemental survey module. Each respondent completes a supplement corresponding to one insurance product, with the supplement they received randomly determined.[4] These additional measures, which are described below, are designed to provide additional insights into the sources and characteristics of perceived nonpayment risk.[5]

First, for annuities and LTC insurance, we ask respondents to report the probability of full default conditional on a policy not paying out in a given year for a qualifying event. For example, if you are in need of Long-Term care, but your LTC insurer does not pay out what is promised this year, what is the probability that the insurer defaults on all future commitments.

Second, we ask respondents to report the probabilities of full-default and the perceived payment distribution conditional on an economic downturn. In particular, we prompt respondents to consider a scenario where the stock market has decreased by a randomly determined amount: say 10/20/30%. Respondents then report full-default probabilities and the subjective distribution of perceived payments in the next year through a series of survey questions

[4]Specifically, respondents receive the supplement that corresponds to the third insurance product that they respond about in the core, randomly ordered, survey modules. Collecting supplemental measures about perceived nonpayment risk in this order ensures that our core measures are not influenced by these additional subsequent questions.

[5]See Appendix A for examples on how these additional questions are framed in the survey.

nearly identical to those used to construct our core risk measures, with the conditioning language presented above added.

Finally, we ask respondents for more details about the nature of the perceived nonpayment risk. For example, we ask respondents how much work it will be to receive payments from their insurance policy.

3.3 Other Survey Measures

After responding to survey modules for annuities, life, and LTC insurance, the survey concludes with a final section that asks respondents to reflect on their survey responses. The questions in this final section, as well as the randomization in the design of the survey, primarily serve to provide metrics to check survey response quality. In Section 4.5 we discuss how responses to these questions and other patterns in responses inform the credibility of the constructed risk measures.

We first ask respondents to rank life, annuities, and LTC insurance in order from least to most risky. This question primarily serves as a check for internal consistency, as we would expect respondents to rank products with higher reported nonpayment risk metrics as most risky. In practice, we define a set of responses as not internally incoherent if at least one of the core risk metrics is largest for the policy type rated as "most risky" in this question.

We next ask respondents to report the factors that influenced their responses (e.g., personal experiences, family/friends' experiences, government intervention, policy complexity, the aggregate economy, trust in salespeople, other/free response). Finally, respondents were asked to report how confident they were in their responses, and how much thought they had given to insurance policies and chances of nonpayment prior to the survey.

4 Survey Results

In this section we present measures of perceived risk, relate them to individual charactersitics and behaviors, and consider response credibility. All figures and analyses in this section use population weights provided by the UAS, and thus may be interpreted as indicating measures for a representative sample of American's over age 45.

4.1 Perceived Nonpayment Risk Measures

Our core survey modules provide three key metrics of perceived nonpayment risk: the probability of full default (i.e., the policy becomes worthless and makes no future payments), the

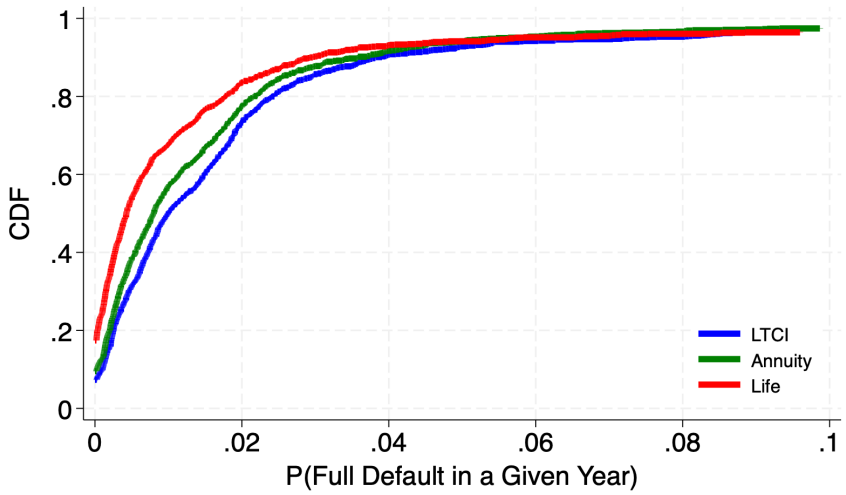

Figure 2: **Annual Probability of Full Default.**

expected value of payments conditional on a qualifying event, and the standard deviation of the perceived distribution of these promised payments.

Figure 2, Panel (a) presents the population CDF of reported probabilities of full-default. Although there is substantial heterogeneity in perception of this risk, the vast majority of respondents assign positive probability to a full default occurring. Only 19% of respondents report zero probability of a full default in life insurance markets, with even less doing so for annuities (11%) and LTCI (10%).There appears to be a clear ordering of perceived risk, with respondents viewing life insurance as least likely to default and LTCI as most likely to default. The median respondent expects life insurance to default in any given year with probability 0.3 percent, annuities with probability 0.5 percent, and LTCI 1 percent. These probabilities are far higher than the full default probabilities observed empirically.

Figure 3 presents the population CDFs for the mean (Panel (a)) and standard deviations (Panel (b)) of the perceived payments in a given year where the policy holder qualified for a claim. For annuities a qualifying event is being alive, for life insurance it is dying, and for LTCI it is needing help with 2 or more activities of daily living. In each figure, the measures are normalized by the full claim amount. Thus, the variables should be interpreted as the fraction of the promised claim that is expected to actually be paid. In Panel (a) we again observe that many respondents perceive substantial payment uncertainty. For life insurance, although almost half of all respondents expect the policy to pay the full promised amount, consumers on average expect to only received 87% of the promised claim amount. Consumers perceive even more risk for annuities and LTCI, reporting mean (median) expected payouts of 82% (90%) and 76% (85%) of the qualified claim amounts, respectively. Panel (b) shows that in addition to expecting to receive a lower payout than promised, consumers perceive

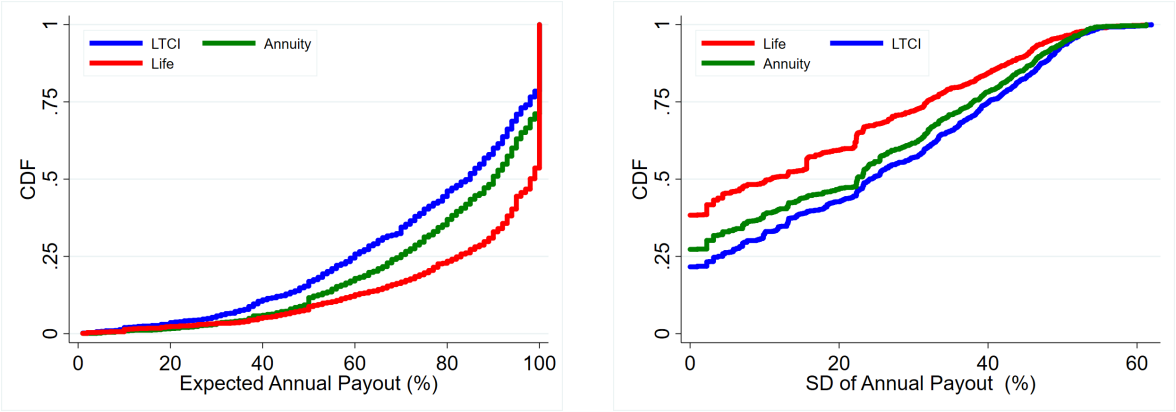

(a) Expected Annual Payout

(b) Standard Deviation of Annual Payout

Figure 3: **Subjective Annual Payout Risk** Conditional on a qualifying event occurring, Panel (a) presents the CDF of the mean of the expected payment distribution and Panel (b) presents the CDF of the standard deviation of the expected payment distribution.

substantial uncertainty in the payout they will receive in each market. On average (at median), perceived payment distributions imply standard deviations of 17% (11%), 19% (19%), 20% (20%) of the promised payment amount for life, annuity, and LTC insurance policies, respectively.

Combined, these figures document our first key result: In all three insurance markets considered, and by all three risk metrics constructed, consumers perceive substantial uncertainty in insurance product payouts. Whether or not there is actually such risk in the market does not change the fact that consumers view life-cycle insurance products as risky. Additionally, there appears to be a clear risk ranking of insurance types, with life insurance being relatively safe, LTC insurance being relative risky, and annuities in the middle.

To provide insight into how consumers value this uncertainty, we next examine the population distributions of certainty equivalent measures in Figure 4. Each point on these CDFs indicate the fraction of the population that would prefer a certain payout at a lower share of the promised payout instead of holding the risk associated with the actual policy. We first observe that certainty equivalent measures are relatively similar across types of insurance policies. At the median, respondents would trade an actual annuity, life insurance, and LTCI policy for a risk-free policy offering 75-80% of the promised payments. This indicates that typical respondents dislike policy risk, and would be willing to forego a substantial share of the promised policy payment to remove this risk.

Table 2 formalizes this point. In this Table, Column (1) presents the average expected payout (as share of promised payout), Column (2) presents the mean certainty equivalent implied by the CDFs in Figure 4, while Column (3) presents the risk premia (defined as

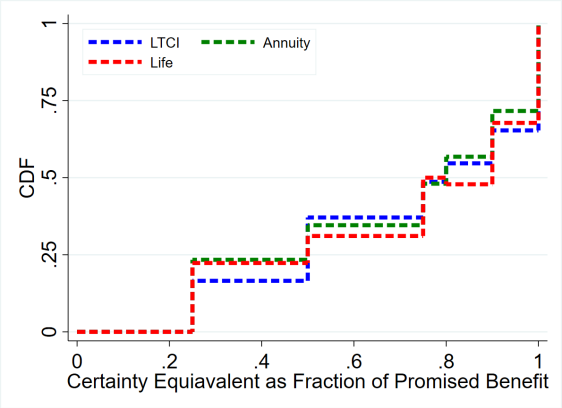

Figure 4: **Certainty Equivalent Distributions.**

Table 2: **Average Expected Payouts, Certainty Equivalent Measures, and the Implied Risk Premia for Life, Annuities, and LTC Insurance**

	Mean Expected Value (1)	Mean Certainty Equivalent (2)	Risk Premia (3)
Life	87.16	81.43	5.72
Annuity	81.51	73.79	7.72
LTCI	76.17	72.90	3.27

the difference between Column (1) and Column (2)) associated with each type of insurance policy. We find that the risk premia associated with each type of insurance policy is quite large. For annuities, in particular, the risk premium is nearly 8%, a level that is comparable to the well-documented, puzzlingly high risk premia associated with publicly traded equities. Overall, Table 2 suggests that not only do respondents view insurance policies as risky, but they require high risk premia to hold them.

4.2 Other Survey Measures

In addition to our three main measures of nonpayment risk, we also analyze several other measures of risk. First, Figure 5 presents the perceived annual payments in different, randomly determined, aggregate states that are proxied by different stock market returns. Relative to the baseline, we find that respondents expect annuity and life insurance payments to decrease in economic downturns, as indicated by the shift in the CDFs to the left conditional on a stock market decrease of 10% and even further left for a decrease of 20%. Interestingly, we observe no corresponding shift for LTC insurance payments.These patterns suggest that the sources of risk associated with different insurance policies varies. Speculatievely, perhaps LTC insurance risk is generally associated with claim denial due to subjective state

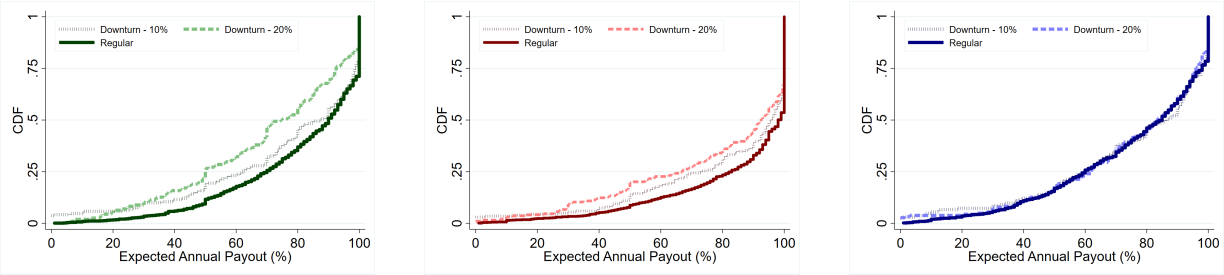

(a) Annuity (b) Life Insurance (c) LTC Insurance

Figure 5: **Expected Payout Distribution Conditional on Different Aggregate Stock Market Conditions.** This figure reproduces Figure 3a in "Regular" times, and also reports survey responses for times when the price of the aggregate stock market is down by 10 percent in a year or down by 20 percent.

verification by insurance companies, while annuity and life insurance risk is more associated with the balance sheet of insurance companies. At the very least, this pattern of responses indicates a certain level of sophistication for survey respondents.

Overall, our supplemental measures of nonpayment risk provide several insights into the situations and manners in which individuals expect insurance companies to meet their contractual obligations.

4.3 Nonpayment Risk and Insurance Ownership

Our survey shows that perceived nonpayment risk in insurance products is large. However, in isolation they do not establish whether perceived risk affects consumer demand. To explore the relationship between demand and perceived risk, in Table 3 we regress indicators of ownership on our risk and other behavioral measures. In particular, for each product $k \in \{\text{Annuity, LI, LTCI}\}$, we run the following regression across all respondents i in the survey:

In particular we estimate a probit model at the respondent-product (i, k) level:

$$\Pr\left(\text{Own}_{i,k} \,\middle|\, E_i\left[\omega^k\right], q_i^k, \sigma_i(\omega^k), X_i\right) = \Phi\left(\alpha_k + \beta_1^k E_i\left[\omega^k\right] + \beta_2^k q_i^k + \beta_3^k \sigma_i(\omega^k) + X_i'\gamma\right)$$

where $\text{Own}_{i,k}$ indicates ownership of product k by respondent i, $E_i\left[\omega^k\right]$ is the expected annual payment, q_i^k is the annual probability of full default, $\sigma_i(\omega^k)$ is the standard deviation of annual payments, and X_i includes the Demographic and Other controls; $\Phi(\cdot)$ is the standard

normal CDF.[6] [7] All the demographic and other controls come directly from the UAS via previous survey responses. We report average marginal effects with standard errors clustered at the respondent level (survey weights applied).

We find that our risk measures are highly predictive of product ownership. For annuities (Column (1)), we find that a higher default probability is associated with lower ownership probability, as is a higher standard deviation of perceived payments. Both of these effects are significant at p-values<.001. For life insurance (Column (2)), a higher expected value of payout upon death is highly predictive of ownership, with the effect again significant at p-values<.001. Additionally, although not significant, higher full default probabilities are again associated with lower ownership probabilities for life insurance. For LTC insurance (Column (3)) we again find that a full default probability predicts lower ownership at p-values<.001, while higher standard deviations of perceived payouts is associated with lower ownership probabilities, although the effect is not significant. In Appendix B we show this is robust to excluding the large set of other individual characteristics measures. These measures have little predictive power in excess of our risk measures.

Table 3: **Regression of Ownership on Risk Metrics**

	Own Annuity	Own Life	Own LTCI
	(1)	**(2)**	**(3)**
Payment Exp. Val ($E[\omega^k]$)	-0.0005	0.0045**	0.0006
	(0.373)	(0.003)	(0.181)
Full Def. Prob (q^k)	-0.0020***	-0.0013	-0.0022***
	(0.000)	(0.142)	(0.000)
Payment SD ($\sigma(\omega^k)$)	-0.0029***	-0.0002	-0.0010
	(0.000)	(0.896)	(0.136)
N	1055	1046	1040
R^2	0.268	0.218	0.179
Demographic Controls	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In addition to significantly predicting ownership, the magnitudes of these effects are

[6]The full list of Demographic controls are: age, wealth, income, education, marital status, race, health, household size, and retirement. These controls are not shown due to space limitations but predict ownership in intuitive ways: for example, wealth, income, and education level are all found to significantly predict ownership of all three products.

[7]The full list of Other Controls are: trust, cognitive score, financial literacy score, numeracy score, experienced fraud, risk aversion, propensity to plan and early stock returns. Conditional on measures of nonpayment risk, all are insignificant except for the Financial Literacy Score on Life insurance ownership.

Table 4: **Counterfactual Predictions of Probit Regressions Under No Risk.**

	P(Own)	P(Own|No Risk)
	(1)	(2)
Annuity	.12	.24
Life	.57	.66
LTCI	.10	.23

economically meaningful, as documented in Table 4. Column (1) presents the predicted ownership rates for our survey sample. Most importantly, Column (2) shows the predicted ownership probability if all respondents were assumed to perceive the insurance products as risk free.

