NBER WORKING PAPER SERIES

RISKY INSURANCE:
LIFE-CYCLE INSURANCE PORTFOLIO CHOICE WITH INCOMPLETE MARKETS

Joseph S. Briggs
Ciaran Rogers
Christopher Tonetti

Working Paper 35122
http://www.nber.org/papers/w35122

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2026

We thank Mariacristina De Nardi, Francisco Gomes, Glenn Harrison, Sasha Indarte, Kieran Larkin, Francois LeGrand, Kim Peijnenburg, Kjetil Storesletten, Johannes Stroebel, Lesley Turner, Joe Vavra, Moto Yogo, and seminar participants at many conferences, universities, and research centers for useful comments and discussions. We are especially grateful to Andrew Caplin for his insightful counsel and his financial support in fielding the survey. This research received much appreciated financial support from the Alfred P. Sloan Foundation through the NBER Household Finance Small Grant program. The project described in this paper relies on data from surveys administered by the Understanding America Study, which is maintained by the Center for Economic and Social Research (CESR) at the University of Southern California. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of USC or UAS. The analysis and conclusions set forth here are those of the authors and do not necessarily indicate a concurrence by the Goldman Sachs Group, Inc. or the National Bureau of Economic Research.

Risky Insurance: Life-Cycle Insurance Portfolio Choice with Incomplete Markets
Joseph S. Briggs, Ciaran Rogers, and Christopher Tonetti
NBER Working Paper No. 35122
April 2026
JEL No. D14, E21, G22

ABSTRACT

We study consumer demand for savings, life insurance, annuities, and long-term care insurance using novel survey data and a structural life-cycle model. We document that individuals perceive substantial insurance nonpayment risk, and these beliefs predict ownership. Embedding elicited beliefs into an incomplete-markets model alongside additional real-world insurance features, we match empirical patterns of low participation. Relative to a no-insurance benchmark, access to existing imperfect insurance reduces median wealth by 16% and generates a modest 0.6% welfare gain. Eliminating nonpayment risk would substantially increase insurance ownership, yield a further 11% decline in median savings, and generate an additional 1.7% welfare gain.

Joseph S. Briggs
Goldman Sachs
jsbriggs87@gmail.com

Ciaran Rogers
HEC Paris
rogers@hec.fr

Christopher Tonetti
Stanford University
Graduate School of Business
and NBER
tonetti@stanford.edu

1 Introduction

Americans have access to a range of financial products to insure against major life-cycle health, longevity, and mortality risks. Three primary products are: (i) *life insurance*, which insures against premature mortality by increasing bequests to surviving beneficiaries; (ii) *annuities*, which insure against outliving one's wealth by providing a guaranteed income stream for the duration of life; and (iii) *long-term care insurance*, which covers costs related to assistance with the activities of daily living, e.g., nursing homes.

Despite the severity and frequency of these risks, only 3% of individuals aged 60–65 own a private annuity, 60% own life insurance, and 12% own long-term care insurance (LTCI). Why individuals choose not to insure against these risks, instead partially self-insuring through savings, remains a central question in household finance and macroeconomics. Low ownership may reflect either limited intrinsic demand for resources in these states or crowding out from public insurance; alternatively, individuals may desire more insurance, but find the products available in the private market undesirable.

We depart from the literature, which typically models individual idealized products, by evaluating joint portfolio choice over savings and insurance with realistic features. We make two key contributions. First, using a novel, purpose-designed survey, we document that individuals view insurance itself as inherently risky, harboring substantial concerns that policies will not fully pay out. Second, we embed these empirical nonpayment beliefs—alongside other insurance market frictions—into an incomplete-markets life-cycle model to quantify how these frictions determine portfolio choices, saving, and consumer welfare.

Individuals expect to receive 76%–87% of promised payouts conditional on a qualifying event, and perceive annual full-default probabilities ranging from 1.3% to 1.8% across products. Life insurance is consistently viewed as the safest and LTCI as the riskiest. Validating these measures, we find that elicited beliefs covary with actual insurance ownership. Model counterfactuals reveal that nonpayment risk and high pricing loads severely depress insurance demand, encourage precautionary saving, and reduce consumer welfare. Baseline access to existing, imperfect insurance products generates a modest welfare gain of 0.6% of lifetime consumption and drives a 16% decline in median wealth. Making insurance risk-free substantially increases insurance participation rates—from 55% to 73% for life insurance, 13% to 43% for LTCI, and 2% to 8% for annuities. Consequently, the welfare gain from these idealized insurance products rises to 2.3% and median wealth falls an additional 11% below baseline. Finally, we decompose the sources of these effects and we document how public insurance programs like Medicaid and Social Security shape this environment.

Risky Insurance Survey Details. Conditional on a qualifying event, individuals expect to receive only 87.1% of the promised payout for life insurance, 81.5% for annuities, and 76.2% for LTC insurance. Furthermore, their elicited payout distributions reveal substantial uncertainty around these averages. We additionally measure the annual full-default probability—defined as the likelihood that an insurer permanently ceases all future payments. Consumers believe these full-default probabilities are 1.3% for life insurance, 1.4% for annuities, and 1.8% for LTCI. Auxiliary survey measures reveal that this perceived risk includes not only concerns over insurer solvency, but also broader fears about contract disputes, delays, and procedural complexity in claims processes.

In support of these measures, we show that subjective nonpayment risk predicts actual insurance ownership. While standard demographic controls and behavioral measures struggle to predict individual-level insurance demand, our elicited risk measures robustly predict non-participation. The implied magnitudes are economically large: counterfactual probit estimates indicate that annuity and LTCI ownership rates would approximately double if individuals perceived these products as risk-free.

Structural Model Details: Insurance Ownership, Savings, and Welfare. We embed multiple insurance products into an otherwise standard heterogeneous-agent incomplete-markets life-cycle model of consumption and saving. Individuals face health and mortality risk and borrowing constraints. The model incorporates realistic insurance market features, including subjective nonpayment risk, pricing loads, illiquidity (wedges between purchase and resale prices), and quantity restrictions such as age cutoffs and short-sale limits. The model also features means-tested public long-term care support (Medicaid) and public annuities (Social Security).

We calibrate the model to match empirical moments of the joint distribution of age, sex, income, wealth, health, mortality, and average insurance ownership rates, alongside real-world insurance product characteristics. In addition to the targeted moments, the calibrated model replicates untargeted life-cycle patterns of insurance ownership, with persistently low participation in annuities and long-term care insurance, and high, but age-declining, participation in life insurance. It also matches the correlation of ownership across products.

With the quantified model, we show that product-market frictions play a central role in shaping insurance participation. Eliminating either perceived nonpayment risk or pricing loads yields large increases in ownership. For example, making insurance risk-free raises LTCI participation from 13% to 43%, while removing both frictions pushes it above 50%, with life insurance exhibiting similarly large responses. As individuals substitute from self-insurance toward state-contingent market insurance, liquid wealth falls and savings are drawn down more aggressively in old age. These results on ownership do not depend on whether nonpayment beliefs are rational, since only the beliefs themselves drive behavior.

We then turn to consumer welfare analysis, where the distinction between beliefs and the true probability of realized outcomes becomes important. Because there are no reliable comprehensive measures of true nonpayment risk we consider two cases, always holding subjective beliefs fixed at their measured values: (1) beliefs are rational, meaning nonpayment risk is realized as perceived, and (2) contracts always honor their commitments, so beliefs are too pessimistic. Under rational beliefs, the baseline welfare gain from the existence of private insurance is modest at 0.6% of lifetime consumption on average. However, eliminating nonpayment risk raises the gain to 2.3%, and removing both nonpayment risk and price loads increases the gain to 4.1%. In case (2) with non-rational pessimistic beliefs, the baseline welfare gain is mechanically higher at 1.7%—because individuals always receive insurance payments they are due—but remains below the 2.3% achievable if they correctly perceived insurance as risk-free and accordingly purchased more insurance.

Decomposing these welfare gains reveals that LTCI accounts for over half of the total welfare gain from all products, despite substantially greater ownership of life insurance. Furthermore, the risk of annual partial nonpayment suppresses participation and welfare significantly more than the risk of full default, especially for LTCI and life insurance, which together deliver over 90% of total welfare gains.

Finally, we demonstrate the central role of public insurance. Means-tested public long-term care (Medicaid) substantially crowds out private LTCI; while reducing the generosity of this safety net spurs private insurance demand, it imposes severe, regressive welfare losses. Furthermore, replacing Social Security's public annuity with an upfront lump-sum wealth transfer does not spur private annuitization. Instead, individuals use the liquidity to self-insure. Welfare improves for poorer credit-constrained individuals as they can better smooth consumption, while welfare declines for richer individuals, who lose access to a valued annuity stream that has only an imperfect private-market substitute.

In summary, the phenomenon of risky insurance plays a crucial role in determining insurance ownership, life-cycle savings, and the welfare effects of insurance. Eliminating insurance nonpayment risk—whether actual or incorrectly perceived—yields substantial welfare gains by allowing individuals to insure major life-cycle health and mortality risks. Consequently, the welfare cost of not owning risky insurance is much lower than the cost of not owning idealized products that complete markets.

In the remainder of the paper, Section 2 introduces the data, Section 3 analyzes the elicited distributions of insurance nonpayment risk, Section 4 presents the structural model and its calibration, Section 5 presents model results, and Section 6 concludes.

1.1 Relation to the Literature

Life-Cycle Savings. We contribute to the literature on life-cycle saving, which emphasizes that mortality risk, medical expense risk, long-term care risk, and bequest motives are primary drivers of wealth accumulation. It is well established that the desire to leave a bequest affects life-cycle saving behavior. See, e.g., (Hurd, 1989; De Nardi, 2004; Kopczuk and Lupton, 2007; Love, Palumbo and Smith, 2009; Lockwood, 2012). It is also well-understood that late-in-life medical cost and long-term-care health risks affect life-cycle saving. See, e.g., (Hubbard, Skinner and Zeldes, 1994; Palumbo, 1999; Anderson, French and Lam, 2004; Kopecky and Koreshkova, 2014; Braun, Kopecky and Koreshkova, 2017; De Nardi, French and Jones, 2010; Lockwood, 2012, 2018; Ameriks, Briggs, Caplin, Shapiro and Tonetti, 2020).

Models in this literature often treat late-in-life health shocks as privately uninsurable, restricting the asset space to risk-free bonds and forcing elderly households to self-insure through precautionary saving or rely on government care. We build on this framework by explicitly modeling the private insurance contracts designed to cover these risks. Our contribution is to incorporate insurance portfolio choice and to measure and model nonpayment risk.[1] Consequently, we show that access to private insurance reduces precautionary saving, even though frictions in these markets generate a reliance on self-insurance that would be absent with complete markets.

Insurance Demand and Nonpayment Risk. A large literature examines the demand for annuities (e.g., Yaari (1965); Brown (2001)), life insurance (e.g., Bernheim (1989); Hong and Ròs-Rull (2012)), and LTC insurance (e.g., Brown and Finkelstein (2007, 2008, 2011)). We contribute to the study of the demand for life-cycle insurance products by measuring and modeling imperfect insurance, with an emphasis on the risk of insurance nonpayment.

As discussed by Mayers and Smith (1983) and emphasized in Gomes, Haliassos and Ramadorai (2021), the demand for insurance products should be studied jointly due to the covariance in their payouts.[2] A very useful benchmark for our results is provided by Koijen, Van Nieuwerburgh and Yogo (2016), who evaluate the joint demand for annuities, life insurance, and LTC insurance in a complete-markets framework. We complement this approach by analyzing insurance portfolio choice in a quantitative incomplete-markets model.

In addition to incorporating insurance portfolio choice, we connect the life-cycle savings literature to research on how counterparty risk affects insurance demand.[3] We contribute by

[1] For tractability and computational feasibility, we abtract from other features recently explored in frontier models, such as a more general modeling of family and labor supply (De Nardi, French, Jones and McGee, 2025; Borella, De Nardi, Torres Chain and Yang, 2026).

[2] See also Inkmann, Lopes and Michaelides (2011) and Hubener, Maurer and Rogalla (2013).

[3] Doherty and Schlesinger (1990) show that nonperformance risk alters standard predictions of insurance

measuring subjective insurance nonpayment risk and documenting its connection to actual insurance ownership.[4] Additionally, we embed this risky insurance into a structural model to show how nonpayment risk limits insurance ownership and drives life-cycle savings. Finally, we examine the consumer welfare gains from the existing, imperfect insurance available in the market, and contrast these gains to those achievable from risk-free insurance.

Impact of Social Insurance on Savings and Welfare. We also contribute to the study of how Social Security and Medicaid affect life-cycle savings and welfare. Previous research highlights how means-tested programs like Medicaid provide implicit insurance against medical expenses, rationalizing the low asset holdings of the poorest households (De Nardi, French and Jones, 2016). Additionally, Braun, Kopecky and Koreshkova (2019) explore how public insurance crowds out private insurance. (Conesa and Krueger, 1999) study the effects of Social Security on savings and welfare over the life cycle. Recent work further explores how social insurance expansions alter household borrowing and financial resilience (Bornstein and Indarte, 2023) and affect inequality (Catherine, Lu and Paron, 2025). We add to this literature by examining how the public provision of insurance via Social Security and Medicaid interact with the demand for a portfolio of imperfect private insurance products.

Consumer Beliefs and Economic Behavior. Finally, our work connects to literature that measures consumer beliefs to understand economic behavior (Manski, 2004). In household finance, Giglio, Maggiori, Stroebel and Utkus (2021) link survey-elicited beliefs to administrative data, showing that beliefs about expected returns and rare disaster probabilities predict equity portfolio allocations. Additionally, Bailey, Cao, Kuchler and Stroebel (2018) show how subjective beliefs affect housing tenure decisions. Survey-elicited subjective expectations have also been shown to affect household consumption and labor supply decisions (Wiswall and Zafar, 2021; Conlon, Pilossoph, Wiswall and Zafar, 2018; Weber, D'Acunto, Gorodnichenko and Coibion, 2022; Coibion, Georgarakos, Gorodnichenko, Kenny and Weber, 2024). See the handbook chapter by Koşar and O'Dea (2023) for an overview. Our method of measurement builds on Luttmer and Samwick (2018), who study uncertainty surrounding Social Security benefits. We contribute to the literature by measuring beliefs about insurance nonpayment risk and demonstrating that these subjective expectations covary with actual ownership of life-cycle insurance products. By embedding these elicited beliefs into a structural model, we quantify the welfare and saving implications of these consumer expectations.

demand, and several studies note that small default probabilities can limit annuity market participation (Lopes and Michaelides, 2007; Pashchenko, 2013; Jang, Keun Koo and Park, 2013).