Comparing Columns (1) and (2) shows that perceived risk has significant effects on the overall ownership rates. For annuities, if products were perceived as risk free this exercise suggests that ownership probability would double from .12 to .24, while life insurance ownership rates would increase from .57 to .66. Most strikingly, this exercise suggests extremely large increases in LTC insurance ownership probabilities, from .10 to .23. Of course, this analysis comes with the caveat that these estimated relationships are only locally identified and this model is not the reduced form of any structural model we have developed. In Section 6 we revisit this issue when we examine predictions from our structural model under different assumed risk perceptions.

4.4 What Determines Perceptions of Nonpayment Risk in Insurance Markets?

Having established that beliefs of risks explain behavior, we next turn to examining the sources of these beliefs. To do so, we regress each of our three risk measures for each product on demographic and other individual characteristics. Because our measures are all bounded between zero and one-hundred (with the exception of the standard deviation of payments, for which only the lower bound applies), we estimate tobit regressions with the appropriate lower and upper bounds. The results from these regressions are presented in Table 5.

Several intuitive patterns emerge from these regressions. First, we find that households with higher cognitive (*Cog. Score*) and financial literacy (*Finlit. Score*) scores perceive lower risk. This suggests that some of the negative risk perceptions of insurance products reflects a lack of understanding or belief in the inability to successfully negotiate with insurance companies should an issue arise. In addition, we find that having experienced fraud (*Exp. Fraud*)

in the past is associated with higher perceived risk, especially for full default probabilities. We also find that propensity to plan (*Plan Ahead*) predicts perceived risks, as does stock returns experienced between ages 18-25 (*Cohort Return*). Interestingly, we do not find that self-assessed trust (*Trust*)and risk aversion (*Risk Aversion*) predict beliefs of nonpayment risk, although it is unclear whether this reflects noise in these proxies.

Table 5: **Tobit regressions of risk measures on demographic and behavioral controls.**

	(1) $E[\omega^{\text{Ann}}]$	(2) $E[\omega^{\text{LI}}]$	(3) $E[\omega^{\text{LTC}}]$	(4) $\sigma(\omega^{\text{Ann}})$	(5) $\sigma(\omega^{\text{LI}})$	(6) $\sigma(\omega^{\text{LTC}})$	(7) q^{LI}	(8) q^{Ann}	(9) q^{LTC}
Trust	1.43	2.29	0.28	-1.38	-1.11	-0.40	-0.45	0.16	-0.86
	(0.165)	(0.113)	(0.787)	(0.202)	(0.313)	(0.662)	(0.673)	(0.846)	(0.374)
Cog. Score	0.30	0.36	0.09	-0.15	-0.13	0.09	-0.32*	-0.38**	-0.39**
	(0.059)	(0.096)	(0.565)	(0.346)	(0.493)	(0.501)	(0.039)	(0.006)	(0.007)
Finlit. Score	6.81***	3.58	4.40**	-2.65	-1.46	-1.05	-5.19*	-2.44	-2.12
	(0.000)	(0.145)	(0.010)	(0.147)	(0.437)	(0.469)	(0.017)	(0.180)	(0.236)
Num. Score	0.35	1.71	2.53*	0.38	-1.14	-2.71*	-2.25	-1.07	-2.24
	(0.775)	(0.293)	(0.047)	(0.747)	(0.415)	(0.012)	(0.078)	(0.324)	(0.051)
Exp. Fraud	-0.12	-2.17	-3.24	0.83	2.47	0.69	10.12	7.45*	10.61**
	(0.980)	(0.618)	(0.505)	(0.838)	(0.565)	(0.854)	(0.058)	(0.046)	(0.003)
Risk Subj.	0.80	-0.16	-0.38	-0.69	0.11	0.14	-0.01	0.25	0.54
	(0.158)	(0.795)	(0.498)	(0.163)	(0.839)	(0.760)	(0.983)	(0.612)	(0.282)
Plan Ahead	-2.67*	-1.58	-3.25**	2.65*	2.95**	0.73	0.02	0.92	-0.68
	(0.017)	(0.208)	(0.008)	(0.012)	(0.009)	(0.407)	(0.989)	(0.396)	(0.538)
Cohort Return	-69.23***	-33.01	-47.31*	74.73***	59.35**	35.28	56.52*	44.60*	43.78*
	(0.001)	(0.277)	(0.042)	(0.000)	(0.009)	(0.060)	(0.012)	(0.018)	(0.042)
Cons.	126.78***	99.20**	114.95***	-48.93*	-43.84	-9.50	-3.85	-0.43	20.41
	(0.000)	(0.007)	(0.000)	(0.031)	(0.101)	(0.667)	(0.884)	(0.984)	(0.399)
N	1058	1055	1044	1058	1051	1044	1081	1081	1082

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.5 Response Quality and Credibility

Many of the findings in this study rely on the credibility of the subjective belief measures. The survey was designed to allow for ex-post testing of response quality. For example, the order of survey modules and the policy amounts for individuals that did not already own a given type of insurance, were randomly assigned. Reassuringly, we find that this random order and random framing is insignificant when computing statistics in Tables 3 and 5.

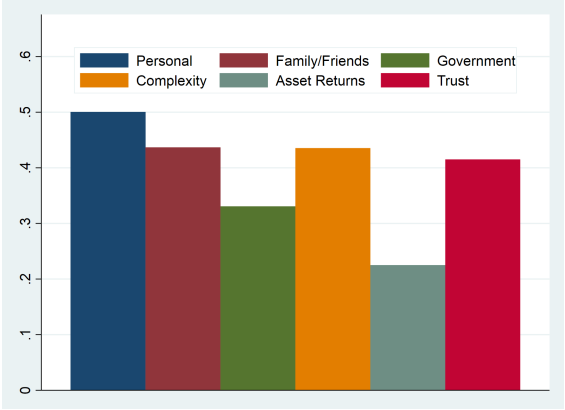

Figure 6: **Factors considered when responding to survey**

A second key built-in check of response quality is our qualitative question, as described in Section 3.3, that asks respondents to rank policies by perceived riskiness. Although such qualitative questions are sometimes difficult to interpret, this question provides a natural check on internal consistency. We define a set of responses as internally coherent if the most risky rated insurance product according to the qualitative question is not rated as the least risky insurance product by all three quantitative risk measures and similarly for the least risky rated product. By this metric, we find that 84% of all respondents are not inconsistent. Furthermore, in Appendix B we analyze this subsample separately, and find that results in Tables 3 and 5 are unchanged.

In addition, several features of survey responses provide evidence of answer quality. For example, we find clear, consistent ordinal rankings of the perceived riskiness of insurance products across individuals. More importantly, our finding that perceived risk measures correlate with product ownership in meaningful ways establishes external validity of our measures, linking subjective expectations questions to observed behavior. Similarly, conditional expectations suggest that respondents perceive aggregate risk as affecting annuity and life insurance payouts, but not LTC insurance. This is consistent with the greater interest rate risk insurance companies face in annuity and life insurance markets.

Finally, self-evaluation measures suggest that respondents were confident in their responses. Figure 6 suggests that respondents considered a number of reasonable factors when responding to the survey. Furthermore, only 17% of respondents report being unconfident in their answers, while only 9% of respondents report having not given any thought to the issues asked about in our survey prior to responding. Overall, these patterns suggest respondents took their charge seriously and answered questions in meaningful ways.

5 Model

This section develops a structural model of insurance product demand that we use to quantify how non-payment risk and other market features shape insurance ownership and consumer welfare for U.S. households. We adopt a representative household perspective: all households are modeled as having identical preferences and beliefs. In particular, we assign households a common probability distribution over insurance payouts that is disciplined by our survey data, such that the representative agent's beliefs capture the cross-sectional distribution of responses. This approach allows us to focus on the aggregate effects of perceived risks and frictions, rather than heterogeneity in preferences or expectations.

Similar to Koijen et al. (2016), we analyze the joint demand for life insurance, annuities, and LTC insurance over the life cycle. Our model departs from their complete-markets framework along several important dimensions. A key methodological contribution is the explicit incorporation of perceived non-payment risk tied to our survey evidence. In addition, the model embeds other real-world product features such as high loads above actuarially fair prices, wedges between purchase and resale prices, and quantity restrictions (e.g., age cutoffs for purchases and short-selling constraints).

Agents in the model differ only in exogenous state variables: age (t), health status (s_t), sex (f), and the quintile of their income-age profile (Y_t). Health and mortality are the main risks, captured through stochastic transitions in s_t. At the start of each period, individuals hold a quantity of risk-free bonds (B_t) and insurance contracts (D_t^k) of type k, with associated annual premia for life insurance and LTCI. They then choose how much to consume, save, and invest in insurance products, subject to the market features described above.

5.1 Model Details

Demographics and Health An agent aged t_0 will live to be at most T years old, with sex denoted $f = 1$ if female and $f = 0$ if male. Each period, the agent realizes its health status s which evolves stochastically between $s = 0$ if normal, $s = 1$ if bad, $s = 2$ if the agent requires LTC, and $s = 3$ if the agent dies this period. Agents are defined to need LTC if they need help with activities of daily living, such as bathing, eating, dressing, getting in/out of bed, and walking.[8] The health status evolves according to a Markov transition matrix $\Gamma_{t,f,Y}$ which is a function of age, sex and income quintile. We allow for transitions to be different between those in the bottom three income quintiles ("poor") and the top two income quintiles

[8]More concretely, a household is defined as being in an LTC state if the household actually receives help in at least two of the five categories. The variable used from the HRS is adl5h.

("rich"). See Appendix C.3 for details on how we estimate this transition matrix.

Health status also determines the exogenous health costs ($HC_{t,s,f}$) that an agent must pay when alive. These costs, which are dependent on age, sex, and current health status, are lognormally distributed and are realized between periods $t-1$ and t and are essentially a stochastic negative shock to liquid wealth. Health costs are distributed according to CDF $F_{HC|t,s,f}$ and increase in expected value with age and health status. See Appendix C for details on how we estimate the health cost shocks and the health transition variables using HRS data.

Preferences Households have time-separable, health-state dependent non-homothetic preferences for consumption C_t and have a warm-glow bequest motive. Flow utility ν_s is

$$\nu_s(C_t) = \theta_{s,t}^{-\sigma} \frac{(C_t + \kappa_s)^{1-\sigma}}{1-\sigma}$$

This follows the specification of utility in Ameriks et al. (2020), where the key innovation was to allow for nonhomothetic utility in the LTC health state i.e. $\kappa_2 \neq 0$. This is analogous to how bequest utility is modeled following Nardi (2004). κ_s controls the degree to which the expenditure when in health state s is a luxury or a necessity, while θ_s a the multiplier on marginal utility of an additional unit of consumption in state s relative to when in good health.

The new feature in this paper relative to Ameriks et al. (2020) is that $\theta_{s,t}$ varies with age. In practice, we allow $\theta_{3,t}$ to vary with age to capture, in a parsimonious way, an age-varying desire to leave a bequest, perhaps driven by the changing economic needs of children as they age. For agents aged above 80, their bequest motive is $\theta_{3,t} = \bar{\theta}_3$, while for those aged between 45 and 80 we derive their bequest motive from the following formula:

$$\theta_{3,t}^{-1} = \theta_{3,45}^{-1} + \left(\bar{\theta}_3^{-1} - \theta_{3,45}^{-1}\right) \left(\frac{\exp(mx)) - 1}{\exp(m) - 1}\right) \tag{1}$$

where $x = \frac{t-45}{80-45}$. $(m, \theta_{3,45}^{-1})$ will be part of the set of calibrated parameters. $\theta_{3,45}^{-1}$ reflects the bequest motive of the youngest agents, while m captures the shape of the life cycle profile of bequest motives.[9]

Overall, we allow for differences in preferences across health states that will be determined

[9]When m tends to 0, the bequest motive becomes linear in age. As m rises, the bequest motive strengthens in middle age, leaving the decline in the motive to be more concentrated among older cohorts in their 70s.

through our calibration strategy described in Section 5.3.

Income Agents have a fixed retirement income age-profile that pays an amount Y_t each year they are alive. This income is the sum of all labor income, social security payments, and payments from defined benefit plans. We model agents as having one of 5 potential paths, estimated from a quintile regression on earnings in the HRS. Details on how we estimate the path for each quintile is given in Section C.2.

Because the focus of this model is asset and insurance product demands and because retirement income (especially social security and defined benefit plan payments) is generally not as risky as prime-age labor income, in our baseline model we do not allow for income uncertainty. This allows us to reduce the state space of the model and facilitates computation.

Assets Agents can invest in four different assets. First, an agent can invest in a bond (B_t) that offers a return of r. Agents are unable to borrow in bond markets, so $B_t \geq 0$.

In addition, agents can invest in three insurance products, indexed by k: a life annuity (Ann), a whole life insurance policy (LI), or long-term care insurance ($LTCI$). Agents can purchase any of the insurance products up until a maximum age $t_{max,k}$. This captures a real world feature that after a certain age insurers, perhaps concerned with adverse selection, will not sell to individuals.

The insurance products differ in the health states s in which they promise to pay out. Let $\bar{D}^k$ be a 4×1 vector of dividends of asset k, where row s is defined as the payout in state s from owning one unit of insurance asset k: $\bar{D}^{Ann} = [1, 1, 1, 0]'$; $\bar{D}^{LI} = [0, 0, 0, 1]'$; $\bar{D}^{LTCI} = [0, 0, 1, 0]'$.

Annuities are purchased by making a lump-sum payment of bonds. We denote the net transactions in annuities at age t as W_t^k. Thus, the law of motion on holdings D_t^{Ann} can be written as:

$$D_{t+1}^{Ann} = D_t^{Ann} + W_t^{Ann}. \tag{2}$$

Individuals owning life insurance and LTC insurance pay annual premia. For individuals who want to reduce their insurance holdings, annuities can be sold for some cash value, while individuals owning life insurance and LTC insurance can simply stop paying the premia.

Purchase prices and annual premia are set as a markup λ_+^k over the actuarially fair price. Annuities are sold at a markdown below the acturially fair price λ_-^{ANN}. The acturially fair price, denoted as $p_{t,s,f}^k$, is the price at which a risk neutral insurance company earns zero expected profits by selling a unit of insurance product k to an individual aged t_0, in health

state s_0, of sex f.[10] We allow for the rate r_I used by insurers to discount future cash flows to differ from the rate of return the household receives on its risk free bonds.[11] However, we assume that income is not observable to the insurer, meaning that the health transitions they use for pricing, $(\Gamma_{j,f})$ are unconditional on income.

For an annuity, the actuarially fair lump-sum unit cost is computed as:

$$p^{Ann}_{t,s,f} = \sum_{t'=t+1}^{T} \frac{1}{(1+r_I)^{t'-t}} \times v_s' \left(\prod_{j=t}^{t'} \Gamma_{j,f} \right) \bar{D}^{Ann} \tag{3}$$

where v_s is a 4×1 vector with a one in row s and zeros elsewhere.

For life insurance and LTC insurance, the price paid is in the form of a constant annual premium. This annual premium is computed such that insurers expect to earn zero profits under the assumption that the product will be held until death (no lapse). This payment can be represented as the consumer shorting an annuity to the insurer with yearly payment:

$$p^{k}_{t,s,f} = \frac{\left[\sum_{t'=t+1}^{T} \frac{1}{(1+r_I)^{t'-t}} \times v_s' \left(\prod_{j=t}^{t'} \Gamma_{j,f} \right) \bar{D}^{k} \right]}{p^{ANN}_{t,s,f}} , \quad \text{for } k \in \{LI, LTCI\} \tag{4}$$

For life insurance and LTCI the annual premium is locked in at the time of purchase, and is thus a state variable. To keep the state space computationally tractable, we only allow individuals to hold at most one life insurance and one LTCI policy at a time. If they want to change their LI or LTCI holdings, they must let their old contract lapse and can purchase a new contract at current prices.

Additionally, we do not allow for short-sales of insurance products, so $D^k_{t+1} \geq 0$.

Default Probabilities We allow for two types of default for each insurance product. First the insurance company might fully default on the value of the contract. Practically, this means that with probability q^k the holdings of the contract (D^k_{t+1}) are set to zero in the next period without any sale proceeds:

$$D^k_{t+1} = \begin{cases} D^k_t + W^k_t & with\ prob\ \ 1-q^k \\ 0 & with\ prob\ \ q^k \end{cases} \tag{5}$$

Second, the insurance company may partially default on the promised payout in a given

[10] This is also under the assumption that the insurer holds the contract until they die.

[11] We can interpret this as the insurer's cost of capital.

period. Let $\hat{D}^k_{t+1}$ be a 4×1 vector conveying the realised payout per unit of product k owned. Knowing that $\bar{D}^k$ is the promised payout per unit of product k, we have that:

$$\hat{D}^k_{t+1} = \omega^k_{t+1} \bar{D}^k \quad \text{where} \quad \omega^k_{t+1} \sim G^k(.)$$

where ω^k_{t+1} reflects the amount paid out in product k as a fraction of what is promised. We calibrate the distribution $G^k(.)$ by discretising it into four equally likely outcomes so that the expected value and the standard deviation matches the expected payment distribution's mean and variance from the survey.

Government Insurance Consumption is insured by the government through a minimum consumption level $\bar{C}_s$ that varies depending on the health state s. To receive government insurance, an agent must give all of its wealth to the government and in return consumes $\bar{C}_s$. Following the rules of Medicaid, agents then enter next period with zero bond and insurance holdings, i.e., $B_{t+1} = D^k_{t+1} = 0$ for all k. If an agent decides to use government insurance when in the good ($s = 0$) or bad ($s = 1$) health state, the agent receives a floor $\bar{C}_s = \bar{C}$.

If the agent is in the LTC ($s = 2$) state and decides to take government insurance, the government provides a consumption level $\bar{C}_2 = \psi_G$. ψ_G parameterizes the value agents place on government care that is loosely based on the institutions of Medicaid.

Intertemporal Budget Constraints We now collect the results from above and specify the intertemporal budget constraint that tracks the evolution of bonds.

The law of motion for risk-free bonds can be written as:

$$\begin{aligned} (6) \qquad \frac{1}{(1+r)} B_{t+1} =& B_t + Y_t + v'_{s_t} \sum_k D^k_t \hat{D}^k_t - C_t - HC_t - p^{LTCI}_{i,t} D^{LTCI}_t - p^{LI}_{i,t} D^{LI}_t \\ & - W^{ann}_t p^{ann}_{t,s,f} \left(1 - \lambda^{ann}_- \mathbb{I}_{W^{ann}_t < 0} + \lambda^{ann}_+ \mathbb{I}_{W^{ann}_t > 0}\right) \end{aligned}$$

The individual starts the period with some bond and insurance holdings, earns income, receives payments from insurance products if a qualifying event occurs and the insurance product pays, consumes, pays health costs, pays premia on life insurance and LTCI, and receives or pays income from the sale or purchase of annuities. $p^{LTCI}_{i,t}$ and $p^{LI}_{i,t}$ are set at the time of purchase of the contracts and are constant unless choosing to change the quantity of insurance owned as described above.

Full Decision Problem We now specify the full decision problem solved by each agent. Let $X_t = \{s, Y_t, f\}$ and $X_{t+1} = \{s', Y_{t+1}, f\}$ denote exogenous state variables. The consumer decision problem is:

$$V_t(B_t, D_t^{ANN}, D_t^{LI}, D_t^{LTC}, X_t, p_i^{LTCI}, p_i^{LI}) = \max\{V_t^{GOV}(B_t, D_t^{ANN}, D_t^{LI}, D_t^{LTC}, X_t, p_i^{LTCI}, p_i^{LI}), \\ V_t^{PRIV}(B_t, D_t^{ANN}, D_t^{LI}, D_t^{LTC}, X_t, p_i^{LTCI}, p_i^{LI})\}$$

where if they take up government care:

$$V_t^{GOV}(B_t, D_t^{ANN}, D_t^{LI}, D_t^{LTC}, X_t, p_i^{LTCI}, p_i^{LI}) = \nu_s(\bar{C}_s) + \beta\mathbb{E}\left[V_{t+1}(B_{t+1}, 0, 0, 0, X_{t+1}, 0, 0)\right] \\ \text{s.t.} \quad B_{t+1} = Y_{t+1} - HC_{t+1,s',f}$$

i.e. they consume minimum consumption $\bar{C}_s$ and give up all their insurance holdings going into period $t+1$, and otherwise

$$\begin{aligned} V_t^{PRIV}(B_t, D_t^{ANN}, D_t^{LI}, D_t^{LTC}, X_t, p_i^{LTCI}, p_i^{LI}) = & \max_{W_t^{ANN}, D_{t+1}^{LI}, D_{t+1}^{LTC}, C_t} \nu_s(C_t) \\ & + \beta\mathbb{E}\left[V_{t+1}(B_{t+1}, D_{t+1}^{ANN}, D_{t+1}^{LI}, D_{t+1}^{LTC}, X_{t+1}, p_{i,t+1}^{LTCI}, p_{i,t+1}^{LI})\right] \\ \text{s.t.} \quad \frac{1}{(1+r)} B_{t+1} = & B_t + Y_t + v'_{s_t} \sum_k D_t^k \hat{D}_t^k - C_t - HC_t - p_{i,t}^{LTCI} D_t^{LTCI} - p_{i,t}^{LI} D_t^{LI} \\ & - W_t^{ann} p_{t,s,f}^{ann} \left(1 - \lambda_-^{ann} \mathbb{I}_{W_t^{ann}<0} + \lambda_+^{ann} \mathbb{I}_{W_t^{ann}>0}\right) \end{aligned}$$

where annual and full default probabilities evolve according to

$$\hat{D}_{t+1}^k = \omega_{t+1}^k \bar{D}^k \quad \text{where} \quad \omega_{t+1}^k \sim G^k(.)$$

$$D_{t+1}^k = \begin{cases} D_t^k + W_t^k & \textit{with prob} \;\; 1 - q^k \\ 0 & \textit{with prob} \;\; q^k \end{cases}$$

and rebalancing is subject to the following conditions:

$$\begin{aligned} & D_{t+1}^k \geq 0 \\ & D_{t+1}^k \leq D_t^k \quad \text{if} \quad t > t_{max} \\ & p_{i,t+1}^k = \begin{cases} p_i^k & \text{if} \quad D_{t+1}^k = D_t^k \\ p_{t,s,f}^k & \text{if} \quad D_{t+1}^k \neq D_t^k \end{cases} \quad \text{for} \quad k \in \{LI, LTC\} \end{aligned}$$

with stochastic processes

$$s' \sim \Gamma_{t,f,Y|s}$$
$$HC_{t+1,s,f} \sim F_{HC|t+1,s,f}$$

and prices

$$p^{Ann}_{t,s,f} = \sum_{t'=t+1}^{T} \frac{1}{(1+r_I)^{t'-t}} \times v'_s \left(\prod_{j=t}^{t'} \Gamma_{j,f} \right) \bar{D}^{Ann}$$
$$p^{LI}_{t,s,f} = \left[\sum_{t'=t+1}^{T} \frac{1}{(1+r_I)^{t'-t}} \times v'_s \left(\prod_{j=t}^{t'} \Gamma_{j,f} \right) \bar{D}^{LI} \right] / p^{Ann}_{t,s,f}$$
$$p^{LTCI}_{t,s,f} = \left[\sum_{t'=t+1}^{T} \frac{1}{(1+r_I)^{t'-t}} \times v'_s \left(\prod_{j=t}^{t'} \Gamma_{j,f} \right) \bar{D}^{LTCI} \right] / p^{Ann}_{t,s,f}$$

5.2 Risks, Costs, and Incomplete Markets

In this model, individuals, depending on their preferences, may have a precautionary saving motive arising from three distinct sources: longevity risk, long-term care (LTC) health risk, and bequest motives. The available financial products allow them to offset these risks: annuities help insure against longevity risk, LTC insurance (LTCI) against LTC health risk, and life insurance against bequest risk. In a complete markets environment - such as the one studied in Koijen et al. (2016) - these insurance products would, in principle, eliminate the need for precautionary savings.

In contrast, our model operates in an incomplete markets setting, incorporating several non-insurable risks, costs, and frictions that can discourage agents from purchasing insurance and instead lead them to self-insure with personal savings. Because capturing these features and analyzing their impact on insurance demand is a core contribution of this study, it is useful to outline our modeling choices and discuss how they shape insurance demand before presenting the main results.

First, on all insurance products, we impose a maximum age at which agents are able to purchase new or add to previously purchased policies. Although insurance companies could always in principle offer appropriately priced insurance products, insurers are generally unwilling to sell new policies to older individuals, possibly due to difficulty accounting for adverse selection. For example Hendren (2013) notes that LTCI applications for individuals over age 80 are automatically rejected, while MetLife notes that 98% of annuity purchases

occur before age 85, and most life insurance companies stop issuing new policies to individuals between ages 80 and 90.

A second source of imperfection in the market we model is price wedges that require that an agent pay a more than actuarially fair price to purchase the product and, in the case of annuities, receive a less than actuarially fair price when selling the product. Theoretically, these price wedges have two effects. First, agents are less willing to purchase the product since it is both more expensive and worth less in the future should they decide to sell it. Second, these price wedges produce regions in the space of asset holdings where individuals are unwilling to buy or sell the product, and potentially generate the infrequent adjustment of insurance holdings that we observe empirically. Importantly, the spread between buying and selling, $\lambda_{+}^{Ann} - \lambda_{-}^{Ann}$, generates different marginal payoffs between states that can not be insured with the modeled assets.

A third significant feature of the insurance we model is the nonpayment risk, captured by the probabilities that the insurance company will default on the total value (q^k) and the uncertainty over the amount paid each year as a fraction of what is promised (ω_t^k). Allowing for these nonpayment probabilities provides two deterrents to insurance purchase. First, because the product price defined in equation 3 does not factor in default, the price of the insurance policy is worse than actuarially fair. More importantly, however, allowing for nonpayment risk introduces uninsurable model states that feature potentially large losses of wealth for agents that invest in these policies, and rational agents may choose to forgo product purchase to avoid holding this risk.

A fourth modeled deterrent to insurance product purchase is uninsurable health risk. Out-of pocket health costs are known to generate a significant precautionary savings motive (see De Nardi et al. (2010), among others), and consumers may be reluctant to commit wealth to insurance products that don't cover these expenses.

In short, the model developed in this paper provides a rich environment to explore how different sources of market incompletion, costs, and risks affect demand for the three considered insurance products, and most importantly, better understand whether the low levels of insurance are rational when one accounts for a range of potential deterrents to insurance product purchase.

5.3 Model Calibration

We estimate 5 income-age profiles by running a quintile regression on HRS data, separately for men and women. Income includes labor income, defined contribution pension plans, and social security payments. We also estimate the sex and age dependent health-state

transition matrix and the out-of-pocket medical health expenditure shocks using HRS data. See Appendix C for details. We also set consumption when using government care and not in need of LTC to $\bar{C}_0 = \bar{C}_1 = \$5K$, in line with the minimum level of consumption chosen in De Nardi et al. (2016).

We also set independently the parameters that govern insurance product characteristics. We firstly truncate the responses for each product at the bottom and top 2% to avoid sensitivity to outliers. We then use the mean of our original survey responses to calibrate the full default probabilities q^k. For the distribution of annual payments as a fraction of what is promised, $G^k(.)$, we discretize into four equally likely outcomes, whose values are chosen to match as closely as possible the distribution of survey responses for each product.[12]

We calibrate prices following Brown and Finkelstein (2011): observed quotes are modelled as actuarially fair values (computed conditional on age, health and health transitions using Γ) scaled by product-specific gross margins λ^k. We calibrate λ^k to match empirical price-to-actuarial ratios ("loads") for life insurance, annuities, and LTCI reported in the literature (e.g., Hong and Rìos-Rull (2012); Brown and Finkelstein (2008)), which are computed under the same pricing structure we use. We interpret λ^k as reduced-form gross margins that summarize supply-side forces - administrative and distribution costs, capital charges, adverse selection, expected lapsation revenues, etc. - allowing us to replicate data-consistent prices without modeling insurer costs explicitly; in counterfactuals we vary λ to study implications for ownership, welfare, and producer surplus.

The nonhomothetic health-state dependent utility functional form is taken from Ameriks et al. (2020), from which we also draw a subset of the preference parameters, $(\theta_2, \bar{\theta}_3, \kappa_2, \kappa_3)$ as well as the value of government-provided LTC, (ψ_G). Their values are also shown in Table 7.

We estimate the remaining parameters of the model, $\Theta := (\beta, \theta_{3,45}, \sigma, m, r, r_I)$ using the simulated method of moments to match 3 sets of moments. We target the 25th, 50th, and 75th percentile of the wealth distribution by 5-year age bins for ages 45–84. We also target the average ownership rate of annuities, life insurance, and long-term-care insurance for the full sample of households. Finally, we target the 9 strategic survey question (SSQ) mean responses from Ameriks et al. (2020). These SSQs place individuals in hypothetical scenarios and ask them to make choices that are designed to be highly revealing of preferences. For each combination of parameters, we solve the model, derive the consequent moments, and

[12]For each insurance product, we draw from the distribution of survey responses of each product the following percentiles: 12.5, 37.5, 62.5, 87.5. It turns out that, by taking these points, and assuming equal probabilities of each outcome, we match almost exactly the mean and standard deviation of the survey distribution for each product. The set of possible outcomes are $\omega_t^{ANN} \in \{.50, .78, .94, 1\}$, $\omega_t^{LI} \in \{.58, .92, 1, 1\}$, and $\omega_t^{LTCI} \in \{.48, .71, .88, 1\}$.

Table 6: **Baseline Calibration - Insurance products**

	Annuities	Life	LTCI
Full Default (q^k)	.014	.013	.018
Annual Payout Fraction ($E[\omega]$)	.815	.871	.761
Annual Payout Uncertainty ($\sigma(\omega)$)	.194	.174	.195
Price Wedge, buying (λ_+^k)	.2	.25	.32
Price Wedge, selling (λ_-^k)	.15	–	–
Max Purchase age ($t_{max,k}$)	80	80	80

then iterate across parameter combinations until we optimize the moment condition. See Appendix D for details on our estimation strategy.

Model Solution In their model with complete markets, Koijen et al. (2016) provide a set of necessary conditions for optimal insurance portfolio choice as well as an analytical characterization of the welfare loss due to sub-optimal insurance product choice. Unfortunately, analytic solutions do not exist for the incomplete markets framework considered in this paper. We thus solve the model quantitatively, namely by backwards induction, utilizing a 6-dimensional grid search over the state variables $(B_t, D_t^{ANN}, D_t^{LI}, D_t^{LTCI}, p_i^{LTCI}, p_i^{LI})$. To simulate the model, we set as an initial condition the distribution of state variable values drawn from UAS survey respondents. In particular, for each respondent, we take their age, sex, income, liquid + housing wealth, and measured insurance product ownership.[13] Then we simulate the model forward, allowing individuals in the model to re-optimize their portfolios and savings over the life cycle. See Appendix D.1 for more details on how the model is simulated.

Parameter Estimation Results The estimated parameters are reported in column 2 of Table 7. A value of $\sigma = 2.5$ implies a positive but moderate degree of household risk aversion. The bequest motive for younger households, $\theta_{3,45}^{-1}$, is calibrated so that marginal utility at death at age 45 is scaled by 29.5 relative to marginal utility in good health. This generates a strong incentive for younger households to hold life insurance to protect against the rare event of early death and leaving a bequest. With $m = 2$, the shape of the bequest function ensures that this motive remains strong through middle age and begins to taper only after about age 70.

The level of β also plays a key role in shaping insurance demand. A relatively high

[13]We presume that the LTC and life insurance annual premiums they locked in for their observed UAS survey holdings are the ones they would have if they purchased at age 45 when healthy.

Table 7: **Baseline Calibration - Preferences and Returns**

(Ameriks et al. (2020))	*Estimated*
LTC motive - $\theta_2 = 0.67$	**Time preference** - $\beta = 0.9575$
Bequest motive - $\bar{\theta}_3 = 1.09$	**Youth bequest motive** - $\theta_{3,45}^{-1} = 29.5$
LTC luxury - $\kappa_2 = -37.44$	**Risk aversion** - $\sigma = 1.5$
Bequest luxury - $\kappa_3 = 7.83$	**Bequest shape** - $m = 2.5$
Gov LTC care - $\psi_G = 77.43$	**Bond return** - $r = 0.05$
Min. consumption - $\bar{C}_0 = \$5\text{k}$	**Insurance discount rate** - $r_I = 0.025$

β reduces incentives to insure against earlier-life health shocks, while placing more weight on longevity risk and consumption later in life. Combined with an interest rate r such that $\beta(1+r) > 1$, this encourages significant wealth accumulation earlier in the life cycle. Appendix D.4 offers a more detailed discussion surrounding the sensitivity of model outcomes to these parameters.