[4] Beshears, Choi, Laibson, Madrian and Zeldes (2014) document that survey respondents rank nonpayment risk as an important consideration limiting annuity purchases.

Figure 1: **Age Profile of Insurance Product Ownership Rates**

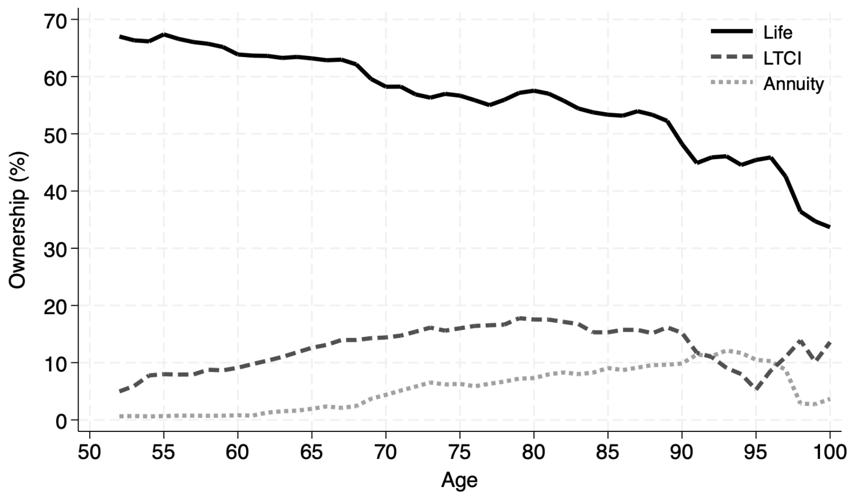

Note: Ownership rates for life insurance (solid black), long-term care insurance (dashed dark gray), and annuities (small-dash light gray) for US individuals. Source: Health and Retirement Study 2014.

2 Data

We use two sources of data in our analysis. First, we use the Understanding America Study (UAS), a nationally representative internet panel of approximately 6,000 adults hosted by the University of Southern California. In May-June 2018 we fielded a survey about life-cycle insurance products to 1,104 UAS respondents who were at least 45 years old (detailed in Section 2.3). Second, we use the Health and Retirement Study (HRS), a nationally representative panel of more than 37,000 Americans older than age 50 with rich measures of health, income, assets, and insurance ownership from 1991 to the present.

We use the two sources in complementary ways. The UAS allows us to measure beliefs about insurance nonpayment risk and to correlate beliefs with insurance ownership and a rich set of individual characteristics. We also use the UAS to set initial values for the state variables of age, sex, income, health, wealth, and insurance ownership, in the structural model. In that sense, our structural model provides an analysis of the lives of our UAS survey respondents. The HRS is commonly used in the literature due to its high quality measurements and large sample size with a long panel dimension. Thus, we use the HRS to obtain empirical target moments to calibrate the structural model. After appropriate reweighting, both the UAS and HRS samples are representative of older Americans, so we utilize both sources of data without significant external-validity concerns.

2.1 Life-Cycle Patterns of Insurance Ownership

Figure 1 plots the cross-section of insurance participation rates by age from the 2014 HRS. Two patterns stand out. First, ownership of annuities and LTCI is rare: at age 65, roughly 12% of individuals own private LTCI and fewer than 5% own a private annuity, consistent with the classic annuity puzzle (Yaari (1965), Modigliani (1986)) and LTCI puzzle (Ameriks, Briggs, Caplin, Shapiro and Tonetti (2018); Braun et al. (2019)). Second, life insurance participation is much higher, around 60% at age 65, but declines steadily with age, consistent with earlier term policies purchased to insure children's welfare maturing over time (Hong and Ròos-Rull (2012)). Despite substantial health and mortality risks, Figure 1 indicates limited demand by older Americans for the annuity or LTC insurance products available in the market.

2.2 The UAS Dataset

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

Our survey was fielded in May 2018. Among invited panelists, the response rate was 82%[5]. In total, 1,104 started and 1,085 completed the survey; average completion time was 17 minutes. Item non-response was minimal: 1,045 answered all life insurance items, 1,054 all annuity items, and 1,039 all LTCI items. To maximize sample size, we do not restrict analyses to a common respondent set across products.

Table 1 compares summary statistics of the UAS and corresponding HRS samples. Some differences versus the HRS are mechanical (English-only sampling in the UAS delivers fewer Hispanic and more White respondents). Others are more substantive (on average UAS respondents are slightly younger and less retired, with less wealth and higher incomes). Importantly, conditional moments and key insurance ownership patterns are similar across both datasets.

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;

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

Table 1: **HRS and UAS Summary Statistics**

	HRS	**UAS**
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
Retired	0.46	0.34
Income ($000s)	69.7	100.6
Wealth ($000s)	582	461
Insurance ownership		
Annuity	0.04	0.12
Life insurance	0.61	0.58
LTCI	0.12	0.10
N	18,747	1,085

Note: We report sample means for each variable. Column (1) uses the HRS sample of individuals aged 50 and over; Column (2) uses respondents aged 45 and over from our UAS insurance survey. Health categories are self-reported. Insurance ownership variables are indicators for owning a private annuity, life insurance policy, or long-term care insurance policy. Observations are weighted using survey weights.

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. These variables appear as controls in selected specifications alongside standard demographics.

2.3 Insurance Survey Description

This section describes UAS survey 118 that we fielded to measure beliefs regarding insurance nonpayment risk.[6] We adapt the elicitation methodology developed by Luttmer and Samwick (2018) for Social Security benefits, to evaluate perceived risk in private insurance markets.

[6]The survey and collected data are available at `https://uasdata.usc.edu/survey/UAS+118`.

Section 2.3.1 presents our two core measures. First, we measure the subjective distribution of partial nonpayment, defined as the fraction of the contracted payout received during a qualifying event. Second we measure the subjective full-default probability, defined as the likelihood that an insurer permanently ceases all future payments. Section 2.3.2 then presents certainty-equivalent valuations and auxiliary measures exploring the sources of these nonpayment beliefs.

2.3.1 Core Measures of Subjective Nonpayment Risk

Survey Design Overview. Our survey features three modules to elicit perceived nonpayment risk in annuity, life insurance, and LTC insurance products. We randomize module order to mitigate and test for potential anchoring effects. Each module branches based on the respondent's current ownership status. If the respondent owns the product, we collect their specific policy details (e.g., benefit size) and ask them to evaluate the nonpayment risk of their actual policy. For non-owners, we provide a standardized product description, randomly assign a hypothetical benefit size, and ask them to evaluate the best policy they believe they could currently purchase in the market. With a specific reference policy established, we then elicit subjective nonpayment risk. To illustrate the survey instrument, this section details the annuity module; the life and LTC insurance modules are implemented analogously.

Full-Default Probabilities. 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 $[20,000] 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 in a given year the annuity becomes worthless due to no fault of your own. 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 input their perceived probability using a 0-to-100 slider, accompanied by dynamic text that explains the meaning of their current selection.

Annual Partial Nonpayment Distributions. We then elicit the perceived distribution of payouts conditional on a valid qualifying event. Following the bins-and-balls methodology developed by Delavande and Rohwedder (2008), respondents use a visual interface to allocate balls that represent probability mass across bins that define potential outcomes. On the initial screen, respondents distribute 20 balls across three broad bins: receiving no payment, a partial payment, or the full promised amount:

> 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 then reallocate any probability mass assigned to the partial-payment event across five sub-bins (ranging from 1–20% to 81–99%). These bins represent the realized payout as a percentage of the contracted amount.

> 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 procedure elicits the full subjective payout distribution, allowing us to compute the mean and standard deviation of perceived annual payouts as a fraction of the contracted amount.[7]

[7] Moment calculations assume all probability mass is located at the midpoint of each bin.

2.3.2 Auxiliary Measures of Subjective Nonpayment Risk

Next, we elicit certainty equivalents by asking respondents whether they would trade their risky insurance policy for a risk-free contract offering a guaranteed fraction of the promised payout. Respondents who perceive positive nonpayment risk make this choice exactly once for a randomly assigned guaranteed payout of 25%, 50%, 75%, 80%, or 90% of the promised amount (see Appendix A for the survey instrument). This single-shot elicitation means we do not observe certainty equivalents at the individual level, precluding the recovery of the full risk premia distribution. Instead, the aggregate acceptance rates across guaranteed payout thresholds map the population cumulative distribution function (CDF) of certainty equivalents. From this CDF, we compute population-level summary statistics, including the mean certainty equivalent for each product and their implied risk premia.

To complement the core risk metrics, a supplemental module elicits further dimensions of nonpayment risk. Each respondent completes this supplement for a single insurance product after completing the core modules.[8] One component of the supplement measures perceived nonpayment risk conditional on an aggregate economic downturn. We prompt respondents to evaluate a scenario in which the stock market drops by a randomly assigned magnitude (10%, 20%, or 30%). Using an elicitation structure identical to the core survey, respondents then report both the full-default probability and the subjective partial payout distribution conditional on this macroeconomic shock. We also explore the underlying drivers of perceived nonpayment risk. For example, we elicit respondents' expectations regarding the administrative burden and procedural effort required to claim their contracted benefits.

The survey concludes with a final module asking respondents to reflect on their answers. These questions, alongside the randomized survey design, provide metrics to evaluate overall response quality. To test for internal consistency, we first ask respondents to qualitatively rank life, annuity, and LTC insurance from least to most risky. We define a set of responses as internally inconsistent if none of the core quantitative risk metrics is largest for the insurance type qualitatively ranked as the most risky. Next, respondents identify the factors that influenced their beliefs personal experiences, family/friends' experiences, government intervention, policy complexity, the aggregate economic conditions, trust in salespeople, other/free response). Finally, we elicit respondents' confidence in their answers and the amount of prior thought they had given to both insurance purchases and nonpayment risk.

[8]To prevent framing effects from contaminating the primary data, the supplement corresponds to the final product in the respondent's randomized sequence and is administered only after all core modules are complete. See Appendix A for details.

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

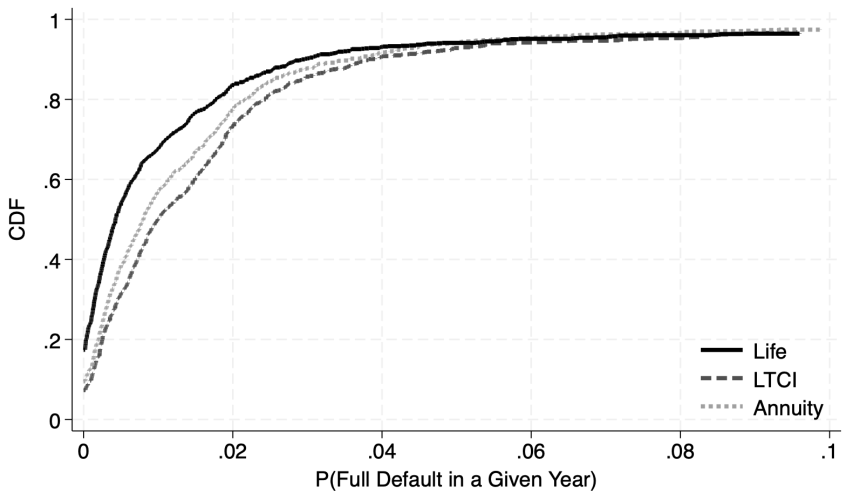

Note: Cross-sectional CDFs of subjective annual full-default probabilities, defined as the probability that an insurer permanently ceases all future payments.

3 Insurance Survey Results

This section analyzes our survey data in three steps. First, we document the elicited measures of perceived nonpayment risk. Second, we examine how these beliefs relate to individual characteristics and predict actual insurance ownership. Finally, we evaluate the credibility and internal consistency of the survey responses. All analyses apply UAS population weights to yield nationally representative estimates for Americans over age 45.

3.1 Subjective Nonpayment Risk Measures

We document two primary metrics of perceived nonpayment risk: the annual probability of full-default and the annual expected partial payout conditional on a qualifying event. A qualifying event is defined as survival for annuities, death for life insurance, and needing help with 2 or more activities of daily living for LTCI.

Figure 2 presents the cross-sectional CDFs of subjective full-default probabilities. The data exhibit pervasive nonpayment concerns, with the vast majority of respondents assigning a positive probability to a full default. Only 19% of respondents report a zero full-default probability for life insurance, alongside 11% for annuities and 10% for LTCI. Responses reveal a clear ordinal ranking, with life insurance viewed as the least likely to default and LTCI as the most risky. At the median, the expected annual full-default probability is 0.3 percent for life insurance, 0.5 percent for annuities, and 1.0 percent for LTCI.

Figure 3: **Subjective Annual Partial Payout Distribution**

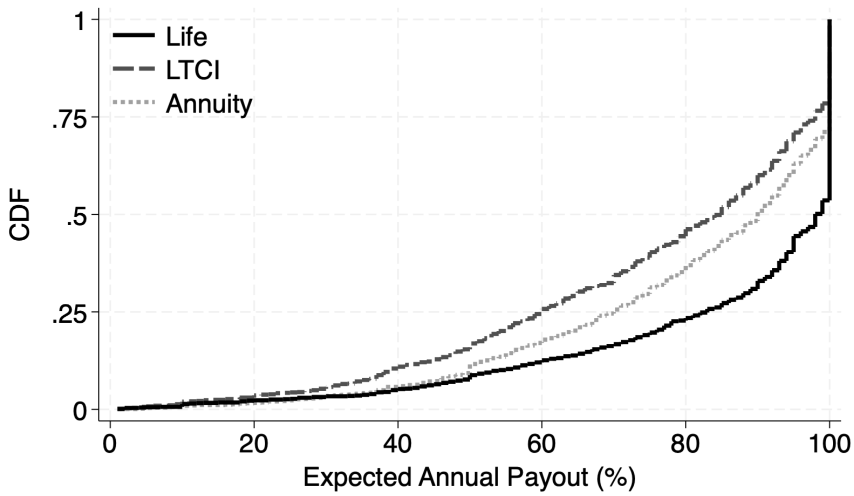

Note: Cross-sectional CDFs of the mean of the subjective payout distribution conditional on a qualifying event. Payouts are expressed as a fraction of the contracted payout.