6 Model Results

We now present results from our estimated model. Figure 7 shows that the model replicates both the level and the hump-shaped life-cycle profile of wealth observed in the HRS across percentiles. Figure 8 further demonstrates that we closely match participation rates of all insurance products across the life cycle. Another important success of the model is that it also captures household portfolio composition: both the number of insurance products owned (Table 8) and the correlations in participation across products (Table 9). This is especially noteworthy since only the aggregate participation rates of insurance products were targeted in the estimation.

Table 8: **Number of Insurance Products Owned**

# Products	**0**	**1**	**2**	**3**
Model	41%	46%	12%	1%
Data	35%	55%	9%	1%

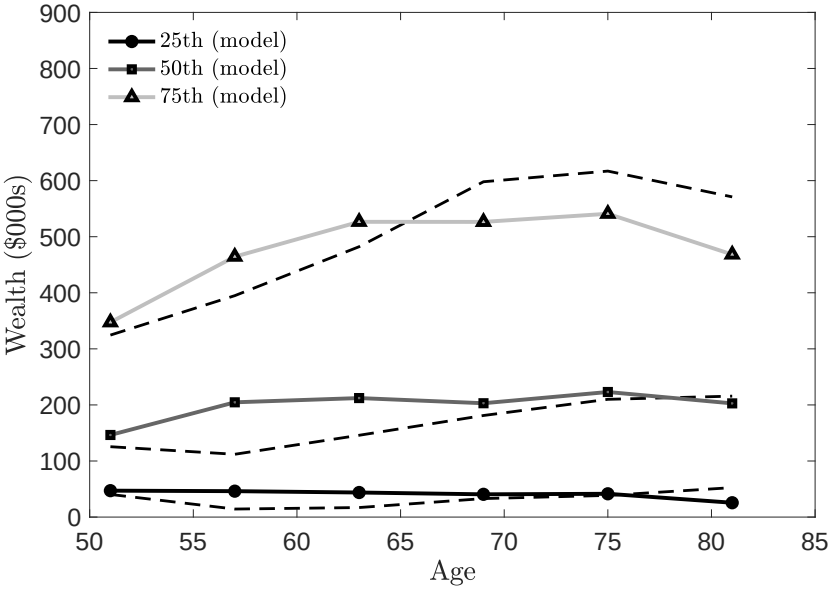

Figure 7: **Wealth Moments 25p, 50p and 75p Model and Data**

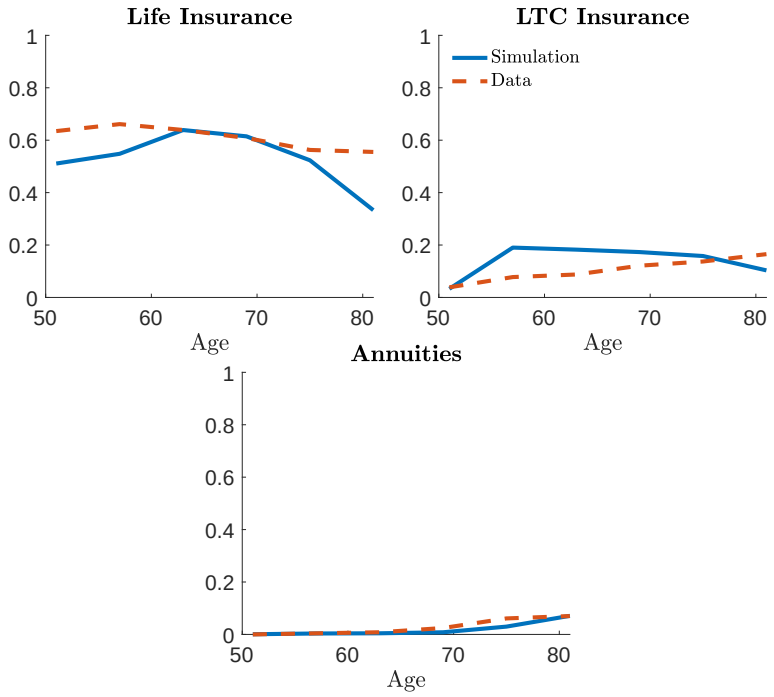

Figure 8: **Insurance Ownership: Model and Data**

Table 9: **Correlation of Insurance Product Ownership**

Participation Correlations	**(LI,LTC)**	**(LI,Ann)**	**(LTC,Ann)**
Model	0.15	-0.03	0.10
Data	0.08	-0.02	0.06

To understand what drives these portfolio choices, we begin with life insurance - the most widely purchased product. The main driver of life insurance demand is the strong estimated bequest motive for young and middle-aged households. This explains why participation remains high through the 70s, only declining once the bequest motive weakens at older ages. Importantly, participation is not limited to the wealthy. Table 10 shows that even the majority of households in the bottom 25% of the wealth distribution own life insurance. They are able to do so because life insurance requires only annual premium payments, rather than up-front payments that confront their zero borrowing constraint. Instead, households allocate a share of future labor income to the contract - on average 18% as reported in Table 11. In contrast, wealthy households are least likely to participate (29%), as they prefer to self-insure, finding available products unattractive due to market incompleteness and deciding instead to self-insure.

Table 8 confirms that nearly 90% of households either hold no insurance at all or just a single life insurance product. When they do add another product, it is most often LTC insurance, consistent with the positive correlation between Life and LTC participation in Table 9. Households are motivated to take up LTC insurance for two reasons. First, there is an incidental bequest motive: LTC raises wealth in states where leaving a bequest becomes more likely. In practice, however, this role is already largely filled by life insurance. Second, LTC is a necessity state, making actuarially fair LTCI potentially highly valuable. Still, actual holdings remain modest because transaction costs and non-payment risk make the products very expensive. Participation in annuities is essentially zero. While annuities provide longevity protection, their high cost again leads households to self-insure through wealth accumulation.[14]

Overall, Table 10 shows that non-participation in any product is widespread across the wealth distribution. This leaves many households under-insured against longevity, health,

[14]Do note, however, that annuity participation rises towards 10% as households reach the purchasing age limit of 80. A key reason for not holding annuities earlier in life is that the discount rate, r_I, is below the return on bonds, r, which matters a lot for long-maturity insurance because of compounding interest. This becomes less of a concern later in age when conditional life expectancy drops.

Table 10: **Participation Rates: By Wealth Percentile Group**

	0-25th	**25-50th**	**50-75th**	**75-100th**
A. *Model*				
LI	60%	73%	54%	29%
LTC	10%	27%	15%	12%
Ann	0%	2%	3%	5%
Non-Part	40%	23%	39%	62%

Table 11: **Promised Premia: Fraction Lifetime Labor Income**

	0-25th	**25-50th**	**50-75th**	**75-100th**
LI	18%	29%	17%	8%
LTC	3%	6%	3%	3%

and bequest risks, generating a strong precautionary savings motive. Figure 7 reflects this: later in life, households across the wealth distribution do not draw down wealth but instead retain it as self-insurance. Absent these risks, one would expect households to run down savings in order to smooth consumption against the steep decline in labor income (see Appendix C.2). Indeed, as we show later in this section, once product-market frictions are removed, households take on more insurance and draw down wealth much more aggressively in old age.

In the remainder of this section, we demonstrate that transaction wedges and non-payment risk are central determinants of insurance participation; that these frictions impose significant welfare costs, even if non-payment beliefs are mistaken; that profit margins vary across products and scenarios; and that social security and publicly provided LTC have substantial effects on both ownership and welfare.

6.1 Ownership and Product Market Frictions

In this section we examine the role of product market frictions in shaping households' portfolio choices between insurance products and wealth. Columns (3)-(6) of Table 12 report results under four scenarios: the "Baseline," our benchmark specification; "No Price Wedges", which eliminates markups over actuarially fair prices; "No Non-Payment Risk", which sets default

probabilities to zero; and "No Price Wedges and Non-Payment Risk", which combines both adjustments. Hereafter, we interchange Scenario 3 for "No Price Wedges", Scenario 4 for "No Non-Payment Risk", and Scenario 5 for "No Price Wedges and Non-Payment Risk".

Table 12: **Insurance Demand and Welfare Costs of Sub-optimal Portfolio Holdings: Rational Expectations**. Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	Data	**No Insurance**	**Baseline**	**Scenario 3**	**Scenario 4**	**Scenario 5**
A. *Insurance Ownership*						
LI	61%	0.0%	55%	71%	73%	84%
LTC	12%	0.0%	13%	29%	43%	55%
Ann	4%	0.0%	2%	3%	8%	10%
B. Δ *Personal Savings*						
25th Perc.	–	–	−37%	−45%	−52%	−60%
50th Perc.	–	–	−15%	−21%	−32%	−40%
75th Perc.	–	–	−5%	−12%	−21%	−29%

Panel A of Table 12 reports aggregate ownership rates for annuities, life insurance, and LTC insurance. Column (3) confirms that the Baseline matches aggregate participation rates in the data. Comparing columns (4)-(6) highlights the impact of frictions. For life insurance, the 25% load above actuarially fair pricing significantly suppresses demand: removing the markup raises ownership to 71%. Non-payment risk has an even larger effect, with ownership rising to 73% when this risk is eliminated. Together, zero markups and zero non-payment risk increase life insurance participation to 84%. The pattern is even starker for LTC insurance: removing non-payment risk raises ownership from 13% to 43%; removing markups alone raises it to 29%; and removing both pushes it to 55%. For annuities, by contrast, the effects are more muted - ownership rises only modestly to 8% without non-payment risk, 3% without markups, and 10% when both frictions are removed.

Even under risk-free, actuarially fair pricing, not all households purchase insurance: 9% of households opt out of all three products. This non-participation is heavily concentrated among the poorest households - 41% of those in the bottom wealth decile remain uninsured. This reflects the role of implicit government care, which provides a fallback against longevity and LTC risks and thus weakens the value of private insurance for the poorest households.

Panel B of Table 12 shows how wealth holdings adjust across scenarios, relative to a setting with no insurance. We report percentage changes at the 25th, 50th, and 75th percentiles. In the Baseline (column 3), wealth falls substantially for poorer households - by 37% at the 25th percentile - because they have the strongest precautionary motive to save.

Without insurance, these households face the risk of reaching adverse health states with insufficient resources (e.g., dying young or entering LTC). Access to insurance allows them to substitute toward state-contingent assets. Wealthier households, who are more likely to have adequate savings across states, show smaller reductions (15% at the median and 5% at the 75th percentile).

Comparing columns (4)-(6) shows that the decline in wealth is more pronounced as frictions are removed. The largest drop occurs when both price wedges and non-payment risk are eliminated. Two mechanisms drive this result. First, insurance products become more attractive investments when cheaper and safer, leading households to shift funds out of wealth and into insurance. Second, improved insurance weakens the precautionary motive to hold wealth, since state-contingent protection is now better provided by markets.

Figure 9 further illustrates how frictions shape the wealth distribution across percentiles and ages. The figure shows that wealth holdings decline relative to the Baseline across the distribution, with the decline especially pronounced at older ages. Intuitively, as households age, the likelihood of entering LTC or leaving a bequest rises, so precautionary saving motives are strongest later in life - and fall the most once insurance becomes more complete.

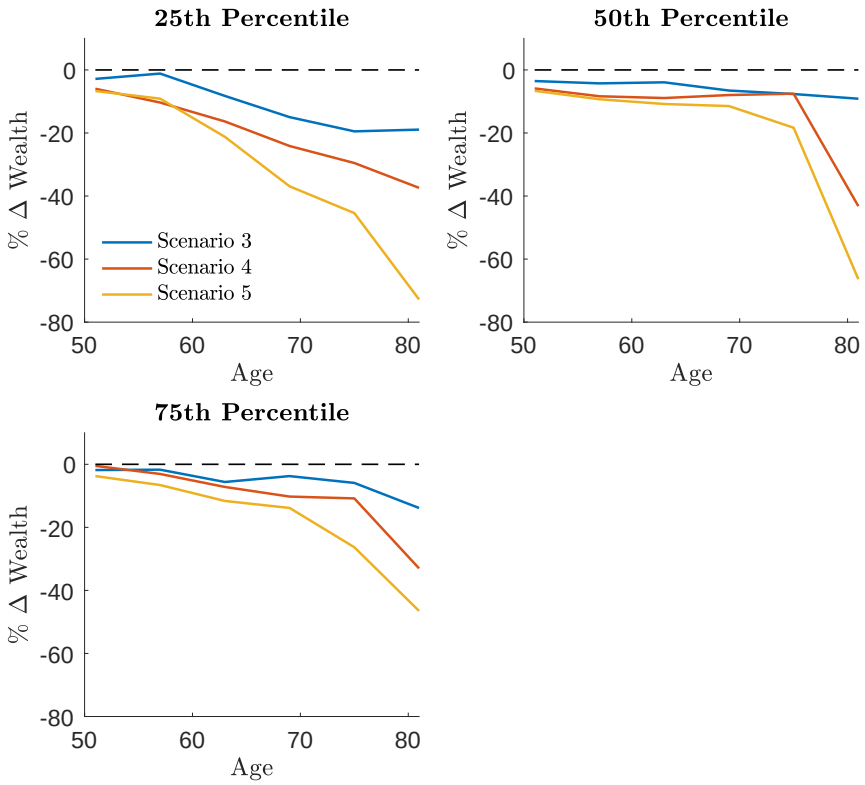

Figure 9: **Wealth Holdings: Comparison Across Scenarios**

6.2 Welfare Results: Rational vs. Non-Rational

As of now, we have yet to take a stand on whether or not beliefs are rational because households' portfolio choices depend only on their beliefs about insurance payouts, not on whether these beliefs are correct. For welfare, however, the distinction matters: if beliefs are mistaken, realized outcomes diverge from what households expected. To assess this, we compare two settings while holdings beliefs fixed: (i) Rational: beliefs are correct, so non-payment risk is realized as perceived, and (ii) Non-Rational: households still act on the same beliefs, but in fact insurers always pay i.e. there is no true non-payment risk.

We evaluate welfare across the same scenarios as before - Baseline, No Price Wedges, No Nonpayment Risk, and No Price Wedges/Nonpayment Risk - using a consumption-equivalent measure aggregated across households. This measure is the % increase in consumption required across all states of the world when there are no insurance products in order to be indifferent to living in the world described by each scenario.

Table 13: **Welfare Costs of Sub-optimal Portfolio Holdings Across Scenarios: Rational vs. Non-Rational Expectations** Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	No Insurance	**Baseline**	**Scenario 3**	**Scenario 4**	**Scenario 5**
A. *Rational*					
Consumption Equivalent	0%	0.6%	1,8%	2.3%	4.1%
B. *Non-Rational*					
Consumption Equivalent	-0%	1.7%	3.4%	2.3%	4.1%

Panel A of Table 12 shows that, under the Rational setting, baseline welfare gains are modest at 0.6% of lifetime consumption. This reflects the fact that nearly half of households optimally choose not to own any insurance (Table 8) relying instead on self-insurance and public transfers, and those who own insurance are left with products that sometimes do not pay when promised. Removing frictions, however, produces substantial welfare improvements: eliminating price wedges nearly triples welfare gains to 1.9%; eliminating non-payment risk raises them to 2.3%; and removing both simultaneously yields a 4.0% gain.[15]

[15]Non-payment risk takes two forms: annual payment risk and full default risk. It turns out that 75-80% of welfare costs of nonpayment risk are caused by annual payment risk. Intuitively, for Life and LTC insurance, you pay for it gradually using annual premia, and if they default, the household simply ceases to pay premia in the future. As a result, it is not so costly, except that they may lose the preferential premium price that they locked in on the products.

This pattern underscores the central result of the model: large welfare gains arise when insurance products are available at actuarially fair prices and are free of default risk because households can effectively hedge away substantial health risks. Instead, a market with insurance products containing large mark-ups and risk of non-payments prevents individuals from insuring against these risks using these products, at a substantial welfare cost.