Figure 3 presents the cross-sectional CDFs for the mean of subjective annual payouts conditional on a qualifying event. Expected payouts are expressed as a fraction of the contracted payout. For life insurance, although around half of all respondents expect full payment, respondents on average expect to receive only 87% of the promised amount. Payout expectations are even lower for annuities and LTCI. For annuities, the mean expected payout is 82% and the median is 90%, while for LTCI the mean is 76% and the median is 85%.[9]

Together, these figures document our first key result: across all three products individuals perceive life-cycle insurance products as risky. Regardless of the objective nonpayment probabilities, it is these subjective beliefs that ultimately govern their behavior. Furthermore, we find life insurance is consistently viewed as the safest, LTCI as the riskiest, and annuities as intermediate in risk.

3.2 Auxiliary Measures of Subjective Nonpayment Risk

To evaluate how consumers value this uncertainty, we next examine the population distributions of certainty equivalents in Figure 4. Each point on these CDFs indicates the fraction of individuals who prefer a risk-free contract with that payout share to their current risky insurance policy. These measures are relatively similar across insurance types. At the median, respondents would trade a risky annuity, life insurance, or LTCI policy for a risk-free policy

[9] Appendix Figure B.1 reveals that respondents perceive substantial uncertainty around these payout expectations. On average, subjective payouts have a standard deviation of 17% of the promised amount for life insurance, 19% for annuities, and 20% for LTC insurance.

Figure 4: **Certainty Equivalent Distributions**

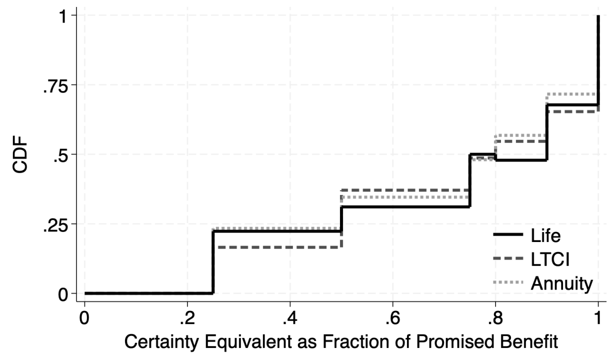

Note: CDFs of respondents' certainty equivalents, defined as the guaranteed fraction of the promised payout that renders a respondent indifferent to the risky policy. Each point on the CDF indicates the fraction of respondents who prefer that certain payout share to the risky policy.

Table 2: **Insurance Product Risk Premia**

	Expected Value	**Certainty Equivalent**	**Risk Premia**
Life	87.16	81.43	5.72
Annuity	81.51	73.79	7.72
LTCI	76.17	72.90	3.27

Note: This table reports mean expected payout shares, certainty equivalent shares, and the implied risk premia for each insurance product. Payout values are expressed as percentage shares of the promised payouts. The risk premium equals the expected payout share minus the certainty equivalent share.

that guarantees 75–80% of the promised payout. This indicates that typical respondents are risk averse and willing to forgo a substantial share of the promised payout to eliminate nonpayment risk.

Table 2 reports the average expected payout and certainty equivalent for each insurance type. The difference between these measures is the implied risk premium. We find that risk premia are economically significant across all policies; the annuity risk premium, for instance, is nearly 8%, comparable to the well-documented equity risk premium. While measured subjective nonpayment risks are substantial, the associated risk premia are comparable to those found in other asset classes. These results confirm that individuals view insurance as risky and require significant compensation to bear that risk.

To examine the drivers of risk perceptions, Figure 5 presents subjective annual payout distributions conditional on different aggregate financial conditions, proxied by stock market returns. Relative to the baseline, we find that respondents expect annuity and life insurance payments to decrease in economic downturns: the payout CDFs shift leftward conditional on

Figure 5: **Expected Payout Conditional on Aggregate Financial Conditions**

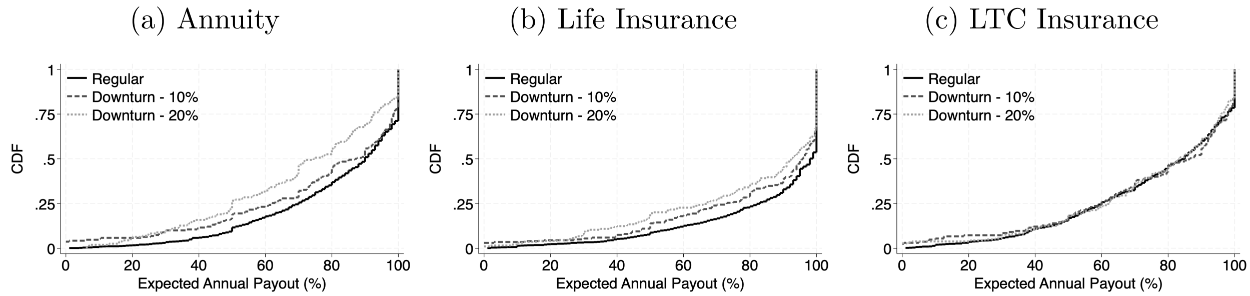

Note: Cross-sectional CDFs of expected mean payout shares conditional on aggregate financial conditions, proxied by annual stock market returns. Beliefs are reported for "regular" times and for 10% and 20% annual stock market declines.

a 10% market decline, with a more pronounced shift at 20%. In contrast, the subjective LTC insurance payout distribution is invariant to market conditions. These divergent patterns suggest that perceived sources of risk vary by insurance type. Speculatively, LTCI risk may be associated with claim denials stemming from subjective state verification by insurance companies, whereas annuity and life insurance risks may be associated with the solvency of insurers. At a minimum, this heterogeneity indicates a degree of sophistication in survey responses and nonpayment beliefs.

3.3 Subjective Nonpayment Risk and Insurance Ownership

The survey documents substantial perceived nonpayment risk. We next examine whether these subjective beliefs predict actual insurance ownership by regressing ownership indicators on our elicited risk metrics, controlling for demographic and other individual characteristics. Specifically, for each product $k \in \{\text{Ann, LI, LTC}\}$, we run the following probit regression across all respondents i:

$$\Pr\left(\text{Own}_i^k\right) = \Phi\Big(\alpha^k \;+\; \beta_1^k\, \mathbb{E}_i\left[\omega^k\right] \;+\; \beta_2^k\, \delta_i^k \;+\; \beta_3^k\, \sigma_i(\omega^k) \;+\; X_i'\gamma\Big) \tag{1}$$

where $\text{Own}_{i,k}$ is an indicator for ownership of product k; $\Phi(\cdot)$ is the standard normal CDF; $\mathbb{E}_i\left[\omega^k\right]$ is the expected annual payout percentage share; $\sigma_i(\omega^k)$ is the standard deviation of the subjective payout distribution, δ_i^k is the annual full-default probability; and X_i is a vector of demographic and behavioral controls available in the UAS.[10] Table 3 reports the estimated

[10] Demographic controls include age, wealth, income, education, marital status, race, health, household size, and retirement status, which have an intuitive relationship to ownership. Other controls include trust, cognitive score, financial literacy score, numeracy score, experienced fraud, risk aversion, propensity to plan, and early-life stock market returns. Conditional on the core risk metrics, these other covariates are generally insignificant. See Appendix B for an expanded table that reports estimated coefficients for all controls.

Table 3: **Regression of Insurance Ownership on Risky Insurance Beliefs**

	Own Annuity	**Own Life**	**Own LTCI**
Payout Expected Value ($\mathbb{E}_i[\omega^k]$)	-0.0002	0.0040**	0.0007
	(0.770)	(0.010)	(0.131)
Payout Standard Deviation ($\sigma_i(\omega^k)$)	-0.0025***	-0.0006	-0.0008
	(0.001)	(0.733)	(0.246)
Full-Default Probability (δ_i^k)	-0.0021***	-0.0012	-0.0022***
	(0.000)	(0.167)	(0.000)
N	1054	1045	1039
R^2	0.267	0.214	0.175
Demographic Controls	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes

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

Note: Average marginal effects predicting product ownership via probit regression (1). Regressors include the subjective expected payout percentage share $\left(\mathbb{E}_i[\omega^k]\right)$, the payout distribution standard deviation $\left(\sigma_i(\omega^k)\right)$, and the annual full-default probability $\left(\delta_i^k\right)$. All specifications include demographic and supplementary individual controls (control coefficient estimates reported in Appendix Table B.1). Robust standard errors are reported in parentheses.

average marginal effects, with robust standard errors.

The regression results suggest that perceived nonpayment risk is a deterrent to insurance ownership. For annuities, both a higher full-default probability and a larger standard deviation of perceived payouts significantly reduce the probability of ownership ($p < 0.001$). For life insurance, a higher expected payout predicts ownership ($p < 0.01$), and for LTC insurance, the full-default probability predicts lower ownership ($p < 0.001$). Although the remaining risk measures generally exhibit the expected signs, they are not statistically significant.[11] Since beliefs were measured after ownership decisions were made, these probit estimates cannot definitively establish causality. However, if we assume that subjective nonpayment beliefs are persistent, a causal interpretation remains plausible. Regardless of the causal direction, this robust correlation provides strong evidence that our elicited beliefs capture economically meaningful variation rather than random survey noise.

Table 4 quantifies the economic magnitude of these effects by comparing baseline predicted ownership to a counterfactual where insurance is perceived as risk-free ($\mathbb{E}_i\left[\omega^k\right] = 100$; $\sigma_i(\omega^k) = \delta_i^k = 0$). Eliminating perceived risk doubles participation rates for both annuities (from 12% to 24%) and LTC insurance (from 10% to 23%), and increases life insurance

[11]Appendix B documents that these results are robust to excluding the extensive set of other controls, which offer little predictive power beyond our nonpayment risk measures.

Table 4: **Probit Counterfactual Ownership Rates with Safe Insurance**

	P(Own)	**P(Own | Safe Insurance)**
Annuity	.12	.24
Life	.57	.66
LTCI	.10	.23

Note: The first column reports predicted ownership probabilities using estimates from regression equation (1). The second column reports counterfactual results if all respondents perceive insurance products as risk-free, i.e., using the predicted values from regression equation (1) setting $\mathbb{E}[\omega^k] = 100$, $\sigma(\omega^k) = 0$, and $\delta^k = 0$.

participation from 57% to 66%. In sum, the survey data and probit estimates suggest that subjective nonpayment risk is substantial and meaningfully depresses insurance ownership.

However, these statistical counterfactuals implicitly hold all other individual characteristics constant. In practice, eliminating nonpayment risk would trigger broader behavioral shifts; for example, increased reliance on safe insurance might diminish the need for precautionary savings. This endogenous reduction in wealth would, in turn, feedback into insurance demand. To capture the joint determination of savings and insurance portfolios, and to evaluate the welfare implications of risky insurance, Section 5 revisits these counterfactuals using a structural model.

3.4 What Determines Subjective Nonpayment Risk?

We next explore the underlying determinants of these subjective nonpayment risk beliefs. Table 5 reports estimates from regressing each of the three risk measures on demographic and supplementary individual characteristics. Because these risk metrics are naturally bounded, we estimate Tobit models with appropriate censoring limits.

The estimates reveal several intuitive drivers of risk perception. First, higher cognitive and financial literacy scores predict significantly lower perceived risk. This suggests that negative beliefs may partially reflect a lack of product understanding or concerns about successfully interacting with insurance companies, either when purchasing the product or navigating the claims process. Second, prior experience with fraud is associated with higher perceived risk, particularly for full-default probabilities. Furthermore, the propensity to plan and cohort stock returns from ages 18–25 significantly predict nonpayment beliefs. In contrast, self-assessed trust and risk aversion lack predictive power, although it is unclear to what degree noise in these proxy measures drives insignificance.

Table 5: **Regression of Risk Measures on Demographic and Behavioral Controls**

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

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

Note: Tobit regression estimates predicting the expected annual payout share, $\mathbb{E}[\omega^k]$, standard deviation of the subjective payout distribution, $\sigma(\omega^k)$, and annual full-default probability δ^k. All specifications include demographic and supplementary individual characteristics. Robust standard errors are reported in parentheses.

3.5 Survey Response Quality and Credibility

The validity of our findings relies on the credibility of the elicited subjective beliefs. First, the survey design allows us to rule out some framing effects: the randomized order of modules and hypothetical payout amounts yields no significant differences in the core estimates presented in Table 3. Second, we verify internal consistency using a qualitative ranking of insurance riskiness as described in Section 3.2. We classify 84% of respondents as internally consistent—meaning their qualitative ranking of insurance products align with their quantitative elicitations—and show in Appendix B that our results are robust to restricting analysis to this subsample.

In addition, several features of survey responses provide evidence of answer quality. For example, there is a consistent ordinal ranking of the perceived riskiness of insurance

Figure 6: **Factors Considered when Responding to the Survey**

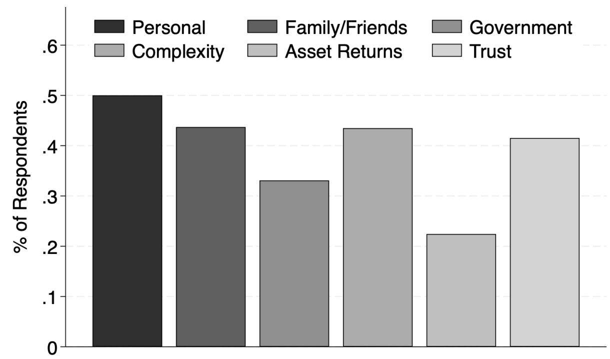

Note: Fraction of respondents citing each factor as a substantive consideration in their nonpayment beliefs.

products across individuals. More importantly, the ability of these subjective measures to predict actual insurance ownership establishes their external validity, directly linking elicited beliefs to observed behavior. Similarly, the elicited conditional expectations reveal that respondents associate aggregate financial risk with annuity and life insurance payout risk, but not with LTCI risk. This distinction accurately reflects the greater interest rate risk insurance companies face in annuity and life insurance markets.

Finally, self-evaluation metrics indicate a high degree of confidence in individuals' responses. As reported in Figure 6, respondents considered a number of reasonable factors when formulating their answers. Moreover, only 17% of the sample reported lacking not being confident in their answers, and only 9% indicated having given no prior thought to the core concerns of the survey. Overall, these patterns suggest that the survey captures considered, economically meaningful, subjective beliefs.

4 Model of Life-cycle Saving and Insurance Demand

This section develops a structural model of life-cycle insurance demand to quantify how nonpayment risk and institutional frictions shape saving, portfolio choice, and consumer welfare. We embed multiple insurance products into an otherwise standard incomplete-markets life-cycle model of consumption and saving with health and mortality risk and borrowing constraints (cf. De Nardi et al. (2010), Lockwood (2018), Ameriks et al. (2020)). The model incorporates realistic insurance market features, including subjective nonpayment risk, pricing loads, illiquidity (wedges between purchase and resale prices), and quantity restrictions such as age cutoffs and short-sale limits. We discipline nonpayment beliefs using our UAS insurance survey data, while remaining product characteristics are calibrated to statistics of market contracts taken from the literature. Standard life-cycle dynamic variables—including income, wealth, and health and mortality risks—are calibrated to match HRS data. We first outline the model environment and calibration, before presenting the quantitative results in Section 5.