Panel B shows the results under Non-Rational beliefs. Policy rules remain unchanged, but realised welfare differs when there are nonpayment beliefs because insurers in fact always pay. In the Baseline, welfare rises to 1.7% (compared to 0.6% under Rational) simply because insurance payouts are always made. More interestingly, comparing this 1.7% with the 2.3% in the "No Nonpayment Risk" column highlights the role of beliefs: when households know there is no default risk, they decide to purchase significantly more insurance, and welfare rises accordingly because they are now using available products more effectively. The characteristics of the products remain unchanged - it is driven by households now understanding correctly that they are in fact better insurance products.

Overall, across both perspectives, the message remains the same: the removal of product market frictions - eliminating price wedges and either removing actual or perceived nonpayment risk - produces large welfare gains by allowing households to insure major life-cycle health risks.

In the following sections, we present welfare costs under the Rational setting. As is the case here, the same conclusions are drawn whether we consider the Rational or Non-Rational beliefs, and so we fix on one of the settings. All welfare results in the remainder of the paper are presented under the Non-Rational setting in Appendix E.

6.3 Welfare Results: By Wealth and Age

We documented the welfare effects for households on average. But how do they look if we look at sub-cohorts split by wealth or age? Panel A of Table 14 shows that poorer households benefit the most from relaxing product frictions, while the wealthiest cohorts - being better able to self-insure - gain less. Panel B indicates that younger cohorts (up to age 65) see slightly larger gains than older cohorts, consistent with having more periods over which improved insurance can operate. Therefore, overall, welfare gains are tilted toward poorer households, and otherwise slightly to younger households.

Beyond comparing scenarios, we ask how much each product contributes to welfare and whether there are complementarities across products. Table 15 reports welfare changes when adjusting one product at a time within each scenario; the final row ("All") shows the effect of adjusting all three jointly. The first column compares the Baseline to a world with

Table 14: **Welfare Costs of Sub-optimal Portfolio Holdings: Rational Expectations**. Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	Baseline	**Scenario 3**	**Scenario 4**	**Scenario 5**
A. *Wealth Perc.*				
00–25	0.2%	2.7%	2.5%	4.8%
25–50	1.3%	2.7%	2.8%	4.5%
50–75	0.6%	1.4%	1.8%	3.6%
75–100	0.5%	1.1%	2.0%	3.4%
B. *Age Cohort*				
45–55	0.9%	2.4%	1.9%	3.9%
55–65	0.4%	2.1%	2.9%	5.0%
65–75	0.8%	1.6%	2.1%	3.3%
75–85	0.4%	1.2%	2.2%	4.3%

no insurance; the remaining columns report changes relative to the Baseline. The results suggest that, overall, products are neither strong complements nor strong substitutes: if anything, there is slight substitutability in the Baseline (with frictions), turning into mild complementarity once non-payment risk and wedges are removed.

6.4 The Role of Public Long-Term Care

In our baseline calibration, we observe low participation in private LTC insurance, consistent with the data. Large markups and nonpayment risks make these products unattractive, leading households to insure against LTC spending needs through other means. Wealthier households tend to self-insure using personal savings, explaining their low participation (see Table 10). Poorer households, in contrast, rely heavily on publicly provided LTC insurance. This program guarantees a minimum level of consumption, Ψ_G, but only on the condition that the household relinquishes all personal wealth to the state. As a result, it primarily benefits poorer households. In the model, 27% of households in the bottom wealth quartile use public LTC at some point in their lives, compared to just 8% in the top quartile.[16]

[16]Even in the second wealth quartile, 26% use public LTC at some point, because the guaranteed consumption level $\Psi_G = 77.43$ is high relative to their typical wealth holdings - wealth that would otherwise be depleted within about two years of LTC expenses.

Table 15: **Welfare Effects of Individual Products: Rational Expectations**. Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	Baseline	**Scenario 3**	**Scenario 4**	**Scenario 5**
A. *Insurance Ownership*				
LI	0.1%	0.9%	0.6%	1.6%
LTC	0.5%	0.2%	0.8%	1.0%
Ann	0.0%	0.2%	0.2%	0.7%
All	0.6%	1.2%	1.6%	3.4%

We next ask: how important is the generosity of public LTC for explaining the low take-up of private products? To assess this, we progressively lower the minimum consumption guarantee Ψ_G and examine how private LTC participation responds. Panel A of Table 16 reports the results. As Ψ_G declines, households substitute strongly into private LTC insurance. For example, when Ψ_G is reduced to 40, ownership rises from 13% to 44%. This increase is concentrated in the second wealth quartile, where participation rises by more than 50%, reflecting the fact that these households are the most reliant on public provision and thus most sensitive to its generosity.

Panel B shows the corresponding welfare effects. Lowering Ψ_G has large aggregate costs - about 4% of lifetime consumption when $\Psi_G = 40$ - and these costs are highly concentrated among poorer households. For the bottom half of the wealth distribution, welfare falls by 6%. Wealthier cohorts are far less affected, since they rely primarily on personal savings rather than public care.

Overall, the relatively generous terms of public LTC substantially dampen private LTC participation. Making public insurance less attractive raises private ownership, especially among lower-wealth households, but at the cost of significant welfare losses concentrated on those very households.

Table 16: **Insurance Demand, Welfare Effects of Public LTC**: Ψ_G is the consumption level offered by the government under Public LTC

	Baseline ($\Psi_G = 77.43$)	$\Psi_G = 70$	$\Psi_G = 60$	$\Psi_G = 50$	$\Psi_G = 40$
A. *Insurance Ownership*					
LI	55%	54%	54%	52%	51%
LTC	13%	18%	25%	35%	44%
Ann	2%	2%	2%	2%	2%
B. Δ *Personal Savings*					
25th Perc.	–	2%	6%	5%	14%
50th Perc.	–	1%	2%	2%	3%
75th Perc.	–	1%	1%	1%	2%
C. *Welfare Effects*					
0-25th Wealth Percentile	-	-0.5%	-1.4%	-2.8%	-5.7%
25-50th Wealth Percentile	-	-0.6%	-1.7%	-3.2%	-6.1%
50-75th Wealth Percentile	-	-0.5%	-1.3%	-1.9%	-2.7%
75-100th Wealth Percentile	-	-0.2%	-0.7%	-1.1%	-1.2%

6.5 The Role of Social Security

When estimating labor income life-cycle profiles for each quintile, we include social security payments. These act much like a publicly provided annuity, since benefits are received conditional on being alive. An open question, then, is how important social security provision is for insurance ownership, welfare, and wealth accumulation.

To answer this, we conduct the following experiment: we remove social security payments from labor income and re-define income quintiles using post-tax labor income (see Section C.2 for details). We then credit households with the present value of the social security payments they would otherwise have received.[17] In other words, the counterfactual asks: what if households received their social security entitlements up front as wealth rather than spread over the life cycle?

The results, shown in Table 17, indicate that households rebalance strongly out of life insurance after receiving this wealth windfall. Rather than purchasing life insurance, they self-insure with the additional assets. Because of this, average bequests increase by more than 23%, reflecting both intentional saving and incidental bequests: households hold on to

[17]We discount these payments using the same rate, r_I, that is used to price insurance products. Social security resembles an annuity, so unwinding it should use the annuity discount rate rather than the wealth return.

wealth in later life to replace the lost social security payments, and if they die earlier, these retained funds are passed on.

One might expect households to replace the lost public annuity with private annuities instead of wealth accumulation. They do not, for two reasons. First, annuities are expensive given market frictions such as non-payment risk and transaction costs. If these frictions are removed at the same time as social security, participation rises to 20% overall, and above 60% for those aged 70 and above, while average bequests increase by only 5%. This highlights how strongly annuity market inefficiencies suppress participation. Second, poorer households use the wealth windfall to bring consumption forward. Previously constrained by the zero-borrowing condition, they can now reallocate resources to earlier periods, which is highly beneficial for them.

Welfare results in Table 17 confirm this. Aggregate welfare improves by 2%, driven by large gains for poorer households: welfare rises by more than 10% for the bottom wealth quartile, as they use the windfall to shift consumption forward. For richer households, however, welfare falls, since they effectively liquidate the public annuity at an actuarially fair price but must then rely on costly annuities or self-insurance to protect late-life consumption.

Overall, replacing social security with an up-front wealth transfer reduces life insurance participation but raises welfare on average, driven by very large gains for poorer households who benefit from being able to bring consumption forward.

Table 17: **Insurance Demand, Welfare Costs of Baseline vs. no social security: Rational Expectations**

	Data	**No Insurance**	**Baseline**	**No Social Security**
A. *Insurance Ownership*				
LI	61%	0%	55%	19%
LTC	12%	0%	13%	14%
Ann	4%	0%	2%	4%
B. *Welfare Gains*				
Consumption Equivalent	-	-	0.6%	2.6%

6.6 Gross Margins of Insurers

We now quantify insurers' gross margins across scenarios. As in Section 5.3, observed quotes are modeled as actuarially fair values (conditional on age, health, and health transitions)

scaled by product-specific wedges. Here we measure the gross margin - i.e. the surplus-to-revenue ratio using realized cash-flows in the simulation. For product $k \in \{\text{LI, LTC, Annuity}\}$, define:

$$\text{PM}^k = \frac{\sum_i \sum_{t'=t}^{T} \frac{\text{rev}^k_{i,t'}}{(1+r_I)^{t'-t}} - \sum_i \sum_{t'=t}^{T} \frac{\text{cost}^k_{i,t'}}{(1+r_I)^{t'-t}}}{\sum_i \sum_{t'=t}^{T} \frac{\text{rev}^k_{i,t'}}{(1+r_I)^{t'-t}}}$$

where $(\text{rev}^k_{i,t'}, \text{cost}^k_{i,t'})$ are the payments made from the household i to the insurer and the realized payouts from the insurer to the household at date t'. Total revenue for product k is the present value of all premia/lump-sum payments made to the insurer across simulated households, discounted at r_I.[18] Costs are the present value of realized payouts. For products already held in the current observed period, we impute the associated historical payments so that we capture the payments already made by the household that would otherwise be missing from the simulations. [19]

Table 18 reports gross margins by product and scenario. In column (1) (Baseline), gross margins are 52%, 33%, and 31% for LI, LTC, and annuities, respectively. Setting price wedges to zero (column (2)) substantially lowers gross margins, showing that these wedges account for a large share of producer surplus. Removing non-payment risk (column (3)) further reduces gross margins by eliminating a source of surplus from unpaid promises. With both frictions removed (column (4)), gross margins fall essentially to zero for LTC and annuities, implying no room for costs at those prices.[20] For LI, a positive gross margin of about 9% remains, reflecting lapsation: many households allow policies to lapse as the bequest motive declines with age, generating premium revenue without a claim. Interpreted as a viability check, products are sustainable to the extent that non-claims costs as a share of revenue are below these reported gross margin rates; where the rate is near zero, viability would require negligible costs or alternative revenues.

[18] More concretely, there are 10,850 simulations: 1,085 in the UAS sample, simulated 10 times.

[19] If a household holds an annuity at time t, we impute a purchase price as what they would pay if newly purchased today. For LI/LTC, we impute annual premia from the prices locked in if purchased at age 45, and have been paid annually until the current period. An assumption is required on price because we do not observe in the UAS the price they pay for their LTC/LI insurance. Results are robust to reasonable alternatives.

[20] For annuities a small residual reflects limited take-up (10% in this scenario), yielding few revenue/cost flows.

Table 18: **Profit Margins of Insurers by Product**. Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	Baseline	**Scenario 3**	**Scenario 4**	**Scenario 5**
A. *Profit Margins*				
LI	52%	37%	33%	9%
LTC	33%	26%	25%	3%
Ann	31%	21%	26%	5%

7 Discussion

This paper provides three novel findings. First, we show that many consumers perceive significant risks in purchasing insurance products due to non-payment and default by insurance companies. Second, we provide compelling evidence that these perceived risks matter for insurance product choice: statistical and structural analyses both suggest ownership rates would be much higher if insurance products were risk free and actuarially-fair priced. Third, we show that the consumer welfare costs of these nonpayment risks and high prices are large, totaling over 4% of lifetime consumption. Overall, our study suggests incomplete markets and perceived risks are often overlooked but important determinants of insurance holdings by consumers. In addition, our study highlights the importance of measuring and modeling the products consumers perceive as available when conducting welfare analyses and designing policy to affect consumer choice.

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Perception of house price risk and homeownership," Technical Report, National Bureau of Economic Research 2018.

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti, "The Long-Term-Care Insurance Puzzle: Modeling and Measurement," *Vanguard Research Initiative Working Paper*, 2018.

_ , _ , _ , _ , and _ , "Long-Term-Care Utility and Late-in-Life Saving," *Journal of Political Economy*, 2020, *128* (6), 2375–2451.

Anderson, Kate, Eric French, and Tina Lam, "You Can't Take It with You: Asset Run-Down at the End of the Life Cycle," *Economic Perspectives*, 2004, *28* (3), 40.

Armona, Luis, Andreas Fuster, and Basit Zafar, "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment," *The Review of Economic Studies*, 2018.

Becker, Bo and Marcus Opp, "Regulatory reform and risk-taking: Replacing ratings," Technical Report, National Bureau of Economic Research 2013.

_ and Victoria Ivashina, "Reaching for yield in the bond market," *The Journal of Finance*, 2015, *70* (5), 1863–1902.

Bernheim, B Douglas, "How strong are bequest motives? Evidence based on estimates of the demand for life insurance and annuities," Technical Report, National Bureau of Economic Research 1989.

Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and Stephen P. Zeldes, "What Makes Annuitization More Appealing?," *Journal of Public Economics*, 2014, *116*, 2–16.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, 2008, *98:5*, 1887–1921.

Borella, Margherita, Mariacristina De Nardi, and Fang Yang, "The Importance of Modeling Income Taxes Over Time: U.S. Reforms and Outcomes," *Journal of the European Economic Association*, 2023, *21* (6), 2237–2286.

Braun, R. Anton, Karen A. Kopecky, and Tatyana Koreshkova, "Private Long-Term Care Insurance: Why is the Market so Small and Coverage Denials so Frequent?," *working paper*, 2016.

Brown, Jeffrey R, "Private pensions, mortality risk, and the decision to annuitize," *Journal of public Economics*, 2001, *82* (1), 29–62.

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

_ and _ , "Insuring Long-Term Care in the United States," *The Journal of Economic Perspectives*, 2011, pp. 119–141.

Conlon, John J, Laura Pilossoph, Matthew Wiswall, and Basit Zafar, "Labor market search with imperfect information and learning," Technical Report, National Bureau of Economic Research 2018.

Curcuru, Stephanie, John Heaton, Deborah Lucas, and Damien Moore, "Heterogeneity and portfolio choice: Theory and evidence," *Handbook of financial econometrics*, 2004, *1*, 337–382.

Davidoff, T, Jeff Brown, and P Diamond, "Annuities and Individual Welfare," *American Economic Review*, 2005, *95(5)*, 743–768.

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

Ellul, Andrew, Chotibhak Jotikasthira, Christian T Lundblad, and Yihui Wang, "Is historical cost accounting a panacea? Market stress, incentive distortions, and gains trading," *The Journal of Finance*, 2015, *70* (6), 2489–2538.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso, "Asset market participation and portfolio choice over the life-cycle," *Journal of Finance*, 2016.

_ , Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity in Returns to Wealth and the Measurement of Wealth Inequality," *The American Economic Review*, 2016, *106* (5), 651–655.

Fleck, Johannes, Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante, "Fiscal Progressivity of the U.S. Federal and State Governments," *Staff Report, Federal Reserve Bank of Minneapolis*, 2025, *663.*

French, Eric and John B. Jones, "The Effects of Health Insurance and Self Insurance on Retirement Behavior," *Econometrica*, 2011, *79* (3), 693–732.

Gomes, Francisco and Alexander Michaelides, "Asset pricing with limited risk sharing and heterogeneous agents," *Review of Financial Studies*, 2008, *21* (1), 415–448.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption over the Life Cycle," *Econometrica*, 2002, *70* (1), 47–89.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "Trusting the stock market," *the Journal of Finance*, 2008, *63* (6), 2557–2600.

Guvenen, Fatih, "A parsimonious macroeconomic model for asset pricing," *Econometrica*, 2009, *77* (6), 1711–1750.

Hendren, Nathaniel, "Private Information and Insurance Rejections," *Econometrica*, 2013, *81* (5), 1713–1762.

Hong, Jay H. and José-Vìctor Rìos-Rull, "Life Insurance and Household Consumption," *The American Economic Review*, 2012, *102* (7), 3701–3730.

Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes, "Expanding the Life-Cycle Model: Precautionary Saving and Public Policy," *The American Economic Review*, 1994, pp. 174–179.

Hubener, Andreas, Raimond Maurer, and Ralph Rogalla, "Optimal portfolio choice with annuities and life insurance for retired couples," *Review of Finance*, 2013, p. rfs046.

Hurd, Michael D, "Mortality Risk and Bequests," *Econometrica*, 1989, *57* (4), 779–813.

Hurd, Michael, Maarten Van Rooij, and Joachim Winter, "Stock market expectations of Dutch households," *Journal of Applied Econometrics*, 2011, *26* (3), 416–436.

Inkmann, Joachim and Alexander Michaelides, "Can the life insurance market provide evidence for a bequest motive?," *Journal of Risk and Insurance*, 2012, *79* (3), 671–695.

_ , Paula Lopes, and Alexander Michaelides, "How Deep is the Annuity Market Participation Puzzle?," *Review of Financial Studies*, 2011, *24* (1), 279–319.

Jang, Bong-Gyu, Hyeng Keun Koo, and Seyoung Park, "How Annuity Demand is Affected by Insurer Default Risk," *Seyoung, How Annuity Demand is Affected by Insurer Default Risk (November 5, 2013)*, 2013.

Koijen, Ralph S. J. and Motohiro Yogo, "The Cost of Financial Frictions for Life Insurers," *American Economic Review*, 2015, *105* (1), 445–475.

_ and _ , "Shadow Insurance," *Econometrica*, 2016, *84* (3), 1265–1287.

_ , Stijn Van Nieuwerburgh, and Motohiro Yogo, "Health and Mortality Delta: Assessing the Welfare Cost of Household Insurance Choice," *Journal of Finance*, 2016, *71* (2), 957–1010.

Kopczuk, Wojciech and Joseph P Lupton, "To leave or not to leave: The distribution of bequest motives," *The Review of Economic Studies*, 2007, *74* (1), 207–235.

Kopecky, Wojciech and Tatiana Koreshkova, "The Impact of Medical and Nursing Home Expenses and Social Insurance Policies on Savings and Inequality," *American Economic Journal: Macroeconomics*, 2014, *6* (3), 29–72.

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

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

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

Lopes, Paula and Alexander Michaelides, "Rare events and annuity market participation," *Finance Research Letters*, 2007, *4* (2), 82–91.

Love, David A, Michael G Palumbo, and Paul A Smith, "The trajectory of wealth in retirement," *Journal of Public Economics*, 2009, *93* (1), 191–208.

Luttmer, Erzo FP and Andrew A Samwick, "The welfare cost of perceived policy uncertainty: evidence from social security," *American Economic Review*, 2018, *108* (2), 275–307.

Manski, Charles F., "Measuring Expectations," *Econometrica*, 2004, *72* (5), 1329–1376.

Merrill, Craig B, Taylor D Nadauld, René M Stulz, and Shane Sherlund, "Did capital requirements and fair value accounting spark fire sales in distressed mortgage-backed securities?," Technical Report, National Bureau of Economic Research 2012.

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

Mommaerts, Corina, "Long-Term Care Insurance and the Family," *Long-term care*, 2015.

Nardi, Mariacristina De, "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 2004, *71* (3), 743–768.

Nardi, Mariacristina De, Eric French, and Jonathan Jones, "Differential Mortality, Uncertain Medical Expenses, and the Saving of Elderly Singles," *Journal of Political Economy*, 2010, *118*, 49–75.

_ , _ , and _ , "Medicaid Insurance in Old Age," *American Economic Review*, 2016, *106*, 3480–3520.

Palumbo, Michael, "Uncertain Medical Expenses and Precautionary Saving near the End of the Life Cycle," *Review of Economic Studies*, 1999, *66*, 395–421.

Pashchenko, Svetlana, "Accounting for non-annuitization," *Journal of Public Economics*, 2013, *98*, 53–67.

Poterba, James M, Steven F Venti, and David A Wise, "Health, education, and the post-retirement evolution of household assets," Technical Report, National Bureau of Economic Research 2013.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron, "Asset pricing with idiosyncratic risk and overlapping generations," *Review of Economic Dynamics*, 2007, *10* (4), 519–548.

Wiswall, Matthew and Basit Zafar, "Preference for the workplace, investment in human capital, and gender," *The Quarterly Journal of Economics*, 2017, *133* (1), 457–507.

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

A UAS Survey: Questions from Surveys

This outlines many of the questions that are asked in the survey.

Certainty Equivalent Below is an example of a question concerning annuities that is asked to a respondent in the survey. Each respondent answers such a question once at a randomly chosen certainty equivalent.

> The way you put balls into various bins shows that you expect to receive about [83]% of your annuity payment next year. It also shows that you could receive more or less than [83]% of the promised payment.
>
> Let's call this distribution of possible payments, as described by you using the bins and balls, your "uncertain payments." So, your uncertain payments are whatever payments you think you might receive next year.
>
> We are now interested in how you value having a contract with no uncertainty. Imagine a contract that is guaranteed to pay [75]% of your annuity payment with no risk of the insurance company not paying out as promised. This is like having all 20 balls on this certain percentage. This contract is unbreakable and cannot be changed by anybody. This contract has no risk, but is guaranteed to pay less than the full promised amount of your original contract.
>
> Would you rather have:
>
> - Guaranteed payment equal to [75]% of the annuity payment you are supposed to receive. or
> - Uncertain payments around an expectation of [83]% of the annuity payment you are supposed to receive?

Conditional Probability This is the question asked to derive beliefs on full default *conditional* on not receiving any payout this year.

Earlier you put [5] balls in the bin indicating "I will receive no payment at all". Suppose that this actually happened.

In the case that you receive no payment next year, what is the percent chance that the annuity never makes another payment at any point in the future? This might occur if the insurance company goes out of business, they claim you violated a clause in the contract, or they ruled the policy void for some other reason.

Beliefs Conditional on Aggregate State Respondents are told to imagine a stock market crash by being told the following prompt:

We are now interested in the chance that the insurance company will meet the policy's obligations under different economic conditions.

Specifically, suppose that the stock market decreases by [10/20/30]% next year.

They then answer the same questions as before for eliciting beliefs on non-payment risk and full default risk, but this time conditional on the stock market shock.

Effort to be Paid This is one of the questions asked to elicit the difficulty on claiming insurance nonpayments

In general, how much effort do you think you will need to put in to receive the annuity payments you are promised?

For example, this could include you or your family members doing paperwork, talking with claims officers, talking with doctors, hiring lawyers, or other such activities.

Please choose one of the following:

- No effort at all.
- A small amount of effort.
- A medium amount of effort.
- A large amount of effort.

B UAS Survey: Additional Regression Results

Table B.1: **Regression of Ownership on Risk Metrics**

	Own Annuity	Own Life	Own LTCI	Own Annuity	Own LIfe	Own LTCI
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Payment Exp. Val ($E[\omega^k]$)	-0.0018	0.0046***	0.0007	-0.0005	0.0045**	0.0006
	(0.212)	(0.001)	(0.111)	(0.373)	(0.003)	(0.181)
Full Def. Prob (q^k)	-0.0021***	-0.0015	-0.0023***	-0.0020***	-0.0013	-0.0022***
	(0.000)	(0.129)	(0.000)	(0.000)	(0.142)	(0.000)
Payment SD ($\sigma(\omega^k)$)	-0.0043**	-0.0006	-0.0009	-0.0029***	-0.0002	-0.0010
	(0.002)	(0.686)	(0.195)	(0.000)	(0.896)	(0.136)
Trust				0.0188	-0.0063	0.0162
				(0.091)	(0.758)	(0.241)
Cognitive Score				-0.0007	-0.0033	0.0004
				(0.747)	(0.271)	(0.852)
Financial Literacy Score				-0.0112	-0.0662*	-0.0083
				(0.459)	(0.019)	(0.609)
Numeracy Score				-0.0079	0.0207	-0.0240
				(0.560)	(0.319)	(0.101)
Experienced Fraud				0.0298	0.0545	-0.0031
				(0.549)	(0.375)	(0.941)
Risk Aversion				-0.0072	-0.0160	-0.0015
				(0.252)	(0.072)	(0.776)
Propensity to Plan				0.0137	-0.0013	0.0016
				(0.243)	(0.947)	(0.888)
Early Stock Returns				0.1474	-0.5123	-0.7936
				(0.757)	(0.441)	(0.122)
N	1055	1046	1040	1055	1046	1040
R^2	0.170	0.132	0.129	0.268	0.218	0.179
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.2: **Regression of Ownership on Risk Metrics in our sample with internally consistent responses (Table B)**

	Own Annuity (1)	Own Life (2)	Own LTCI (3)	Own Annuity (4)	Own LIfe (5)	Own LTCI (6)
Payment Exp. Val ($E[\omega^k]$)	-0.0017	0.0046***	0.0009	-0.0002	0.0038*	0.0008
	(0.270)	(0.001)	(0.050)	(0.770)	(0.015)	(0.118)
Full Def. Prob (q^k)	-0.0016**	-0.0015	-0.0022***	-0.0016**	-0.0012	-0.0022***
	(0.008)	(0.129)	(0.000)	(0.004)	(0.174)	(0.000)
Payment SD ($\sigma(\omega^k)$)	-0.0045**	-0.0006	-0.0010	-0.0027***	-0.0006	-0.0010
	(0.004)	(0.686)	(0.208)	(0.001)	(0.701)	(0.163)
Trust				0.0225	-0.0133	0.0214
				(0.058)	(0.551)	(0.151)
Cognitive Score				-0.0013	-0.0028	0.0015
				(0.585)	(0.393)	(0.558)
Financial Literacy Score				-0.0000	-0.0755*	0.0003
				(1.000)	(0.014)	(0.987)
Numeracy Score				-0.0214	0.0199	-0.0348*
				(0.161)	(0.385)	(0.042)
Experienced Fraud				0.0517	0.0800	0.0186
				(0.378)	(0.238)	(0.704)
Risk Aversion				-0.0085	-0.0142	-0.0025
				(0.210)	(0.147)	(0.660)
Propensity to Plan				0.0004	-0.0159	0.0024
				(0.974)	(0.444)	(0.843)
Early Stock Returns				-0.0824	-0.5042	-0.8261
				(0.869)	(0.457)	(0.144)
N	871	1046	856	871	862	856
R^2	0.141	0.132	0.142	0.246	0.234	0.198
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.3: **Tobit regressions of risk measures on demographic and behavioral controls, sample with internally consistent responses.**

	(1) ann_exp_value	(2) never_def_prob_life	(3) ann_sd	(4) life_exp_value	(5) never_def_prob_ann	(6) life_sd	(7) never_def_prob_ltci	(8) ltci_exp_value	(9) ltci_sd
main									
trust	0.1855	-0.2577	-0.7070	1.3690	1.3169	-0.3861	-2.0044*	0.3406	-0.4151
	(0.855)	(0.812)	(0.446)	(0.261)	(0.157)	(0.694)	(0.043)	(0.752)	(0.605)
cog_score	0.2363	0.4034*	-0.0933	0.3813	0.1936	-0.1091	0.4867**	0.1479	-0.0176
	(0.169)	(0.021)	(0.565)	(0.090)	(0.206)	(0.542)	(0.003)	(0.393)	(0.891)
finlit_score2	5.9770**	4.7221*	-2.0708	0.6433	1.9868	0.7930	1.5694	3.6353	-0.5710
	(0.002)	(0.011)	(0.254)	(0.782)	(0.221)	(0.647)	(0.303)	(0.052)	(0.669)
num_score	0.8866	2.1226	-0.3404	1.8751	1.1263	-1.6830	0.8036	2.8365*	-2.7947**
	(0.514)	(0.113)	(0.781)	(0.274)	(0.318)	(0.214)	(0.494)	(0.038)	(0.004)
exp_fraud2	3.0947	-10.1457*	-3.3082	-1.0436	-7.0727*	0.0553	-9.9046**	-1.3876	-2.3892
	(0.553)	(0.043)	(0.454)	(0.818)	(0.045)	(0.989)	(0.005)	(0.798)	(0.530)
risk_subj	0.5937	0.2764	-0.4378	-0.7318	-0.0563	0.8557	-0.1707	-0.2562	0.2187
	(0.311)	(0.580)	(0.367)	(0.275)	(0.911)	(0.099)	(0.730)	(0.660)	(0.610)
plan_ahead	-1.2719	0.3207	1.3106	-0.2452	0.1293	1.5910	1.6552	-1.8149	-0.4275
	(0.316)	(0.785)	(0.247)	(0.860)	(0.906)	(0.161)	(0.129)	(0.183)	(0.645)
N	872	888	872	868	888	864	889	857	857
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C Estimated Inputs for Model

In this section we describe our estimation of health costs, income profiles and health transitions that serve as inputs into the structural model.

C.1 Health Cost Shocks

To estimate the health cost distribution, we define expenses as out-of-pocket medical expenses using the variable r10oopmd in the RAND HRS in 2014. We subtract o-o-p home health services and nursing home payments because they are the type of expenses specific to the LTC state (that we have separately). We then consider the sub-sample of respondents in the 2014 sample who are not part of a couple, and those between the ages of 55 and 100. We estimate the distribution following the process outlined in Ameriks et al. (2020). To estimate the mean of the health cost distribution, we regress the log of medical expenses against a quartic of age, health status, sex and health interacted with age. Using the residuals from this first regression, we regress the squared residuals on the same set of state variables as in the first regression to find the conditional variance of medical expenses. The resulting mean and variance estimates parameterize the lognormal distribution for medical expenses conditional on age, sex, and health status. We then discretize the error term, using the Tauchen method to generate the medical expense process.

Figures C.1 and C.2 plot the mean health cost shock over the life cycle for both men

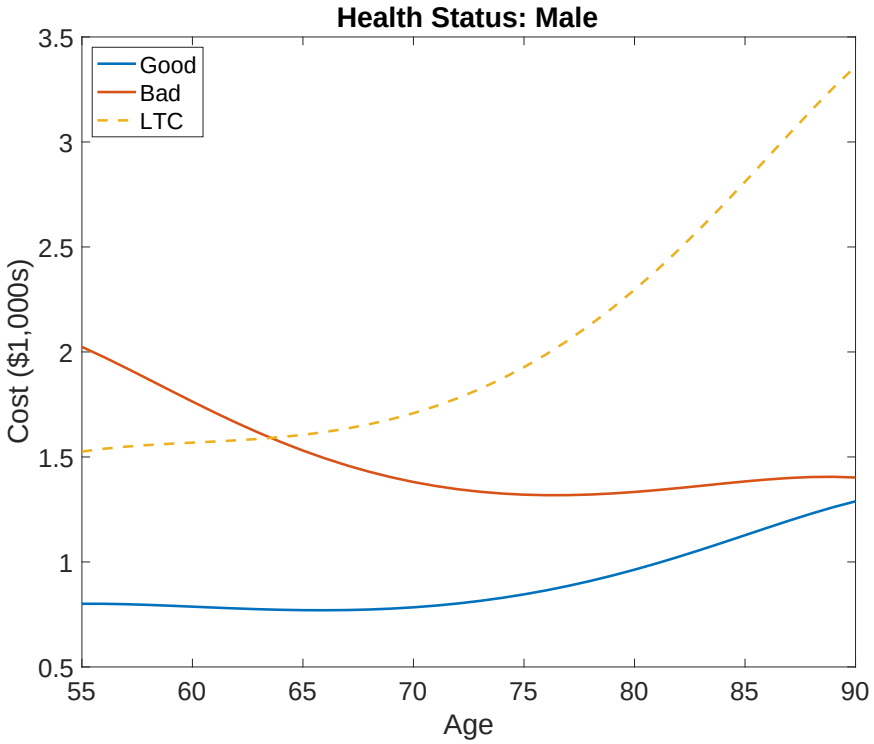

Figure C.1: **Mean Health Cost Shock Profile: Male**

and women. Men in poor health at age 65 spend about $750 more on average than men in good health. Later in life as men age, those in LTC spend on average about $2000 more per year, excluding any nursing home or home health services. As a result, the LTC state can potentially be a very costly health state to arrive at.