Model Overview. We develop a model of individual choice with exogenous prices that are calibrated to match real-world prices. Time is discrete and each period is a year. Individuals are heterogeneous in exogenous state variables: age (t), health status (s_t), sex (f), and an income-age profile (y, with element y_t). Health and mortality are the main nonfinancial risks in the model. An individual's health status can be good, poor, in need of long-term care, or dead. It evolves according to a Markov transition matrix that is a function of age, sex, health status, and a measure of permanent income. Health directly affects flow utility via a nonhomothetic health-state-dependent utility function, out-of-pocket medical expenditures, and insurance prices. Individuals enter each period holding risk-free bonds (b_t) and a portfolio of insurance contracts (x_t^k) for $k \in \{\mathrm{LI}, \mathrm{LTC}, \mathrm{Ann}\}$. The annual premia for life insurance (p_i^{LI}) and LTCI (p_i^{LTC}) are locked in at the time of purchase and depend on contemporaneous individual states. At the beginning of the period health and insurance nonpayment shocks are realized. Given their updated states and shocks, individuals choose how much to consume, save, and invest across a portfolio of insurance products or to use means-tested government services. Individuals do not equalize marginal utility across dates and states due to insurance nonpayment risk, price loads, quantity limits, and borrowing constraints.

4.1 Model Details

We now introduce each aspect of the model in detail before assembling the components to write the complete decision problem.

Demographics and Health. The model is populated by heterogeneous individuals indexed by age $t \in \{45, 46, \ldots, 99\}$ and sex $f \in \{0, 1\}$, with $f = 1$ denoting female. Each period, an individual's health status is given by $s_t \in \{0, 1, 2, 3\}$, corresponding respectively to normal health, poor health, requiring long-term care, and death. An individual requires LTC if they need and receive help with activities of daily living (ADLs).[12] Health evolves according to a Markov transition matrix $\pi(s_{t+1} \mid s_t, t, f, y)$ which is a function of current health, age, sex, and permanent income. We allow for transition probabilities to differ for those in the bottom three permanent-income quintiles versus the top two, capturing the well-documented

[12] An individual requires LTC if they receive help in at least two of five ADL categories (bathing, eating, dressing, getting in/out of bed, and walking), as measured by the adl5h variable in the HRS.

wealth gradient in health and longevity. Additionally, there are out-of-pocket medical costs, $\xi \sim \Xi(\xi \mid s_t, t, f)$. This lognormally distributed cost is a negative shock to liquid wealth, with a mean that increases in age and worse health states.

Preferences. Households have time-separable, health-state-dependent nonhomothetic preferences for consumption, c_t, and a warm-glow bequest motive. Flow utility u_s is

$$u_s(c_t) = \theta_{s,t}^{-\gamma} \frac{(c_t + \kappa_s)^{1-\gamma}}{1-\gamma}$$

Following Ameriks et al. (2020), this specification allows for nonhomothetic utility in the LTC state ($\kappa_2 \neq 0$), analogous to the standard nonhomothetic treatment of bequests developed by De Nardi (2004). The parameter κ_s controls the degree to which expenditure in health state s is a luxury or a necessity, while $\theta_{s,t}$ scales marginal utility. Because this framework nests the more-common model ($\kappa_2 = 0$), the data can inform the importance of nonhomothetic LTC utility, placing long-term care risk on equal footing with the bequest motive to explain measured life-cycle saving patterns.

We depart from the prior literature by allowing the bequest scale $\theta_{3,t}$ to vary with age. This parsimoniously captures age-dependent bequest motives, reflecting, for instance, the changing economic needs of children over the life cycle. For ages $t \geq 80$, the bequest parameter is constant at $\theta_{3,t} = \bar{\theta}_3$. For ages $45 \leq t < 80$ it evolves according to:

$$\theta_{3,t}^{-1} = \theta_{3,45}^{-1} + \left(\bar{\theta}_3^{-1} - \theta_{3,45}^{-1}\right) \left(\frac{\exp(\rho x) - 1}{\exp(\rho) - 1}\right), \tag{2}$$

where $x = \frac{t-45}{80-45}$. The parameter $\theta_{3,45}^{-1}$ anchors the initial strength of the bequest motive, while the parameter ρ controls the curvature of its life-cycle profile.[13]

Income. Surviving individuals receive an exogenous, deterministic annual income, y_t. This income comprises post-tax labor earnings, Social Security benefits, and defined-benefit pension payouts. Using HRS data, we estimate five life-cycle income profiles via quintile regressions and assign individuals to a profile based on their measured income. We abstract from idiosyncratic income shocks primarily to reduce the computational burden of the dynamic portfolio choice problem. Because our analysis focuses on older cohorts whose income depends more on annuitized retirement payouts like Social Security and pensions, muting income risk provides a tractable approximation of their late-in-life economic environment.

[13]As ρ tends to 0 the bequest motive becomes linear in age. Our calibration finds $\theta_{3,45}^{-1} > \bar{\theta}_3^{-1}$ and $\rho > 0$, producing a bequest motive that is decreasing and strictly convex in age, thereby concentrating the decline at older ages.

Assets. The core innovation of our structural model is the joint saving and insurance portfolio decision. Individuals allocate wealth across four assets. The first is a standard risk-free bond, $b_t \geq 0$, which pays a constant return r and is subject to a no-borrowing constraint. For tractability, given the complexity of the insurance asset space, we use this single liquid-wealth variable to encompass all assets other than insurance.

Second, individuals can purchase three types of insurance products, indexed by $k \in \{\text{LI}, \text{LTC}, \text{Ann}\}$: immediate life annuities, whole life insurance, and long-term care insurance. Each year, individuals choose their desired start-of-next-period holdings for each insurance product, denoted by $\tilde{x}^k_{t+1}$.

The insurance products are distinguished by the health states s in which they pay out. Let χ^k be a 4×1 vector of dividends for asset k, where row s specifies the unit payout in state s: $\chi^{\text{Ann}} = [1, 1, 1, 0]'$; $\chi^{\text{LI}} = [0, 0, 0, 1]'$; $\chi^{\text{LTC}} = [0, 0, 1, 0]'$. Thus, annuities pay while alive, life insurance pays upon death, and LTCI pays when in need of LTC. These payout vectors, combined with the stochastic health dynamics, determine the covariance of these insurance asset payouts.

Individuals take exogenous prices as given. Insurance prices are a function of age, health, and sex and are calibrated to match real-world prices in Section 4.3.

Annuities are purchased via a lump-sum payment at unit price $p^{\text{Ann}}_{t,s,f}$. In contrast, life insurance and LTCI require ongoing annual premia. These premia are locked in at initial purchase based on an individual's contemporaneous age, health, and sex at prices $p^{\text{LI}}_{t,s,f}$ and $p^{\text{LTC}}_{t,s,f}$, making the contracted premia endogenous state variables (p^{LI}_{it} and p^{LTC}_{it}). To maintain computational tractability, individuals may hold at most one life insurance and one LTCI policy at a time. To adjust these holdings, they must let their existing contract lapse and purchase a new policy at current prices.

The economic environment features several institutional frictions. Short-selling is prohibited ($\tilde{x}^k_{t+1} \geq 0$), and new policies cannot be purchased after an exogenous maximum age T, reflecting insurer concerns over adverse selection. Additionally, annuities can be sold for cash value, but at a discounted price $p^{\text{Ann}}_{t,s,f}\left(1 - \lambda^{\text{Ann}}_{-}\right)$.

Insurance Nonpayment Risk. Consistent with our survey measurements, insurance contracts are subject to two forms of nonpayment risk. First, an insurer may fully default on the contract with probability δ^k. Absent default, next period's insurance holdings equal the chosen quantities; under a full default, holdings fall to zero without any sale proceeds:

$$x^k_{t+1} = \begin{cases} \tilde{x}^k_{t+1} & \text{with probability } 1 - \delta^k \\ 0 & \text{with probability } \delta^k. \end{cases} \tag{3}$$

For life and LTC insurance, a full default destroys the locked-in premium advantage. To restore the lost coverage, the individual must repurchase the policy at prevailing spot prices, which are higher due to their advanced age. For annuities, future payouts are lost with zero recovery of the initial lump-sum purchase amount.

Second, insurers may partially default on promised payouts. In any given period, the realized payout is only a fraction $\omega_t^k \in [0,1]$ of the promised payout, where $\omega_t^k \sim \Omega^k(\cdot)$.

To preserve computational tractability, we abstract from belief heterogeneity. All individuals share a common subjective probability distribution over these nonpayment shocks, matched to the central tendencies of our survey data.

Government Insurance The government provides a state-dependent consumption floor, $\bar{c}_s$. Approximating government means-testing rules, participating individuals are constrained to enter the subsequent period with zero assets, $b_{t+1} = x_{t+1}^k = 0$. For individuals in normal or poor health, the consumption floor is $\bar{c}_0 = \bar{c}_1 = \bar{c}$. For those who need LTC ($s = 2$), the government provides $\bar{c}_2$. This parameter captures the value individuals place on government-provided LTC services, such as Medicaid-funded nursing homes.

Intertemporal Budget Constraint Collecting the components described above, the intertemporal budget constraint tracking the evolution of bonds is:

$$\frac{1}{(1+r)} b_{t+1} = b_t + y_t + v_{s_t} \sum_k \omega_t^k x_t^k \chi^k - c_t - \xi_t - p_{i,t+1}^{\text{LTC}} \tilde{x}_{t+1}^{\text{LTC}} - p_{i,t+1}^{\text{LI}} \tilde{x}_{t+1}^{\text{LI}} \\ - (\tilde{x}_{t+1}^{\text{Ann}} - x_t^{\text{Ann}}) p_{t,s,f}^{\text{Ann}} \left(1 - \lambda_-^{\text{Ann}} \mathbb{I}_{(\tilde{x}_{t+1}^{\text{Ann}} < x_t^{\text{Ann}})}\right). \tag{4}$$

The right-hand side of the constraint balances available resources against period expenditures. Resources consist of initial bond holdings b_t, income y_t, and total realized insurance payouts. We define v_{s_t} to be a 1×4 indicator vector that identifies the individual's current health state, so $v_{s_t} \sum_k \omega_t^k x_t^k \chi^k$ captures total insurance payouts net of any partial nonpayment shocks. These resources fund consumption c_t, out-of-pocket medical expenses ξ_t, and the annual premia life and LTC insurance policies. The prices for the annual premia are locked-in at the purchase price as long as the individual is not choosing to change the quantity owned ($\tilde{x}_{t+1} = x_t$). The final term captures annuity transactions: adjustments are made at the prevailing spot price $p_{t,s,f}^{\text{Ann}}$, subject to the illiquidity discount λ_-^{Ann} when selling.

Full Decision Problem. The value function for an individual with states $S_t := \{s_t, y, f, b_t, x_t^{\text{Ann}}, x_t^{\text{LI}}, x_t^{\text{LTC}}, p_{i,t}^{\text{LTC}}, p_{i,t}^{\text{LI}}\}$ is $V_t(S_t) = \max\{V_t^{\text{Gov}}(S_t)\ ,\ V_t^{\text{Mkt}}(S_t)\}$. The value of using means-tested government services is:

$$V_t^{\text{Gov}}(S_t) = u_s(\bar{c}_s) + \beta\mathbb{E}\left[V_{t+1}(s_{t+1}, y, f, 0, 0, 0, 0, 0, 0)\right].$$

Alternatively, individuals remaining in the market choose consumption, savings, and desired insurance portfolio (c_t, b_{t+1}, $\tilde{x}_{t+1}^{\text{Ann}}$, $\tilde{x}_{t+1}^{\text{LI}}$, and $\tilde{x}_{t+1}^{\text{LTC}}$) to maximize their value:

$$\begin{aligned}
V_t^{\text{Mkt}}(S_t) = \max\, & u_s(c_t) + \beta\mathbb{E}\left[V_{t+1}(S_{t+1})\right] \\
\text{s.t.} \quad & \frac{1}{(1+r)} b_{t+1} = b_t + y_t + v_{s_t} \sum_k \omega_t^k x_t^k \chi^k - c_t - \xi_t - p_{i,t+1}^{\text{LTC}} \tilde{x}_{t+1}^{\text{LTC}} - p_{i,t+1}^{\text{LI}} \tilde{x}_{t+1}^{\text{LI}} \\
& \qquad - (\tilde{x}_{t+1}^{\text{Ann}} - x_t^{\text{Ann}}) p_{t,s,f}^{\text{Ann}} \left(1 - \lambda_-^{\text{Ann}} \mathbb{I}_{(\tilde{x}_{t+1}^{\text{Ann}} < x_t^{\text{Ann}})}\right).
\end{aligned}$$

Full default forces realized next-period insurance holdings to zero:

$$x_{t+1}^k = \begin{cases} \tilde{x}_{t+1}^k & \text{with} \quad \text{probability } 1 - \delta^k \\ 0 & \text{with} \quad \text{probability } \delta^k. \end{cases}$$

The locked-in annual premia for life and LTCI reset to the spot price if the contract changes quantity, either voluntarily or through default:

$$p_{i,t+1}^k = \begin{cases} p_{i,t}^k & \text{if} \quad \tilde{x}_{t+1}^k = x_t^k \\ p_{t,s,f}^k & \text{if} \quad \tilde{x}_{t+1}^k \neq x_t^k \end{cases} \quad \text{for} \quad k \in \{\text{LI}, \text{LTC}\}.$$

Portfolio choices are bound by no-short-sale and maximum-issue-age constraints, $\tilde{x}_{t+1}^k \geq 0$ and $\tilde{x}_{t+1}^k \leq x_t^k$ for $t > T$. Health, medical costs, and nonpayment risk are distributed according to $\pi(s_{t+1} \mid s_t, t, f, y)$, $\Xi(\xi \mid s_t, t, f)$, and $\Omega^k(\omega^k)$. Prices $p_{t,s,f}^k$ and r are taken as given.

Numerical Solution. Because our incomplete-markets framework precludes the analytical solutions available in complete-market settings, we solve the individual choice problem numerically via backward induction using a grid search over a discretized state space. To simulate the model, we draw initial conditions from the empirical joint distribution of the UAS survey. Specifically, we assign each simulated individual their observed age, sex, income, wealth, and initial insurance holdings. We then simulate the model forward, allowing these individuals to dynamically optimize their consumption, savings, and insurance portfolios over the remainder of their life. Model statistics are computed using UAS population weights. Appendix D.1 details the solution and simulation procedures.

4.2 Discussion of Risks, Costs, and Incomplete Markets

Individuals in our model may have a precautionary saving motive arising from three distinct life-cycle sources: longevity risk, long-term care risk, and the desire to leave a bequest. Under complete markets (e.g., Koijen et al. (2016)), individuals could perfectly hedge these concerns using annuities, life insurance, and LTCI, eliminating the need for precautionary bond holdings. In contrast, our incomplete-market framework introduces several risks, costs, and constraints that may depress insurance demand and push individuals toward self-insurance via personal savings. Before presenting the quantitative results, it is useful to outline how these specific frictions shape insurance portfolio choices.

The most critical friction is insurance nonpayment risk (δ^k and ω^k), which is the main source of market incompleteness. Because insurance is designed to pay out in states with high marginal utility, a surprise default in those exact states is particularly devastating. Anticipating these uninsurable defaults could deter insurance purchases, even when perfect insurance for that state would be highly valued.

The second major friction is insurance pricing. In practice, prices are not actuarially fair; rather, they embed substantial loads above discounted state probabilities. Furthermore, annuity liquidations face steep surrender discounts. Absent nonpayment risk, the insurance products span all health states. Nonetheless, price loads would prevent the equalization of marginal utility across states.

Third, uninsurable out-of-pocket medical expenses generate a precautionary saving motive (see, e.g., De Nardi et al. (2010)). Individuals may be reluctant to lock their wealth into insurance contracts that cannot be easily liquidated to cover immediate medical cost shocks.

Finally, we impose a maximum age at which agents are able to purchase new policies. In practice, insurers routinely reject older applicants, possibly due to adverse selection concerns. For example, Hendren (2013) reports 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 around age 80–90.

By capturing these risks and institutional frictions, our model provides a laboratory to quantify how market imperfections shape life-cycle insurance portfolio choice, explore how insurance ownership affects life-cycle saving patterns, and evaluate the welfare provided by real-world, rather than idealized, insurance products.

4.3 Model Calibration

We first calibrate some parameters of the model externally and then use the simulated method of moments (SMM) to calibrate the remaining parameters.

Standard Parameters External Calibration. We use HRS data to measure income, wealth, health, and medical costs. We estimate 5 income-age profiles using sex-specific quintile regressions. This income measure includes post-tax labor earnings, Social Security benefits, and defined-benefit pension payouts. Bonds, b_t, encompass all non-insurance assets, including financial and housing wealth.[14] We set the government consumption floor for individuals not requiring LTC to \$5,000 ($\bar{c} = 5$), consistent with standard values in the literature (e.g., De Nardi et al. (2016)). See Appendix C for details on constructing y_t, b_t, $\pi(s_{t+1} \mid s_t, t, f, y)$, and $\Xi(\xi \mid s_t, t, f)$.

Insurance Parameters External Calibration. We next calibrate the parameters that determine insurance product characteristics, reported in Table 6. To avoid sensitivity to outliers, we truncate the UAS insurance survey responses at the bottom and top 2%. We then use the mean of these trimmed responses to calibrate the full-default probabilities δ^k. For the partial nonpayment distributions, $\Omega^k(.)$, we discretize the outcome space into four equally likely events and set the associated partial-payment values to match the mean and standard deviation of the empirical subjective annual payout distributions. This yields $\omega_t^{\text{Ann}} \in \{.50, .78, .94, 1\}$, $\omega_t^{\text{LI}} \in \{.58, .92, 1, 1\}$, and $\omega_t^{\text{LTC}} \in \{.48, .71, .88, 1\}$.[15]

To calibrate prices in our model of demand, our objective is to confront individuals with the prices they would face in the real-world market, remaining agnostic to the supply-side determinants of such prices. In practice, prices are not actuarially fair; rather, they embed substantial loads above discounted state probabilities. Brown and Finkelstein (2011) and Hong and Rìos-Rull (2012) report real-world market prices in terms of loads above actuarially fair prices. The actuarially-fair price is computed to equal the discounted expected lifetime payouts of the contract, conditional on the purchasing individual's age, health, sex, and the corresponding health transition probabilities. Therefore, we define a product-specific load, λ^k, such that multiplying the actuarially-fair price by $1 + \lambda^k$ yields the empirical market price. We compute these actuarially-fair prices using a risk-free interest rate of $r_I = 2.5\%$, in line with Brown and Finkelstein (2007, 2011), and Koijen and Yogo (2015).

[14] Financial wealth is the sum of liquid assets (checking, saving, money market, and certificate of deposit accounts) and investment assets (brokerage, mutual fund, education-saving, individual retirement, and defined-contribution retirement accounts).

[15] Outcomes are set to the 12.5, 37.5, 62.5, and 87.5 percentiles of the empirical subjective annual payout distributions. This procedure closely matches the empirical means and standard deviations.

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

	Annuities	**Life**	**LTCI**
Full Default (δ^k)	.014	.013	.018
Annual Payout Fraction ($E[\omega]$)	.815	.871	.761
Annual Payout Uncertainty ($\sigma(\omega)$)	.194	.174	.195
Insurance Loads (λ^k)	.2	.25	.32
Annuity Surrender Discount (λ_{-}^{Ann})	.35	–	–
Max Purchase age (T)	80	80	80

Note: Payout beliefs calibrated to match surveyed subjective nonpayment belief distributions. Insurance loads from Brown and Finkelstein (2011) and Hong and Rìos-Rull (2012). Maximum purchase age from Hendren (2013). Annuity surrender discount from industry report Christian (2025).

To match observed market prices, we apply the loads reported by Brown and Finkelstein (2011) and Hong and Rìos-Rull (2012): 20 percent for annuities ($\lambda^{\text{Ann}} = 0.20$), 25 percent for life insurance ($\lambda^{\text{LI}} = 0.25$), and 32 percent for long-term care insurance ($\lambda^{\text{LTC}} = 0.32$). For annuity liquidations, industry reports (Christian (2025)) suggest surrender discounts of 20–40 percent, so we set the sale penalty to 35 percent ($\lambda_{-}^{\text{Ann}} = 0.35$). We interpret λ^k as reduced-form wedges that summarize supply-side forces, including administrative and distribution costs, capital charges, adverse selection, expected lapsation revenues, and markups. This approach generates an empirically accurate choice environment without explicitly modeling insurer operations.

Remaining Parameters SMM Calibration. Having fixed the exogenous state processes and insurance product characteristics, we jointly calibrate the remaining preference and return parameters using the simulated method of moments (SMM). These parameters include γ, θ_2, κ_2, $\bar{c}_2$, $\theta_{3,45}$, ρ, $\bar{\theta}_3$, κ_3, β, and r. We allow the rate of return on individual savings to differ from the risk-free rate used to price insurance, as the single saving asset in the model proxies for the broader, higher-yielding, portfolio of assets individuals hold in the data.

The SMM procedure targets three sets of empirical moments: (i) the 25th, 50th, and 75th percentile of the wealth distribution by 5-year age bins for ages 45–84; (ii) the average ownership rate of annuities (4%), life insurance (61%), and long-term care insurance (12%); and (iii) mean responses to the nine strategic survey question (SSQ) from Ameriks et al. (2020). These SSQs place individuals in hypothetical scenarios and ask them to make choices that are designed to identify preferences. We target the extensive margin of insurance ownership because survey data reliably captures binary participation status; self-reported face values and benefits levels are likely measured with greater error. The SMM minimizes the distance between these empirical moments and their simulated counterparts (see Appendix

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

LTC luxury	$\kappa_2 = -37.44$	**Time preference**	$\beta = 0.9575$
LTC marginal	$\theta_2 = 0.67$	**Risk aversion**	$\gamma = 1.5$
Bequest luxury	$\kappa_3 = 7.83$	**Govt LTC floor**	$\bar{c}_2 = 77.43$
Old bequest marginal	$\bar{\theta}_3 = 1.09$	**Govt floor**	$\bar{c} = 5$
Young bequest marginal	$\theta_{3,45} = 0.034$	**Household bond return**	$r = 0.05$
Bequest life-cycle shape	$\rho = 1.5$	**Insurer discount rate**	$r_I = 0.025$

Note: Baseline calibration for parameters governing preferences and returns. See text for calibration procedure.

D for calibration details).

To manage the computational burden of solving the life-cycle portfolio problem over a high-dimensional state space and parameter space, we employ a multi-stage optimization routine. We initialize the utility parameters (γ, θ_2, κ_2, $\bar{c}_2$, $\bar{\theta}_3$, and κ_3) at the estimated values from Ameriks et al. (2020) and conduct a grid search over the remaining parameters (β, $\theta_{3,45}$, ρ, r). Then, we hold β, $\theta_{3,45}$, ρ, and r at their best-fit values, and search over the initially-fixed preference parameters to improve the model fit. Because this second stage yielded substantial improvements only via the risk aversion parameter (γ), our final optimization stage jointly calibrates $(\gamma, \beta, \theta_{3,45}, \rho, r)$, holding the remaining parameters fixed at their initial Ameriks et al. (2020) values.

Calibrated Parameter Values. Table 7 reports the calibrated parameter values. A value of $\gamma = 1.5$ implies a moderate degree of risk aversion. The bequest motive for younger individuals, controlled by $\theta_{3,45}^{-\gamma}$, scales the marginal utility of a bequest at age 45 by approximately 180 relative to death at age 80. This generates a powerful incentive for younger individuals to hold life insurance to protect against the rare event of early death. The curvature parameter $\rho = 1.5$ sustains this strong bequest motive through middle age, causing it to substantially taper only after age 70.

Time preference also plays a critical role in shaping insurance demand. A relatively high discount factor β reduces incentives to insure against earlier-life health shocks, while placing more weight on longevity risk, LTC risk that realizes at older ages, and consumption later in life. Combined with a calibrated interest rate r such that $\beta(1 + r) > 1$, patience encourages wealth accumulation earlier in the life cycle. Appendix D.4 provides further details on the sensitivity of model outcomes to these parameter values.

5 Model Results: Insurance Ownership and Welfare

This section presents results from our calibrated structural model. We begin in Section 5.1 by establishing that the calibrated model accurately matches both targeted moments and untargeted patterns of insurance ownership. The remainder of this section details four core findings. First, Section 5.2 quantifies how subjective nonpayment risk and price loads significantly depress insurance ownership and increase wealth due to precautionary saving. Second, Section 5.3 evaluates the consumer welfare gain provided by real-world insurance products relative to a no-insurance benchmark. At baseline, the market yields only modest welfare gains equivalent to 0.6% of lifetime consumption. However, an idealized market offering insurance with zero nonpayment risk and zero loads generates substantial consumer welfare gains of 4.1%, reconciling our baseline findings with complete-market benchmarks. Third, Section 5.4 decomposes these gains, revealing two key insights: LTCI generates the largest welfare improvement despite lower ownership rates than life insurance, and partial nonpayment risk is four times more detrimental to welfare than full-default risk. Finally, Sections 5.5 and 5.6 quantify the impact of government-provided LTCI and annuities, via Medicaid and Social Security, on private insurance ownership and welfare.

5.1 Model Fit

Figure 7 shows that the model replicates both the average level and the life-cycle profile of wealth across the wealth distribution. Although the calibration only targeted average insurance participation rates, the model successfully generates several untargeted features of insurance demand. Figure 8 demonstrates our fit to the life-cycle pattern of insurance ownership. Furthermore, Table 8 shows that the model captures portfolio composition in terms of both the number of insurance products owned and the correlations in ownership across products. Nearly 90% of individuals either hold no private insurance at all or only a single life insurance product. When they do hold a second product, it is most often LTC insurance, consistent with the positive correlation between Life and LTCI ownership.

Figure 7: **Wealth Moments: Model and Data**

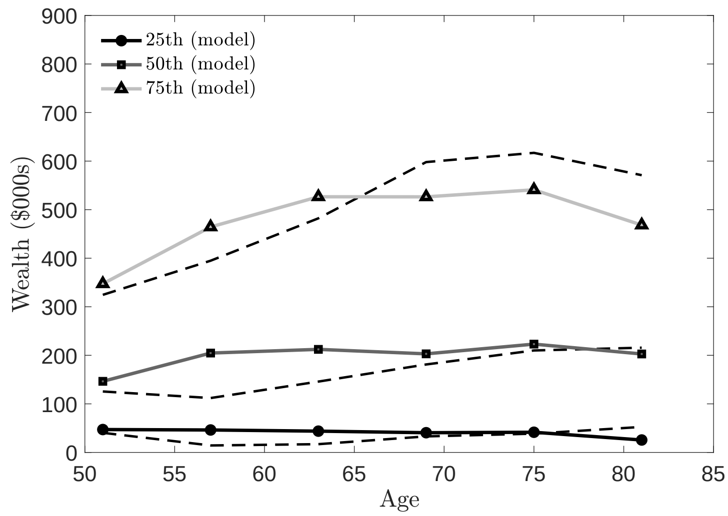

Note: Wealth levels by age for the 25th, 50th, and 75th percentiles of the distribution. Data moments (dashed lines) are from the HRS; model moments (solid lines) are from the baseline calibration.

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

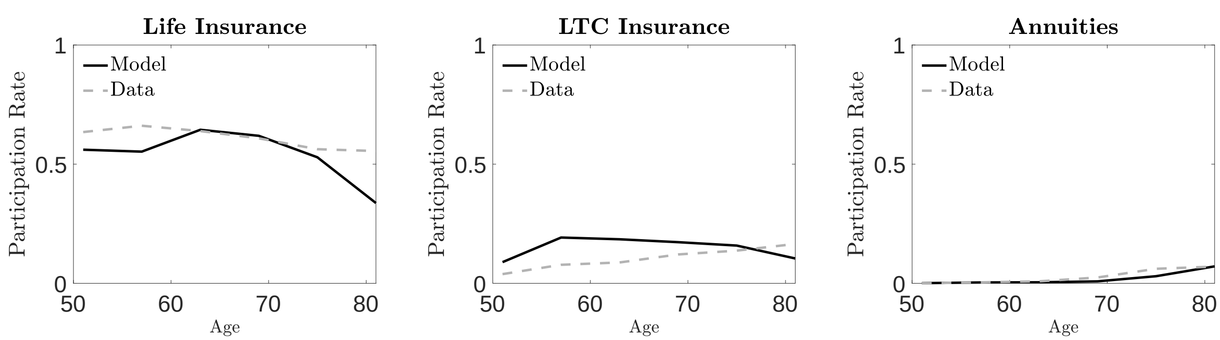

Note: Insurance participation rates by age. Dashed lines are empirical rates from the HRS; solid lines are from the baseline calibrated model.