C.2 Income Process

We estimate a deterministic labor income profile of income for five separate income profiles. We define income using the RAND HRS dataset as the sum of income from earnings, public and private pension and annuities, social security income, disability and retirement, unemployment compensation, and other government transfers.[21] We evaluate this at the household level by summing up the income of both the respondent and the spouse. We take the income for each household every time we observe them between wave 2 (1994) and wave 12 (2014) of the RAND HRS dataset. Income is evaluated at constant 2014 dollars.

[21] More specifically, our measure of labor income is the sum of Respondent and spouse earnings (RwIEARN, SwIEARN), pensions and annuities (RwIPENA, SwIPENA), SSI and Social Security Disability (RwISSDI, SwISSDI), Social Security retirement (RwISRET, SwISRET), unemployment and workers compensation (RwIUNWC, SwIUNWC) and other government transfers (RwIGXFR, SwIGXFR).

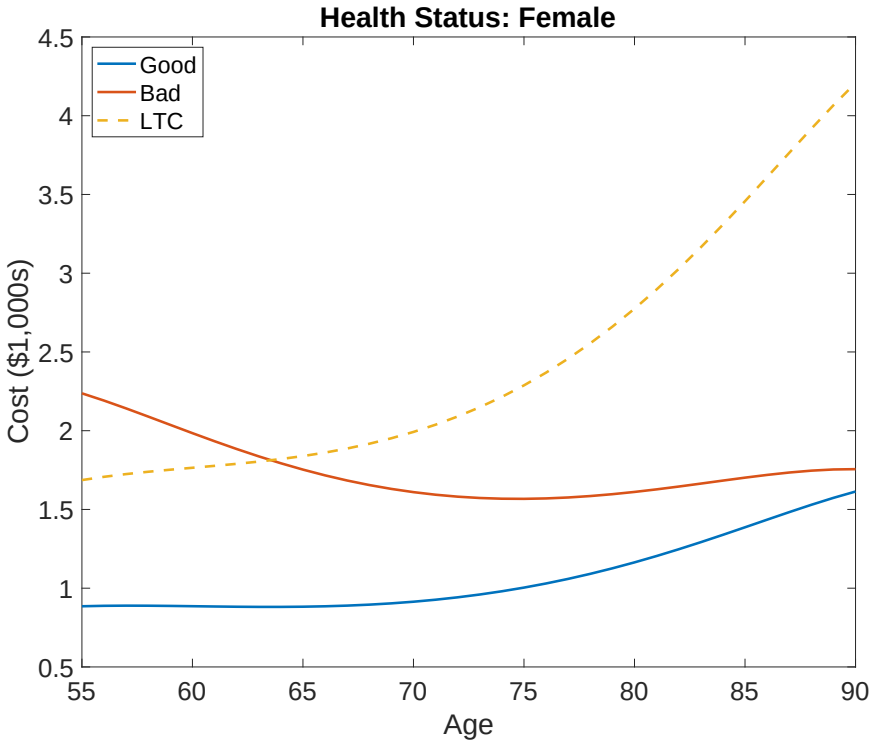

Figure C.2: **Mean Health Cost Shock Profile: Female**

Our first step is to make adjustments for taxes on labor income. This is important because, if taxes are progressive, they reduce the disparity in disposable incomes across quintiles and cohorts. A simple way to do this is to follow Borella et al. (2023), who show that the following tax function can capture the federal taxes paid by households very well:

$$T(Y) = Y - (1 - \lambda)Y^{1-\tau}$$

They show that, using PSID data along with carefully crafted implied tax payments (for example, using TAXSIM by NBER), they can estimate (λ, τ) for each year. The estimated parameters vary across years because of changing federal tax systems in the US. We will use these year-based estimates in evaluating after-tax income for households from the RAND HRS. This does exclude state-income taxes. Fleck et al. (2025) show, however, that state taxes affect the level but not the progressivity of taxes on average. As a result, federal taxes are enough to capture the progressivity of income taxation that we care about.

Next, we want to define the income quintile that a given household i lies in. We start by defining at each wave (waves 2-12) the income quintile a given household lies in at that time. We start by splitting the sample into 5-year age brackets (ages 45-49, 50-54 etc.) and defin-

ing the income quintile thresholds as the $(20, 40, 60, 80)^{\text{th}}$ percentiles for each age bracket. Then, for each household i at wave w, we assign them an income quintile depending on what bracket their income lies in within their age. Then, we take the average income quintile of a household across all waves that we observe them, rounded to the nearest whole number, which we define as household i's *permanent income quintile (PIQ)*.

Armed with this PIQ measure, we run the following panel regression across all waves w

$$\begin{aligned} log(inc_{i,w}) = \alpha_i + \beta_1 Age_{i,w} + \beta_2 Age_{i,w}^2 + \beta_3 Age_{i,w}^4 + \beta_5 Male_i + \beta_6 (Male_i * Age_{i,w}) \\ + \beta_7 PIQ_i + \beta_8 * (PIQ_i * Age_{i,w}) + \epsilon_{it} \end{aligned}$$

where $inc_{i,w}$ is our measure of income for household i in wave w. The key idea here is that we have fixed effects at the individual household level. This is to control for cohort fixed effects that may otherwise confound results. In particular, if aggregate real income grows over time, we will over-estimate the declining profile of labor income.

The final step is to obtain a measure of predicted income for each combination of age $a \in \{45, ..., 99\}$, permanent income quintile $PIQ \in \{1, 2, 3, 4, 5\}$, and gender $g \in \{m, f\}$. We use the regression results in the following way: we take the average fixed effect across all households for a given (a, PIQ, g) combination. We then plug in this measure as a constant on the right-hand side and plug in the combination (a, PIQ, g) into the right-hand side to get a predicted value $\hat{inc}_{a,PIQ,g}$. This is what we use for the deterministic labor income profile.

The profiles that we feed into the model for each quintile are shown on the left-hand side of Figure C.3. As we can see, there is significant heterogeneity across quintiles, and all exhibit a distinct downward-sloping profile in later years.

In one of our experiments, we consider what happens if households instead receive labor income exclusive of social security income. To derive these profiles, we repeat the exact same steps as before, except that at the very beginning we define income as excluding all social security income sources. This produces the profile conveyed on the right-hand side of Figure C.3. As we can see, labor income profiles shift down, with there being a larger proportional effect on the lower income quintiles.

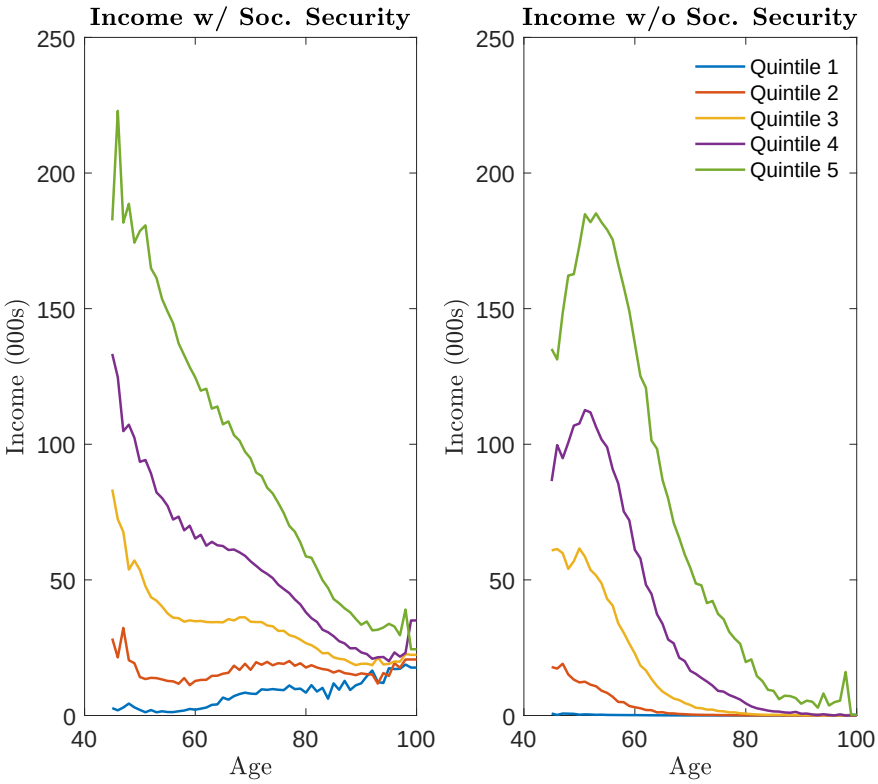

Figure C.3: **Deterministic Labor Income Profiles by Quintile**

C.3 Health Transitions

Health transitions are estimated using the RAND HRS waves 2 to 12. We construct the defined health states $s \in \{0, 1, 2\}$ from two sets of questions. The first makes use of self-reported subjective health-status questions to classify respondents into good ($s = 0$) or bad ($s = 1$) health.[22] The second set of questions is used to determine whether an individual is in the LTC ($s = 2$) health state. This set of questions presents five ADLs and asks whether respondents receive help with any of the five activities. If the respondent answers yes to at least two of these questions, then we define that respondent to be in the LTC health state.

Armed with these well-defined health states, we then estimate a health transition matrix that is conditional on the individual's age, t, sex, f, and permanent income type. In particular, we allow for the health transition to differ between those whose permanent income quintile is 1,2 or 3 ("Poor") and those whose permanent income quintile is 4 or 5 ("Rich").

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

This is to allow the model to reflect the fact that wealthier households exhibit more favorable health transitions in the data.

Figures C.4 and C.5 display the estimated health transitions for the "Poor" and "Rich", respectively. Notice how the "Rich" have slightly more favorable health transitions - they are more likely to stay healthy throughout their life, and conditional on being in an LTC state, they are more likely to stay alive.

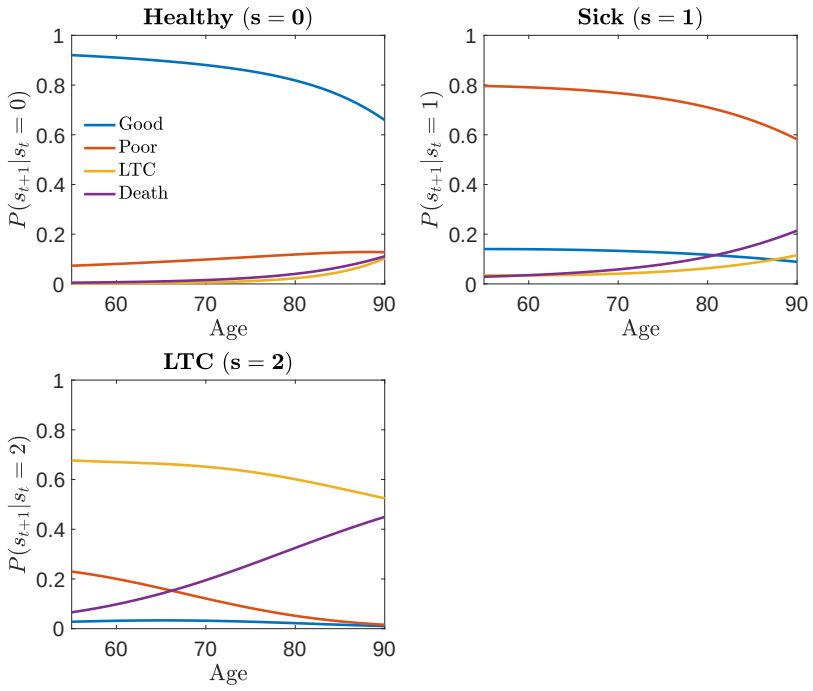

Figure C.4: **Health Transition Profiles: Poor** ($PIQ \in \{1, 2, 3\}$)

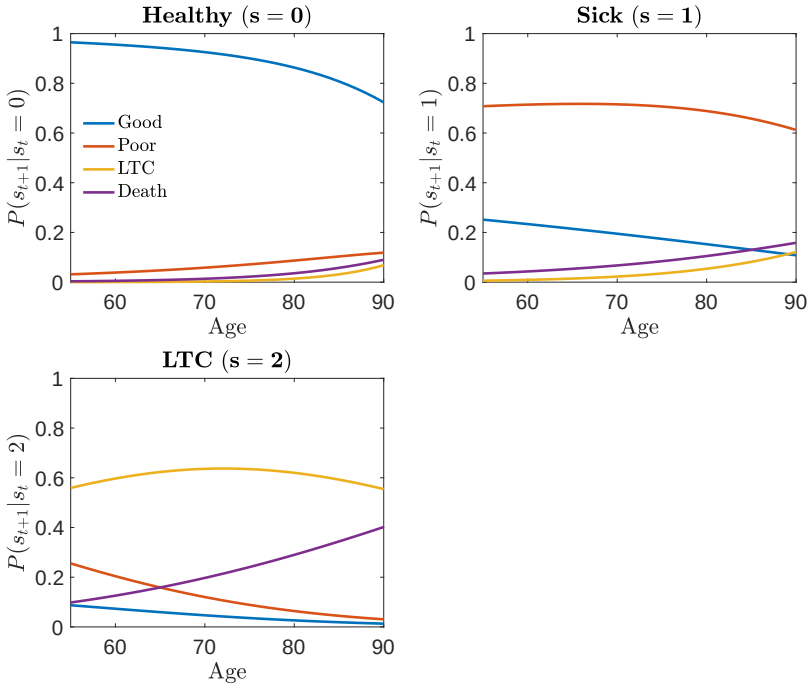

Figure C.5: **Health Transition Profiles: Rich** ($PIQ \in \{4, 5\}$)

D Simulation and Estimation of Jointly-Determined Model Parameters (TBD)

In this section we describe how we simulate the model, and then how we use these simulations to estimate the model by targeting sample moments of wealth, insurance ownership rates, and strategic survey question answers from Ameriks et al. (2020).

D.1 Model Simulations

We firstly solve the model to compute optimal decision rules for wealth, insurance ownership, consumption and government care take-up, given a chosen set of parameters θ. The rules will be a function of individual state variables $\Omega_i = \{B_i, D_i^{ANN}, D_i^{LI}, D_i^{LTC}, X_i, , p_i^{LTCI}, p_i^{LI}\}$, where X_i includes age, health status and permanent income quintile, and the model is solved by conducting a 6-dimensional grid search method over the remaining state variables. Once we have the policy functions defined over Ω_i, we then simulate the model for all 1,085 individuals i in our UAS sample. To set the initial conditions, we take from the UAS sample the empirical counterparts of the state variables, $\hat{\Omega}_i$ for each individual. We simulate each individual 100 times, drawing relevant shocks $\hat{\Xi}$ from the distributions of full default, annual

payment, health cost shocks as well as the health transitions, and simulate the behaviour implied by the optimal policies.[23] We then aggregate these individual behaviours to construct simulated moments from the model, $o(\hat{\Xi}, \theta, \hat{\Omega})$, for which we will compare to empirical targets. When aggregating in the model, we carry the sample weight of agent i with them across their life span in the simulation. As a result, when evaluating life insurance participation rates at age 57 for example, we take the weighted average participation rate across all simulated observations of agents at age 57, where the weight of each observation is the UAS weight of the agent i that is assigned to them at the initial observation date.

Simulating this model is computationally intensive. First, solving the policy functions from ages 45 to 99 via backward induction requires a fine grid search across six dimensions. Then, the model must be simulated 100 times for each of the 1,085 agents. Solving the model for a single parameter set θ takes approximately 36 hours. To manage this, we leverage over 200 computing clusters, enabling extensive parallelization. This setup significantly reduces computation time and allows us to evaluate policy functions across many combinations of θ when estimating the model.

We define below our strategy for choosing optimal parameters θ, and the empirical moments we target.

D.2 Estimation Strategy

We develop an estimation strategy similar to those used in De Nardi et al. (2010), Gourinchas and Parker (2002) and French and Jones (2011) to estimate the parameter set $\Theta := \{\beta, \theta_{3,45}, \sigma, m, r, r_I\}$. All other parameters of the model, outlined in Tables 6 and 7 are externally estimated without the use of the structural model. There are three sets of moments that we use to estimate the model. Firstly, we target wealth percentiles across the life cycle, which are commonly used in life-cycle savings models. Secondly, we target average participation rates for each of the three insurance products that we analyse. Finally, we target a set of moments derived from SSQ responses drawn from Ameriks et al. (2020).

Let $Z = (Z_i)_{i=1}^{I}$ be the collection of measurements for all individuals, which includes behavioural responses (i.e. wealth and insurance holdings), SSQ responses and state variables. The moments we use to estimate the model are the difference between statistics gener-

[23] Results are essentially equivalent if we simulate each individual just 10 times.

ated by the model ($o(\hat{\Xi}, \Theta, \hat{\Omega})$) and empirical data ($s(Z)$, which we define as $g(\hat{\Xi}, \Theta, \hat{\Omega}) = o(\hat{\Xi}, \Theta, \hat{\Omega}) - s(Z)$. We then solve the following problem:

$$\hat{\Theta} = \arg\min_{\Theta} g(\hat{\Xi}, \Theta, \hat{\Omega})' W g(\hat{\Xi}, \Theta, \hat{\Omega})$$

where W is the weighting matrix. We assume that the weights are identical within the set of moments for each of the three categories, where the sum of the weights in each category are equal. We now define the three categories of moments below:

D.2.1 Wealth Moments

The first set of moments consists of wealth percentiles that are conditional on a subset of state variables, in this case age a. We define the moment conditions as the 25th, 50th and 75th percentiles conditional on 5-year age brackets (45-49,50-54,...,80-84). We stop at age 85 as the sample size becomes small. This generates 8 age intervals, and so given three targeted wealth percentiles this produces 24 wealth moments. We then simply compare the model-implied values against the empirical counterparts.

The wealth moments that we construct from the cross-sectional data are taken from the HRS. The argument here is that the UAS is a representative sample of the HRS. As a result, the HRS will be a much less noisy measure of the empirical wealth distribution. However, as we can see in Table 1, the UAS sample is unbalanced, particularly along the dimensions of (i) marital status and (ii) retired status.

In response, in order to better compare across surveys, we should also have the two surveys comparable across these two dimensions. So we considered the following procedure: let's consider the cohort aged 60-65. Then suppose we observe that, in the HRS, $(y_{RET}^{60-65}, y_{MAR}^{60-65})$ are the fractions of these households who are retired and married, respectively, and similarly $(x_{RET}^{60-65}, x_{MAR}^{60-65})$ for the UAS. Then, in the HRS, you make the following adjustments to the weights:

1. (Retired, Married): scale by $\frac{x_{RET}^{60-65} x_{MAR}^{60-65}}{y_{RET}^{60-65} y_{MAR}^{60-65}}$
2. (Retired, Not Married): scale by $\frac{x_{RET}^{60-65}(1-x_{MAR}^{60-65})}{y_{RET}^{60-65}(1-y_{MAR}^{60-65})}$
3. (Not Retired, Married): scale by $\frac{(1-x_{RET}^{60-65})x_{MAR}^{60-65}}{(1-y_{RET}^{60-65})y_{MAR}^{60-65}}$
4. (Not Retired, Not Married): scale by $\frac{(1-x_{RET}^{60-65})(1-x_{MAR}^{60-65})}{(1-y_{RET}^{60-65})(1-y_{MAR}^{60-65})}$

The intuition here is that you are shifting the HRS dataset that is being compared to the UAS from being balanced to being unbalanced, but precisely in the way that the UAS is unbalanced. The outcome of this is that we get Figure 7 as our new comparison of datasets. The key difference is that now the average level of wealth is lower at 526k i.e. closer to the UAS level.

D.2.2 Insurance Participation Rates

This is very similar to the wealth moments. The key difference here is that we are going to simply target the aggregate participation rates of life insurance, annuities and LTC insurance, leaving untargeted the life-cycle profile of insurance ownership in the data. Again, we target the aggregate participation rates from the HRS, and use the adjusted sample weights described above. This generates participation rates in the data of $(61\%, 12\%, 4\%)$ for Life insurance, LTC insurance and annuities, respectively.

D.2.3 SSQ Moments

This final moment set is constructed entirely from SSQ responses of the sample of households in Ameriks et al. (2020). The key difference here is that we fix the parameter values of $(\theta_2, \bar{\theta}_3, \kappa_2, \kappa_3, \psi_G)$ to those estimated by the paper as their posterior means, leaving us only with σ as the estimated variable. The intuition for this is that σ is the least well-identified parameter from these SSQ moments, while the remaining parameters are very well identified by these questions.

As moments, we target the model-implied responses to the empirical mean of each of the nine SSQ variants, as a fraction of the maximum response. The results are shown in Figure D.6 below.

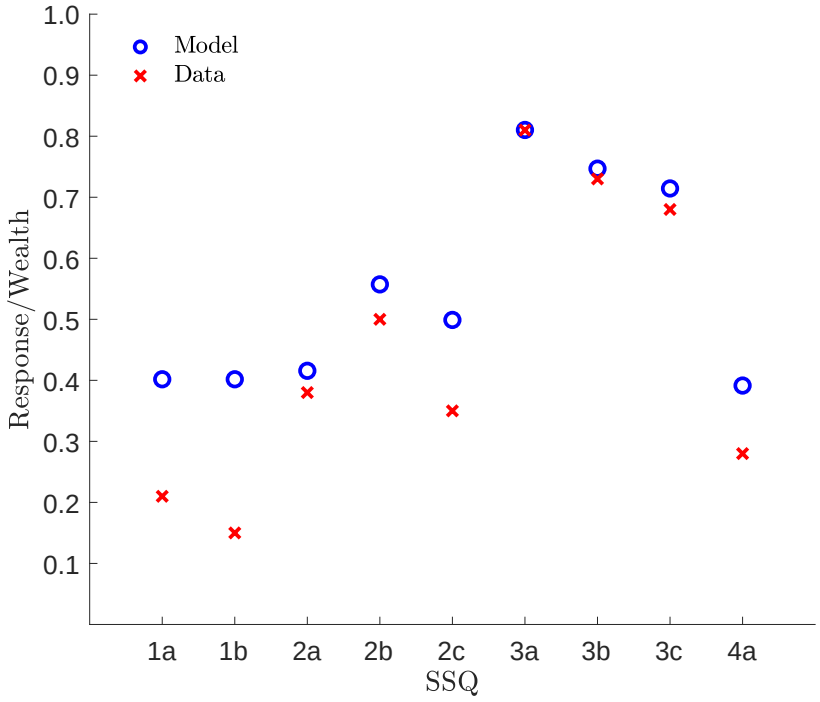

Figure D.6: **Mean SSQ Responses in Model (circles) and Data (crosses)**

D.3 Computing Consumption-Equivalent Welfare

Consumption-equivalent welfare. For scenario $S \in \{A, B\}$ and individual i, let the (simulated) lifetime utility be

$$\widehat{V}_i(S) \equiv \frac{1}{N} \sum_{n=1}^{N} \left[\sum_{t=0}^{T} \beta^t \, v\big(c_{int}^S\big) \right],$$

where n indexes each simulation of health states, health cost shocks and insurance default outcomes. The consumption-equivalent welfare loss of B relative to A for i is the unique $\Delta_i^C \in \mathbb{R}$ that solves

$$\widehat{V}_i(A) \;=\; \frac{1}{N} \sum_{n=1}^{N} \left[\sum_{t=0}^{T} \beta^t \, v\big((1 - \Delta_i^C)\, c_{int}^B\big) \right].$$

Once we compute Δ_i^C for each agent i, we then report the population consumption-equivalent welfare as the UAS-weighted mean $\Delta^C \equiv \sum_i w_i^{UAS} \Delta_i^C \,/\, \sum_i w_i^{UAS}$.

Throughout the paper, we compute welfare effects across various counterfactual scenarios using consumption-equivalent welfare as our evaluation metric. Consider two scenarios, indexed by $\{A, B\}$. For each agent i, we calculate the percentage reduction in consumption

(and, when applicable, bequests), denoted by Δ_i^C, that the agent would be willing to accept in order to be indifferent between scenarios B and A. To do this, we scale consumption by $(1-\Delta_i^C)$ across all simulations in scenario B, then compute the average lifetime utility across those simulations. We choose Δ_i^C such that this average lifetime utility in scenario B equals the average lifetime utility of the same agent across simulations in scenario A. Finally, we aggregate the resulting welfare effects using the UAS weights assigned to each agent.

D.4 Model Sensitivity to Parameters

In conducting the estimation exercise, we learn the sensitivity of our model moments to the parameter set Θ that we are estimating. We discuss each of these parameters below.

D.4.1 Discount Factor - β

The discount factor β is very effective in targeting the wealth distribution. We need it to be relatively high in order to match the flat wealth profiles that do not fall considerably later in life. A higher β also has a strong effect in reducing life insurance holdings. This is because, due to the strong calibrated bequest motive earlier vs. later in life, a higher β means that households place relatively less weight on earlier years where the bequest motive is strongest. A higher β also on the margin reduces the incentive to hold LTC insurance, but this asset is more state contingent and so the effect is weak.

D.4.2 Risk Aversion - σ

The risk aversion coefficient σ very much has a targeted effect on insurance ownership instead of wealth profiles. A higher σ has a strong positive effect on both LTC insurance and life insurance. This is because LTC and life insurance are held for precautionary reasons incase an unlikely event occurs (need long-term care or you die young).

In our calibration, we need σ to be relatively low in order to be close in matching LTC participation rates. Otherwise there is too much demand for LTC. This is because, when you are risk-averse, you really dislike the state-contingent outcome of arriving in the LTC health state ($s = 2$) without any private insurance because you are likely to request government care that is very costly on utility.

D.4.3 Bequest Motive When Young - $\theta_{3,45}$

This controls the bequest motive of young households. The main goal here is to ensure that the bequest motive for young households is strong enough such that their life insurance

participation is high. Doing this requires raising $\theta_{3,45}$. There is a balancing act, however, because this also raises LTC ownership when young. This is driven by the fact that, conditional on being in an LTC state when young, this significantly raises the relative likelihood of dying, as we can see with the conditional probabilities in Figure C.4. As a result, LTC insurance acts similarly to life insurance b providing incidental bequests when young.

D.4.4 Parameter m

This parameter m controls the shape of the bequest motive over the life cycle, governed by equation (1). For m that is very close to 0, it approximates a linear function, but as m increases it tilts the bequest motive when young higher and creates a steeper declining slope for later ages. Our calibration is for a relatively high m of 2.5 because it is very helpful in getting the right shape for the life cycle profile of life insurance ownership. If the shape of the bequest motive was linear (i.e. $m \to 0$), life insurance ownership would fall too quickly with age, hence why we need a higher m.

Of course, again, this does tend to increase LTC ownership when young, so the value of m can't be placed too high. It could also cause life insurance ownership to rise as they age from 45-60 because it will keep the bequest motive strong exactly when death becomes more likely, thereby raising the desire to hold the insurance.

D.4.5 Parameter r

Here we consider what happens when we raise the return on bonds r and the discount rate r_I together (next part isolates change in r_I). This makes savings as a whole more attractive, and helps a lot with achieving an upward trajectory of wealth earlier in life that we see in the data.

D.4.6 Parameter r_I

We now consider what happens if we lower the discount rate r_I while keeping the return on bonds unchanged. From the perspective of the household, this simply raises the price on all insurance products, ceteris paribas. In other words, it is roughly equivalent to an additional transaction cost on all insurance products. This of course then helps to match the average participation rate across all three insurance products.

E Alternative Results: Non-Rational Expectations

This subsection lays out a selection of results from running the same set of experiments as in the main text, but under the alternative assumption that beliefs of non-payment risk are non-rational - in particular, the insurer always ends up paying out what is promised, regardless of what the household believes.

E.1 Portfolio Choice by Scenario: Non-Rational Expectations

Table E.4 repeats the same exercise as under Table 12 but under non-rational expectations. The numbers are very close to each other - what ultimately matters for portfolio choice is the belief of agents whether or not they are true. The slight differences arise because the evolution of wealth in the simulations differ - households in this setting get paid out more on average than what they were expecting. This raises average wealth in the simulations.

Table E.4: **Insurance Demand and Welfare Costs of Sub-optimal Portfolio Holdings: Non-Rational Expectations**. Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	Data	No Insurance	Baseline	Scenario 3	Scenario 4	Scenario 5
A. *Insurance Ownership*						
LI	61%	0.0%	56%	72%	73%	84%
LTC	12%	0.0%	15%	30%	43%	55%
Ann	4%	0.0%	2%	3%	8%	10%
B. Δ *Personal Savings*						
25th Perc.	–	–	−38%	−47%	−53%	−62%
50th Perc.	–	–	−17%	−23%	−33%	−42%
75th Perc.	–	–	−6%	−14%	−23%	−32%

E.2 Welfare Effects by Scenario: Non-Rational Expectations

Table E.5 repeats the same exercises as under Table 14 but under non-rational expectations. The conclusions that we draw are the same - the poorer and younger households benefit disproportionately from the relaxation of frictions in the insurance market: transaction costs and (perceived) non-payment risks.

Table E.5: **Welfare Costs of Sub-optimal Portfolio Holdings: Non-Rational Expectations**. Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	Baseline	**Scenario 3**	**Scenario 4**	**Scenario 5**
A. *Wealth Perc.*				
00–25	1.4%	4.3%	2.5%	4.8%
25–50	2.4%	4.3%	2.8%	4.5%
50–75	1.4%	2.6%	1.8%	3.6%
75–100	1.2%	2.5%	2.0%	3.4%
B. *Age Cohort*				
45–55	1.6%	3.6%	1.9%	3.9%
55–65	2.0%	4.2%	2.9%	5.0%
65–75	1.2%	2.3%	2.1%	3.3%
75–85	1.4%	2.8%	2.2%	4.3%

E.3 Welfare Effects by Product: Non-Rational Expectations

Table E.6 repeats the same exercise as in Table 15 but again under non-rational expectations. The results indicate broadly the same conclusions - the products appear to be neither substitutes nor complements to each other.

E.4 Role of Public Long-Term Care: Non-Rational Expectations

Table E.7 runs the same experiment as in Table 16 but under the assumption of non-rational expectations. Insurance ownership closely mirrors the behavior seen under the rational expectations framework - beliefs primarily drive portfolio choices. Welfare costs continue to rise as public LTC (long-term care) provision deteriorates, with poorer households still bearing the brunt of the burden. The key difference here is that the increase in welfare costs is less steep as public provision worsens. Intuitively, as conditions decline across the Cases, households respond by increasing their LTC insurance holdings, even though they expect only partial payouts. However, because the insurance products ultimately pay out in full, these higher-than-expected payouts help to soften the welfare costs.

Table E.6: **Welfare Effects of Individual Products: Rational Expectations**. Scenario 3: "No Price Wedges"; Scenario 4: "No Non-Payment Risk"; Scenario 5: "No Price Wedges and Non-Payment Risk".

	Baseline	**Scenario 3**	**Scenario 4**	**Scenario 5**
A. *Insurance Ownership*				
LI	0.7%	1.0%	0.1%	1.1%
LTC	0.9%	0.4%	0.3%	0.6%
Ann	0.0%	0.4%	0.2%	0.7%
All	1.6%	1.7%	0.7%	2.4%

Table E.7: **Insurance Demand, Welfare Effects of Public LTC: Non-Rational Expectations** : Ψ_G is the consumption level offered by the government under Public LTC

	Baseline ($\Psi_G = 77.43$)	$\Psi_G = 70$	$\Psi_G = 60$	$\Psi_G = 50$	$\Psi_G = 40$
A. *Insurance Ownership*					
LI	56%	55%	55%	53%	52%
LTC	15%	19%	26%	37%	45%
Ann	2%	2%	2%	2%	2%
B. *Welfare Effects*					
0-25th Wealth Percentile	-	-0.3%	-1.2%	-2.4%	-4.6%
25-50th Wealth Percentile	-	-0.5%	-1.4%	-2.6%	-4.7%
50-75th Wealth Percentile	-	-0.5%	-1.1%	-1.3%	-1.6%
75-100th Wealth Percentile	-	-0.1%	-0.3%	-0.3%	-0.3%

E.5 Role of Social Security: Non-Rational Expectations

This runs the same set of experiments as under Table 17 but under non-rational expectations - what happens to welfare and portfolio choice if we remove social security payments from household income and give it back to them lump-sum today. Results are shown in Table E.8. Evidently, regardless of what we assume for expectations, participation rates are unchanged, and the welfare gain from receiving social security lump-sum today remains approximately 2% of lifetime consumption.

Table E.8: **Insurance Demand, Welfare Costs of Baseline vs. no social security: Non-Rational Expectations**

	Data	No Insurance	Baseline	No Social Security
A. *Insurance Ownership*				
LI	61%	0%	56%	20%
LTC	12%	0%	15%	14%
Ann	4%	0%	2%	4%
B. *Welfare Gains*				
Consumption Equivalent	–	–	1.6%	3.3%