Table 8: **Nontargeted Insurance Ownership Moments**

A. *Number of Products Owned*	**0**	**1**	**2**	**3**
Model	41%	46%	12%	1%
Data	35%	55%	9%	1%

B. *Ownership Correlations*	**(LI, LTCI)**	**(LI, Ann)**	**(LTCI, Ann)**
Model	0.15	-0.03	0.10
Data	0.08	-0.02	0.06

Note: Panel A reports the fraction of individuals owning zero to three insurance products. Panel B reports the pairwise correlation of joint participation between two insurance products. Data are from the HRS; model results are from the baseline simulation.

5.2 Determinants of Insurance Market Participation and Savings

We now quantify how product market frictions shape individuals' saving and insurance portfolio allocations. Table 9 reports outcomes under four scenarios: the *Baseline* calibrated specification; *Safe Insurance*, which sets all subjective nonpayment probabilities to zero; *Zero Loads*, which eliminates the the load markups ($\lambda^k = 0$); and a combined scenario *Zero Loads & Safe Insurance*.[16]

Panel A of Table 9 reports aggregate ownership rates for annuities, life insurance, and LTCI. Our baseline calibrated model matches the empirical aggregate participation rates well. The next three columns quantify how real-world product characteristics depress insurance demand. For life insurance, eliminating nonpayment risk significantly increases participation, with ownership rising to 73%. The 25% load also severely suppresses demand: setting loads to zero raises ownership to 71%. Together, zero loads and zero nonpayment risk increase life insurance participation to 84%. LTCI is even more affected by these frictions: eliminating nonpayment risk more than triples ownership (from 13% to 43%), removing loads doubles it (to 29%), and removing both expands participation to 55%. In contrast, the effects on annuities are muted: ownership rises to 8% without nonpayment risk, 3% with zero loads, and 10% when both frictions are removed.

Panel B of Table 9 quantifies the impact on savings by reporting the percentage change in wealth at the 25th, 50th, and 75th percentiles relative to an economy with no private insurance. In the baseline scenario, access to insurance reduces wealth most dramatically for poorer

[16] We evaluate the zero-load counterfactual without taking a stance on whether these empirical loads generate pure profits or cover necessary operating costs. Rather, this adjustment allows us to contrast our findings with complete market benchmarks, where actuarially-fair prices fundamentally alter the consumer choice problem.

Table 9: **Insurance Ownership and Savings: Impact of Risky Insurance and Loads**

	Data	**Baseline**	**Safe Insurance**	**Zero Loads**	**Zero Loads & Safe Insurance**
A. *Ownership Rate*					
LI	61%	55%	73%	71%	84%
LTCI	12%	13%	43%	29%	55%
Ann	4%	2%	8%	3%	10%
B. Δ *Personal Savings*					
25th percentile	–	−36%	−48%	−42%	−55%
50th percentile	–	−16%	−27%	−21%	−33%
75th percentile	–	−4%	−16%	−10%	−22%

Note: Panel A reports aggregate ownership rates for three insurance products under alternative scenarios. Panel B reports percentage changes in wealth across the distribution (see Figure 7), relative to values from a no-insurance benchmark, averaged across age with population weights. Baseline is the calibrated model; Safe Insurance sets default probabilities to zero ($\omega^k = 1$, $\delta^k = 0$); Zero Loads sets loads to zero; Zero Loads & Safe Insurance combines both adjustments.

individuals (a 36% decline at the 25th percentile), who in the no-insurance counterfactual rely heavily on precautionary savings to hedge against health and mortality risks. In contrast, wealthier individuals, who are more likely to have adequate savings across states, show smaller reductions (16% at the median and 4% at the 75th percentile).

The decline in wealth is more pronounced as frictions are removed. The largest drop occurs when both loads and nonpayment risk are eliminated. Two mechanisms drive this result. First, as cheaper and safer insurance products become more attractive investments, individuals substitute out of risk-free bonds and into formal insurance. Second, improved access to state-contingent protection weakens the precautionary motive to hold liquid wealth. Consequently, when these product market frictions are removed, individuals purchase more insurance and draw down their wealth much more aggressively in old age.

Overall, both subjective nonpayment risk and high loads significantly depress insurance ownership rates and generate substantial increases in precautionary savings.

5.3 Welfare Gains from Life-Cycle Insurance Products

In this section we quantify the welfare gains provided by insurance products with varying characteristics. We report the consumption-equivalent welfare gain relative to a no-insurance benchmark, aggregated across individuals with UAS survey weights.[17]

[17] Consumption-equivalent welfare is the percent increase in consumption across all states of the world in the no-insurance benchmark that makes an individual indifferent between the no-insurance benchmark and the considered scenario (Appendix D.3).

Table 10: **Welfare Gains From Insurance**

	Baseline	**Safe Insurance**	**Zero Loads**	**Zero Loads & Safe Insurance**
A. *Rational Beliefs*				
Δ **Welfare**	0.6%	2.3%	1.8%	4.1%
B. *Non-Rational Beliefs*				
Δ **Welfare**	1.7%	2.3%	3.4%	4.1%

Note: Consumption-equivalent welfare gains relative to a no-insurance benchmark. Panel A assumes individuals are correct in their assessment of nonpayment risk; Panel B assumes zero objective nonpayment risk, and thus, incorrect beliefs. Subjective beliefs are the same across panels, fixed at baseline calibrated values. Baseline is the calibrated model; Safe Insurance sets perceived and objective default probabilities to zero; Zero Loads sets loads to zero; Zero Loads & Safe Insurance combines both adjustments.

Up to now, we have not needed to take a stand on whether individuals' subjective beliefs are rational. An individual's portfolio allocation rule is a function of their state variables and subjective beliefs, regardless of whether those beliefs reflect the true objective probability distribution. Furthermore, because purchases typically occur long before qualifying events, payout realizations do not meaningfully alter the state variables at the time of purchase. Consequently, insurance ownership is virtually identical in a world of rational expectations and in a world in which payouts are always made. For welfare, however, the distinction matters; actually receiving payouts is obviously better than not receiving them. Because we lack empirical measures on realized nonpayment rates, we perform our welfare analysis in two bounding cases. First, we simulate the model under rational expectations, where nonpayment risk realizes according to subjective nonpayment probabilities. Second, we simulate the model with zero actual nonpayment risk. In both cases nonpayment risk beliefs are identical, set to their baseline measured values.

Panel A of Table 10 shows that, when expectations are rational, baseline welfare gains from the existence of these insurance products are modest at 0.6% of lifetime consumption. This reflects the fact that nearly half of individuals optimally choose not to own any formal insurance (Table 8), relying instead on self-insurance and public insurance. Furthermore, those who own insurance are left with products that sometimes do not pay when promised, exactly when individuals have revealed they highly value payouts by their purchase of insurance in the first place. Removing frictions produces substantial welfare improvements. Eliminating nonpayment risk almost quadruples the baseline gains to 2.3%; removing loads nearly triples them to 1.8%; and eliminating both simultaneously yields a 4.1% welfare gain.

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 nonpayment risk because individuals can effectively hedge substantial health risks that they prefer to hedge. Instead, insurance products with high loads and nontrivial risk of nonpayment discourage individuals from purchasing these products at a substantial loss of consumer welfare.

Panel B of Table 10 reports welfare gains when actual nonpayment risk is zero. Because portfolio allocation rules depend solely on subjective beliefs, portfolio choices are identical to the rational expectations case, but ex-post welfare is evaluated using simulations in which promised payouts are always made. This generates a mechanical welfare increase: baseline gains rise from 0.6% under rational expectations to 1.7% simply because individuals experience zero defaults. More interestingly, comparing this 1.7% gain with the 2.3% gain from eliminating both subjective and objective nonpayment risk highlights the role of subjective beliefs. In both scenarios, the objective product characteristics are identical and payouts are guaranteed; however, when individuals accurately perceive that there is no nonpayment risk, they purchase significantly more insurance (Table 9 Panel A). This behavioral shift—utilizing the available products at scale—generates an additional 0.6% welfare gain. Here, the welfare gained simply by correcting beliefs is as large as the rational-expectations baseline value of the products existing in the first place.

Across both perspectives, the core conclusion remains clear: removing product market frictions—whether eliminating loads or removing actual or perceived nonpayment risks—yields substantial welfare gains by allowing individuals to efficiently insure major life-cycle health and mortality risks. Equivalently, the welfare costs of not owning real-world insurance products is much lower than the costs of not owning idealized products that complete markets.

5.4 The Sources of Welfare Gains

To determine how each type of insurance contributes to consumer welfare improvements, Table 11 reports the gains from introducing each product individually, relative to the zero-insurance benchmark.[18] In the baseline economy, the majority of the 0.6% total welfare gain stems from LTCI, even though life insurance is held by roughly four times as many individuals. Isolating life insurance yields only a 0.1% welfare gain, whereas isolating LTCI generates a 0.5% gain. Annuities alone have a negligible contribution. This hierarchy persists when nonpayment risk is eliminated, since LTCI suffers from the highest perceived default risk. However, when price loads are removed, life insurance accounts for a larger share of the total welfare gain. Since life insurance is already perceived as the safest product and has the highest baseline ownership, its welfare contribution share is particularly sensitive to price reductions.

[18] For simplicity of exposition, the remainder of the paper focuses on the rational expectations case. The main takeaways are robust to the case with zero objective nonpayment risk (holding subjective beliefs constant at measured levels); see Appendix E for these results.

Table 11: **Welfare Gains from Individual Insurance Products**

	Baseline	**No Annual Nonpayment Risk**	**Safe Insurance**	**Zero Loads**	**Zero Loads & Safe Insurance**
LI	0.1%	0.6%	0.7%	1.0%	1.7%
LTCI	0.5%	1.1%	1.3%	0.7%	1.5%
Ann	0.0%	0.1%	0.2%	0.2%	0.7%
All	0.6%	1.9%	2.3%	1.8%	4.1%

Note: Consumption-equivalent welfare gains relative to a no-insurance benchmark (rational beliefs case). Each row isolates the introduction of only the specified product. The final row introduces all three products and reproduces Panel A from Table 10. Baseline is the calibrated model; No Annual Nonpayment Risk sets the annual partial nonpayment risk to zero ($\omega^k = 1$), but leaves full-default probabilities δ^k at baseline levels; Safe Insurance sets all default probabilities to zero; Zero Loads sets loads to zero; Zero Loads & Safe Insurance sets all nonpayment probabilities and loads to zero.

Nonpayment risk takes two forms: annual partial nonpayment risk and full-default risk. Table 11 decomposes the welfare cost associated with each type of risk. When all insurance products are available, the annual partial nonpayment risk accounts for 76% of the total gains from removing all nonpayment risk; eliminating only the annual partial nonpayment risk delivers a 1.3% gain over the baseline compared to the 1.7% gain from removing all nonpayment risk. Turning to individual products, life and LTC insurance are paid for gradually by annual premia, so if a policy defaults the individual simply ceases future premia payments. As a result, the primary cost of a full default is the loss of the preferential premium price that was locked-in at initial purchase. In contrast, a full default on an annuity is substantially more costly because the individual loses the entire upfront lump-sum purchase payment. As a result, full-default risk accounts for roughly half of the total nonpayment risk costs for annuities.

In summary, LTCI generates 0.5% of the total 0.6% baseline welfare gain provided by all insurance products. Additionally, removing annual partial nonpayment risk has a four times larger welfare impact than eliminating full-default risk. This is especially true for life and LTC insurance, which are the products that deliver the vast majority of welfare gains across all scenarios.

5.5 The Effect of Medicaid on Ownership and Welfare

Our baseline calibration reflects the empirically low participation rate in the private LTCI market. Because high loads and nonpayment risks depress the value of private insurance, individuals rely on two substitute channels to insure against LTC shocks. The first is Medicaid, which provides a guaranteed consumption floor, $\bar{c}_2$. Due to its fixed care quality and strict

Table 12: **The Effects of Government LTC Provision (Medicaid)**

	Baseline $\bar{c}_2 = 77.43$	$\bar{c}_2 = 60$	$\bar{c}_2 = 40$
A. *Ownership Rate*			
LI	55%	54%	51%
LTCI	13%	25%	44%
Ann	2%	2%	2%
B. *LTCI Ownership Rate*			
0–25th Wealth Percentile	10%	12%	42%
25–50th Wealth Percentile	27%	42%	72%
50–75th Wealth Percentile	15%	29%	56%
75–100th Wealth Percentile	12%	35%	46%
C. *Welfare Effects*			
Aggregate	-	-1.3%	-4.0%
0–25th Wealth Percentile	-	-1.4%	-5.7%
25–50th Wealth Percentile	-	-1.7%	-6.1%
50–75th Wealth Percentile	-	-1.3%	-2.7%
75–100th Wealth Percentile	-	-0.7%	-1.2%

Note: Columns correspond to alternative levels of the public long-term care consumption guarantee ($\bar{c}_2$). Panel A reports model-implied insurance participation rates. Panel B reports private long-term care insurance ownership rates across the wealth distribution. Panel C reports consumption-equivalent welfare changes relative to the Baseline, in aggregate and by wealth percentile, using population weights.

means testing, this public safety net disproportionately benefits poorer individuals and crowds them out of the private LTCI market. Wealthier individuals, who are less likely to rely on Medicaid, instead substitute toward self-insurance by accumulating substantial precautionary savings.[19]

Given these factors, we examine the extent to which public LTC generosity crowds out private insurance demand. Panel A of Table 12 evaluates private LTCI participation as we progressively reduce the public consumption floor ($\bar{c}_2$).[20] As Medicaid LTC support declines, there is a strong substitution into private LTCI. For example, reducing $\bar{c}_2$ to 40 more than triples the LTCI ownership rate from 13% to 44%. As Panel B highlights, this effect is most pronounced in the second wealth quartile, where participation rises from 27% to 72%.

[19]In the model, 27% of individuals in the bottom wealth quartile use Medicaid at some point in their lives, compared to just 8% in the top quartile. Heavy reliance on Medicaid extends into the second wealth quartile, where 26% of individuals eventually use public LTC. This occurs because the consumption floor of $\bar{c}_2 = 77.43$ is high relative to their typical wealth holdings, which would otherwise be exhausted by roughly two years of spending when in need of LTC.

[20]This analysis omits any effects from the mechanical improvements to the government budget from lower public LTC expenditures.

These middle-wealth individuals lack the substantial precautionary savings necessary to fully self-insure, yet they possess enough resources to afford private insurance and avoid diminished public care.

Panel C reports the corresponding welfare effects relative to the baseline economy. Lowering $\bar{c}_2$ has large aggregate costs—approximately 4% of lifetime consumption when $\bar{c}_2 = 40$—which are heavily concentrated among poorer individuals. For the bottom half of the wealth distribution, welfare falls by roughly 6%, whereas wealthier individuals are partially insulated because they rely more on personal savings than public care.

In summary, while the baseline provision of public LTC substantially crowds out private LTCI participation, it generates significant consumer welfare, particularly for individuals in the lower half of the wealth distribution.

5.6 The Effect of Social Security on Ownership and Welfare

Social Security benefits are embedded in our baseline estimates of life-cycle income profiles. Because Social Security functions essentially as a publicly provided annuity, we now evaluate how its existence affects private insurance ownership, wealth accumulation, and welfare. To do so, we simulate a counterfactual where we remove Social Security payments from the baseline estimation of income and instead credit individuals with the present value of the payments they would have otherwise received.[21] This counterfactual asks: how would behavior and welfare change if individuals received their Social Security benefits up front as a lump-sum wealth transfer rather than as a de facto mandatory annuity?

While one might expect individuals to replace the lost public annuity with private annuities, Panel A of Table 13 shows that they do not; private annuity participation remains low at 4%. This occurs because market frictions make private annuities an unattractive asset; when we eliminate both loads and nonpayment risk for annuities, ownership rises to 19%.

Relative to the baseline, the primary portfolio shift under the lump-sum counterfactual is a reduction in life insurance. Instead of purchasing formal life insurance, individuals self-insure using their increased liquid wealth. Panel B highlights this wealth accumulation: aggregate bequests rise by 23%, reflecting increases in both intentional and incidental bequests. Because individuals hold more wealth later in life to self-insure against longevity risk, early mortality generates larger bequests. However, when annuity market frictions are removed this effect is muted, with bequests rising by only 5%.

Panel C of Table 13 reports the welfare implications of the Social Security counterfactual. When we replace Social Security with an upfront lump-sum transfer, aggregate welfare

[21] We discount these payments using r_I, the same rate used to price private insurance products. See Appendix C.2 for details.

Table 13: **The Effects of Social Security**

	Data	**Baseline**	**No Social Security**	**Safe Annuities & No Social Security**
A. *Ownership Rate*				
LI	61%	55%	19%	30%
LTCI	12%	13%	14%	20%
Ann	4%	2%	4%	19%
B. *Aggregate Bequests*				
Δ vs. Baseline	-	-	23%	5%
C. *Welfare Gains*				
Δ Welfare	-	-	2.0%	4.1%

Note: Panel A reports insurance ownership rates in the data and across model scenarios. Panel B reports the percent change in aggregate bequests (both intentional and incidental) relative to Baseline. Panel C reports consumption-equivalent welfare changes relative to Baseline. No Social Security removes Social Security payments and rebates their present value to individuals as a lump-sum transfer.

improves by 2.0% over baseline. These gains are concentrated among poorer individuals, who use the windfall to smooth consumption intertemporally, significantly benefiting from the relaxation of their borrowing constraint. Welfare gains rise by 4.1% when private annuity market frictions are simultaneously removed. This 2.1% marginal improvement is significantly larger than the 0.7% gain from removing annuity frictions in the baseline model (Table 11). Intuitively, without Social Security crowding out demand, individuals benefit more from improved private annuities, reflected in both higher participation rates and welfare.

Overall, replacing Social Security with an up-front wealth transfer substantially reduces life insurance participation and has minimal effects on LTCI and private annuity participation. Welfare is higher on average in this counterfactual, driven by large gains for poorer individuals, but wealthy individuals prefer to have the de facto Social Security annuity, given the perceived nonpayment risk and loads associated with private market annuity alternatives.

6 Conclusion

Older Americans face substantial health, mortality, and longevity risks, yet participation in private insurance markets—particularly for annuities and long-term care insurance—remains puzzlingly low. While previous literature demonstrates that individuals would highly value perfect, complete-market insurance, we show that this lack of participation is largely driven by the imperfect nature of real-world products. A primary contribution of this paper is demonstrating that subjective nonpayment risk is a pervasive concern among consumers and a major deterrent to insurance ownership. Using a novel survey, we document that individuals expect to receive 76%–87% of promised payouts conditional on a qualifying event, and perceive annual full-default probabilities ranging from 1.3% to 1.8% across products. Validating these measures, we find that elicited beliefs predict actual insurance ownership.

Embedding these beliefs into a structural life-cycle model reveals that nonpayment risk and high pricing loads severely depress insurance demand, encourage precautionary saving, and reduce consumer welfare. Baseline access to existing, imperfect insurance products generates a modest welfare gain of 0.6% of lifetime consumption and drives a 16% decline in median wealth. Eliminating nonpayment risk and pricing loads substantially increases insurance participation rates—from 55% to 84% for life insurance, 13% to 55% for LTCI, and 2% to 10% for annuities. Consequently, the welfare gain from these idealized insurance products rises to 4.1% and median wealth falls an additional 17% below baseline. Finally, we document how public insurance programs like Medicaid and Social Security shape this environment.

References

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.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel, "The Economic Effects of Social Networks: Evidence from the Housing Market," *Journal of Political Economy*, 2018, *126* (6), 2224–2276.

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.

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.

_ , _ , Johanna Torres Chain, and Fang Yang, "Why do households save and work?," *working paper*, 2026.

Bornstein, Gideon and Sasha Indarte, "The Impact of Social Insurance on Household Debt," *Working paper*, 2023.

Braun, R. Anton, Karen A. Kopecky, and Tatyana Koreshkova, "Old, Sick, Alone and Poor: A Welfare Analysis of Old-Age Social Insurance Programs," *Review of Economic Studies*, 2017, *84* (2), 580–612.

_ , _ , and _ , "Old, Frail, and Uninsured: Accounting for Features of the U.S. Long-Term Care Insurance Market," *Econometrica*, 2019, *87* (3), 981–1019.

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

_ and Amy Finkelstein, "Why is the Market for Long-term Care Insurance so Small?," *Journal of Public Economics*, 2007, *91*, 1967–1991.

_ and _ , "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.

Catherine, Sylvain, Ellen Jiayang Lu, and James D Paron, "A dynamic model of the racial wealth gap," *Working paper*, 2025.

Christian, Rachel, "How to Sell Annuity Payments in 5 Steps," *Bankrate.com*, September 2025. https://www.bankrate.com/retirement/selling-an-annuity-how-to-do-it.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, Geoff Kenny, and Michael Weber, "The Effect of Macroeconomic Uncertainty on Household Spending," *American Economic Review*, March 2024, *114* (3), 645–77.

Conesa, Juan C. and Dirk Krueger, "Social Security Reform with Heterogeneous Agents," *Review of Economic Dynamics*, 1999, *2* (4), 757–795.

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.

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

_ , 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* (11), 3480–3520.

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

Doherty, Neil A and Harris Schlesinger, "Rational Insurance Purchasing: Consideration of Contract Nonperformance," *The Quarterly Journal of Economics*, 1990, *105* (1), 243–253.

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.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, "Five Facts about Beliefs and Portfolios," *American Economic Review*, May 2021, *111* (5), 1481–1522.

Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai, "Household Finance," *Journal of Economic Literature*, 2021, *59* (3), 919–1000.

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

Health and Retirement Study, "Health and Retirement Study, RAND HRS Longitudinal File public use dataset," [Computer file], Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740), Ann Arbor, MI 2024.

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.

Inkmann, Joachim, 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.

Koşar, Gizem and Cormac O'Dea, "Expectations data in structural microeconomic models," in Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw, eds., *Handbook of Economic Expectations*, Elsevier, 2023, chapter 21, pp. 647–675.

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

_ , 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.

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

_ , "Incidental Bequests and the Choice to Self-Insure Late-Life Risks," *American Economic Review*, 2018, *108* (9), 2513–2550.

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.

Mayers, David and Clifford W. Smith, "The Interdependence of Individual Portfolio Decisions and the Demand for Insurance," *Journal of Political Economy*, 1983, *91* (2), 304–311.

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

Nardi, Mariacristina De, Eric French, John Bailey Jones, and Rory McGee, "Why Do Couples and Singles Save during Retirement? Household Heterogeneity and Its Aggregate Implications," *Journal of Political Economy*, 2025, *133* (3), 750–792.

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.

RAND HRS, "RAND HRS Longitudinal File," [Computer file], Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging (grant numbers NIA U01AG009740 and NIA R01AG073289) and the Social Security Administration, Santa Monica, CA 2024.

Understanding America Study, "Understanding America Study 118: Insurance Product Choices," [Computer file], Center for Economic and Social Research (CESR) at the University of Southern California, Los Angeles, CA 2018.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion, "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications," *Journal of Economic Perspectives*, August 2022, *36* (3), 157–84.

Wiswall, Matthew and Basit Zafar, "Human Capital Investments and Expectations about Career and Family," *Journal of Political Economy*, 2021, *129* (5), 1361–1424.

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

The complete survey and collected data are available at `https://uasdata.usc.edu/survey/UAS+118`. The body of the paper includes the two main modules used to elicit our core nonpayment risk measures. This section highlights some of the other key modules.

Certainty Equivalent. The respondents expected payout is filled in from their previous responses. The guaranteed value is randomized.

> 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?

Beliefs Conditional on Aggregate Financial State.

> 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.

Effort in Collecting Payment.

> 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.

Conversion of of Lifetime Full-Default Probabilities to Annual Rates. In the survey, we randomly assigned respondents to one of two possible framings for the elicitation of the full-default probability: "We are now interested in the percent chance that" (i) "in a given year the annuity becomes worthless due to no fault of your own," and (ii) "the annuity becomes worthless due to no fault of your own at any point before the end of your life." We express all results in annualized terms, which requires converting responses from frame (ii) into annual rates. To do so, first we use the respondent's age, sex, permanent income, and health status, together with the health-state Markov transition matrix $\pi(s_{t+1} \mid s_t, t, f, y)$, to compute their life expectancy. Second, we divide the full-default lifetime probability by their expected years of remaining life.

B UAS Survey: Additional Results

Figure B.1: **Subjective Annual Partial Payout Distribution**

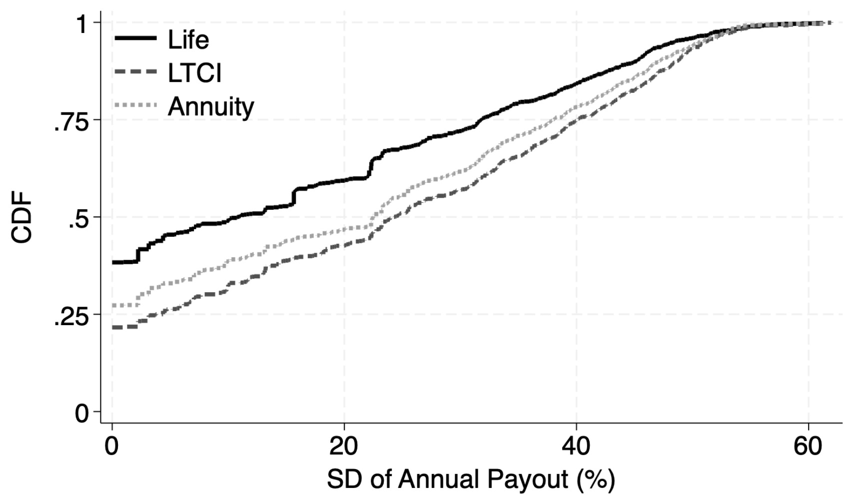

Note: Cross-sectional CDFs of the standard deviation of the subjective payout distribution conditional on a qualifying event. Payouts are expressed as a percentage of the contracted payout.

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

	Own Annuity	Own Life	Own LTCI	Own Annuity	Own Life	Own LTCI
Payment Exp. Val ($E[\omega^k]$)	-0.0018	0.0048***	0.0007	-0.0002	0.0040**	0.0007
	(0.194)	(0.001)	(0.090)	(0.770)	(0.010)	(0.131)
Full Def. Prob (δ^k)	-0.0022***	-0.0015	-0.0023***	-0.0021***	-0.0012	-0.0022***
	(0.000)	(0.129)	(0.000)	(0.000)	(0.167)	(0.000)
Payment S.D. ($\sigma(\omega^k)$)	-0.0043**	-0.0006	-0.0008	-0.0025***	-0.0006	-0.0008
	(0.002)	(0.712)	(0.246)	(0.001)	(0.733)	(0.246)
Trust				0.0194	-0.0053	0.0165
				(0.080)	(0.795)	(0.231)
Cognitive Score				-0.0011	-0.0023	0.0005
				(0.584)	(0.447)	(0.836)
Financial Literacy Score				-0.0149	-0.0696*	-0.0155
				(0.331)	(0.012)	(0.327)
Numeracy Score				-0.0081	0.0225	-0.0223
				(0.550)	(0.289)	(0.122)
Experienced Fraud				0.0369	0.0211	-0.0136
				(0.437)	(0.743)	(0.746)
Risk Aversion				-0.0069	-0.0156	-0.0020
				(0.277)	(0.085)	(0.699)
Propensity to Plan				0.0140	0.0014	0.0030
				(0.233)	(0.944)	(0.792)
Early Stock Returns				0.4653	-0.4527	-0.1782
				(0.140)	(0.363)	(0.629)
Male	0.0528*	-0.0213	-0.0208	0.0479	-0.0272	0.0066
	(0.047)	(0.619)	(0.445)	(0.100)	(0.543)	(0.799)
Age	-0.0074	0.0238	-0.0191	0.0126	0.0039	-0.0297
	(0.641)	(0.304)	(0.271)	(0.427)	(0.897)	(0.229)
Higher Education	0.0447	0.1087	0.1029	0.0539	0.1243	0.1039
	(0.347)	(0.394)	(0.097)	(0.318)	(0.280)	(0.163)
Married	-0.0327	-0.0484	0.0478	-0.0253	-0.0574	0.0404
	(0.259)	(0.401)	(0.220)	(0.393)	(0.313)	(0.299)
Retired	0.0212	-0.0693	0.0576	0.0023	-0.0966	0.0499
	(0.638)	(0.203)	(0.088)	(0.956)	(0.063)	(0.142)
Bad Health	-0.0706	-0.0450	-0.0394	-0.0539	-0.0515	-0.0442
	(0.053)	(0.484)	(0.076)	(0.116)	(0.408)	(0.088)
LTC Health	-0.0901**	0.0520	0.0285	-0.0963*	0.0601	0.0046
	(0.003)	(0.602)	(0.639)	(0.034)	(0.515)	(0.940)
Income ($million)	-0.0502	1.1503**	0.6993**	-0.0623	1.0854**	0.6390*
	(0.872)	(0.005)	(0.008)	(0.842)	(0.004)	(0.013)
Wealth ($million)	-0.0229*	-0.0041	-0.0003	-0.0195	-0.0042	-0.0013
	(0.049)	(0.982)	(0.986)	(0.095)	(0.814)	(0.948)
N	1054	1045	1039	1054	1045	1039
R^2	0.161	0.132	0.129	0.267	0.214	0.175

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

Table B.2: **Regression of Ownership on Risk Metrics: Sample of Internally-Consistent Respondents**

	Own Annuity	Own Life	Own LTCI	Own Annuity	Own Life	Own LTCI
Payment Exp. Val ($E[\omega^k]$)	-0.0019	0.0048***	0.0009*	0.0002	0.0037*	0.0008
	(0.236)	(0.001)	(0.047)	(0.745)	(0.018)	(0.098)
Full Def. Prob (δ^k)	-0.0017**	-0.0015	-0.0022***	-0.0017**	-0.0011	-0.0022***
	(0.007)	(0.129)	(0.000)	(0.004)	(0.228)	(0.000)
Payment S.D. ($\sigma(\omega^k)$)	-0.0046**	-0.0006	-0.0009	-0.0024**	-0.0008	-0.0008
	(0.003)	(0.712)	(0.257)	(0.004)	(0.656)	(0.308)
Trust				0.0227	-0.0117	0.0233
				(0.057)	(0.597)	(0.123)
Cognitive Score				-0.0018	-0.0022	0.0015
				(0.438)	(0.502)	(0.581)
Financial Literacy Score				-0.0041	-0.0793**	-0.0106
				(0.808)	(0.009)	(0.574)
Numeracy Score				-0.0205	0.0240	-0.0312
				(0.174)	(0.305)	(0.064)
Experienced Fraud				0.0549	0.0459	0.0002
				(0.319)	(0.514)	(0.996)
Risk Aversion				-0.0086	-0.0142	-0.0026
				(0.200)	(0.154)	(0.649)
Propensity to Plan				0.0018	-0.0131	0.0063
				(0.882)	(0.527)	(0.624)
Early Stock Returns				0.3934	-0.8636	-0.1302
				(0.231)	(0.088)	(0.760)
Male	0.0516	-0.0213	-0.0338	0.0503	-0.0461	0.0030
	(0.086)	(0.619)	(0.269)	(0.123)	(0.350)	(0.922)
Age	0.0049	0.0238	-0.0325	0.0218	-0.0110	-0.0423
	(0.752)	(0.304)	(0.089)	(0.204)	(0.720)	(0.124)
Higher Education	0.0280	0.1087	0.1145	0.0442	0.2009	0.1160
	(0.561)	(0.394)	(0.104)	(0.429)	(0.109)	(0.167)
Married	-0.0312	-0.0484	0.0619	-0.0229	-0.0562	0.0543
	(0.330)	(0.401)	(0.154)	(0.479)	(0.352)	(0.205)
Retired	0.0516	-0.0693	0.0651	0.0239	-0.1351*	0.0636
	(0.310)	(0.203)	(0.084)	(0.599)	(0.014)	(0.097)
Bad Health	-0.0532	-0.0450	-0.0314	-0.0373	-0.0380	-0.0356
	(0.177)	(0.484)	(0.212)	(0.299)	(0.575)	(0.224)
LTC Health	-0.1055***	0.0520	0.0393	-0.1258*	0.0815	0.0066
	(0.001)	(0.602)	(0.566)	(0.012)	(0.406)	(0.925)
Income ($million)	0.3509	1.6675**	0.7293*	0.3376	1.4445***	0.6152*
	(0.275)	(0.005)	(0.027)	(0.295)	(0.001)	(0.048)
Wealth ($million)	-0.0121	-0.0118	0.0066	-0.0092	0.0006	0.0066
	(0.335)	(0.982)	(0.769)	(0.466)	(0.974)	(0.762)
N	870	861	855	870	861	855
R^2	0.136	0.132	0.146	0.249	0.233	0.195

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

Table B.3: **Regressions of risk measures on demographic and behavioral controls: Sample of Internally-Consistent Respondents**

	$\mathbb{E}\left[\omega^{\text{Ann}}\right]$	$\mathbb{E}\left[\omega^{\text{LI}}\right]$	$\mathbb{E}\left[\omega^{\text{LTC}}\right]$	$\sigma\left(\omega^{\text{Ann}}\right)$	$\sigma\left(\omega^{\text{LI}}\right)$	$\sigma\left(\omega^{\text{LTC}}\right)$	δ^{Ann}	δ^{LI}	δ^{LTC}
Trust	1.0977	0.1183	-0.9709	2.0273	0.3521	-0.8306	0.4604	-0.7580	-0.2009
	(0.308)	(0.920)	(0.397)	(0.197)	(0.705)	(0.456)	(0.667)	(0.486)	(0.816)
Cognitive Score	0.2918	-0.2956	-0.0826	0.4496*	-0.3541**	-0.1473	0.1054	-0.3976**	0.1229
	(0.061)	(0.069)	(0.598)	(0.046)	(0.008)	(0.455)	(0.521)	(0.006)	(0.352)
FinLit Score	6.2454**	-5.7096*	-3.2941	1.0647	-3.4585*	-0.7321	3.8532*	-3.0576*	-1.6330
	(0.002)	(0.011)	(0.085)	(0.659)	(0.033)	(0.712)	(0.025)	(0.038)	(0.257)
Numeracy Score	0.3544	-1.8488	-0.0476	0.8821	-0.8468	-1.2302	2.1549	-1.7082	-2.7506*
	(0.787)	(0.164)	(0.969)	(0.614)	(0.467)	(0.390)	(0.118)	(0.166)	(0.012)
Experienced Fraud	2.3797	11.4553*	-2.9870	-2.1338	9.6440*	0.2470	0.2863	9.7430**	-2.9601
	(0.645)	(0.043)	(0.494)	(0.613)	(0.012)	(0.951)	(0.957)	(0.006)	(0.418)
Risk Aversion	0.6932	-0.3524	-0.6117	-0.3328	-0.1490	0.5058	-0.3963	0.2742	0.1350
	(0.245)	(0.481)	(0.218)	(0.606)	(0.779)	(0.359)	(0.529)	(0.610)	(0.759)
Plan Ahead	-2.3545	-0.0698	2.3770*	-1.4989	0.5972	2.9070*	-2.6377*	-1.1701	-0.1551
	(0.059)	(0.960)	(0.038)	(0.269)	(0.583)	(0.014)	(0.044)	(0.293)	(0.863)
Cohort Return	-61.0428**	41.3907	57.7057**	-14.9360	30.8438	56.7812*	-0.8987	29.8672	7.8004
	(0.007)	(0.103)	(0.003)	(0.651)	(0.119)	(0.015)	(0.971)	(0.193)	(0.655)
N	870	886	870	866	886	862	855	888	855

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

Note: Tobit regression estimates predicting the expected annual payout share, standard deviation of the subjective payout distribution, and annual full-default probability. All specifications include demographic and supplementary individual characteristics. Robust standard errors are reported in parentheses.

C Estimated Inputs for Model

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

C.1 Health Transitions and Cost Shocks

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.

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 with lower permanent income, $y \in \{1, 2, 3\}$, and those with higher permanent income, $y \in \{4, 5\}$. This is to allow the model to capture the well-documented wealth gradient in health and longevity.

Figures C.2 and C.3 display the estimated health transitions. Those with higher permanent income have more favorable health transitions. For example, they are more likely to stay healthy throughout their life and, they are more likely to survive conditional on requiring LTC.

Health Cost Shocks. To estimate the medical 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 types 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 medical cost distribution, we regress the log of medical expenses on a quartic polynomial in 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

[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.

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

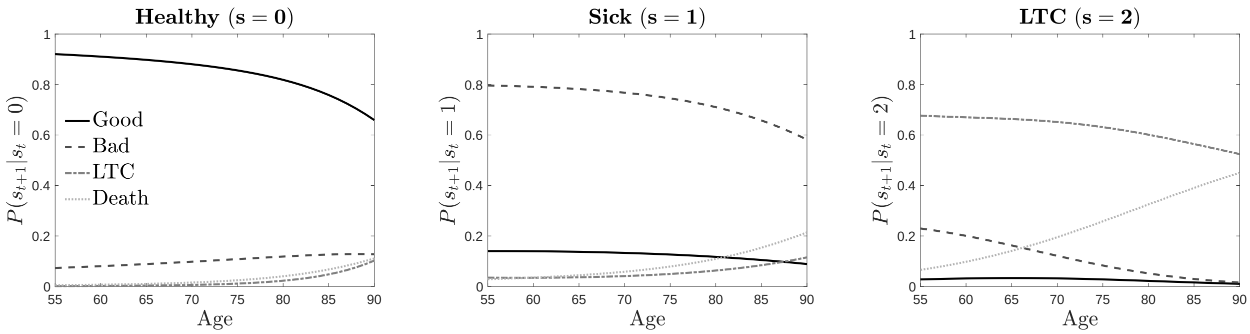

Note: Each panel reports model-implied age-specific health transition probabilities for individuals in the bottom 3 permanent income quintiles. The three panels correspond to current health states Healthy, Sick, and LTC. Within each panel, the four lines show the probability of transitioning to Good, Bad, LTC, or Death.

Figure C.3: **Health Transition Profiles: Higher Income** ($PIQ \in \{4, 5\}$)

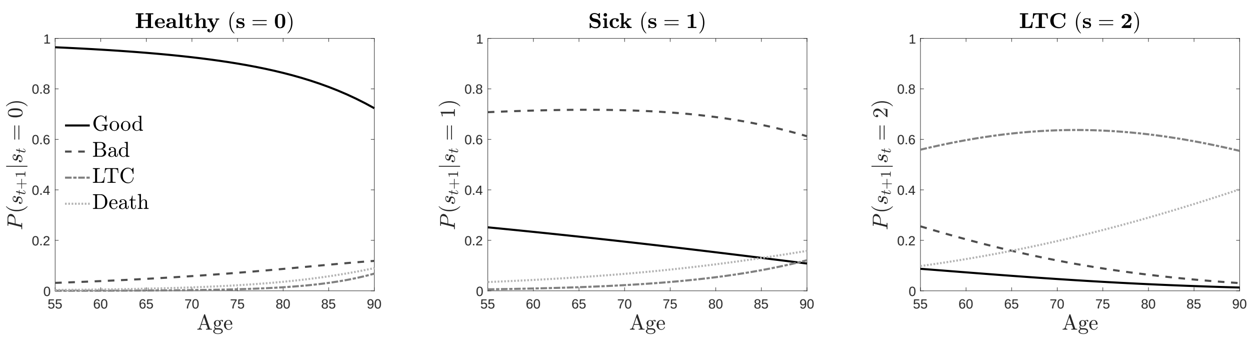

Note: Each panel reports model-implied age-specific health transition probabilities for individuals in the top two permanent income quintiles. The three panels correspond to current health states Healthy, Sick, and LTC. Within each panel, the four lines show the probability of transitioning to Good, Bad, LTC, or Death.

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.4 plots the mean medical cost shock over the life cycle for both men 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.

Figure C.4: **Mean Health Cost Shock Profile**

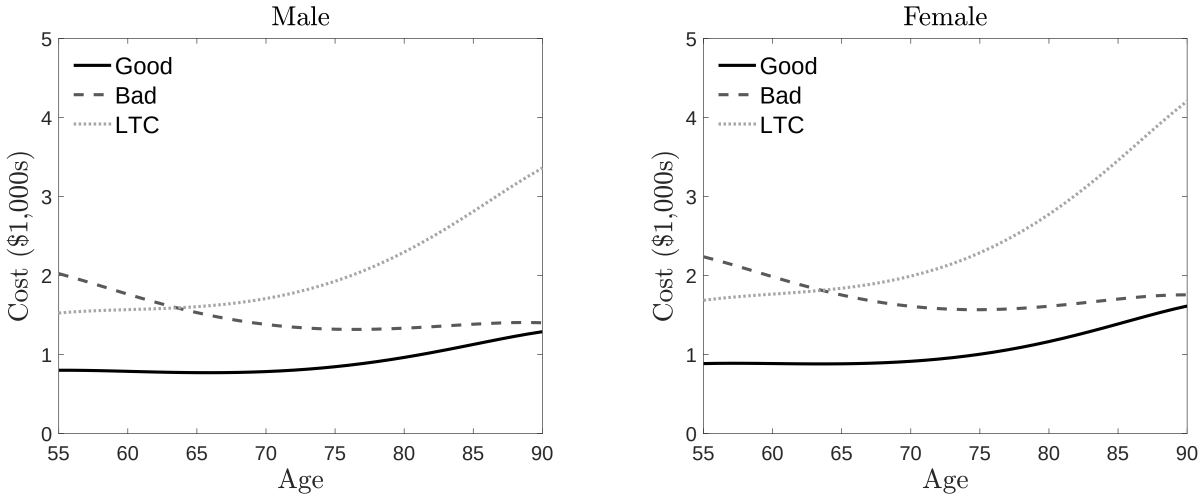

Note: The figure plots the mean of the medical cost shock distribution over the life cycle for males (left panel) and females (right panel).

C.2 Income Process

We estimate a deterministic labor income profile for five separate income quintiles. We define income using the RAND HRS 2014 dataset as the sum of income from earnings, public and private pensions and annuities, social security income, disability and retirement, unemployment compensation, and other government transfers.[23] 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.

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, De Nardi and Yang (2023), who show that the following tax function can capture the federal taxes paid by individuals 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

[23]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).

these year-based estimates in evaluating after-tax income for individuals from the RAND HRS 2014 dataset. This excludes state-income taxes. Fleck, Heathcote, Storesletten and Violante (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 defining 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

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

where $inc_{i,w}$ is our measure of income for individual i in wave w. The key idea here is that we have fixed effects at the individual 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 individuals for a given (a, PIQ, g) combination.[24] 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 $inc_{a,\hat{PIQ},g}$. 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.5. 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 individuals 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

[24]Results are very similar, although a bit smoother, if we do not condition on age a in the combination.

Figure C.5: **Deterministic Labor Income Profiles by Quintile**

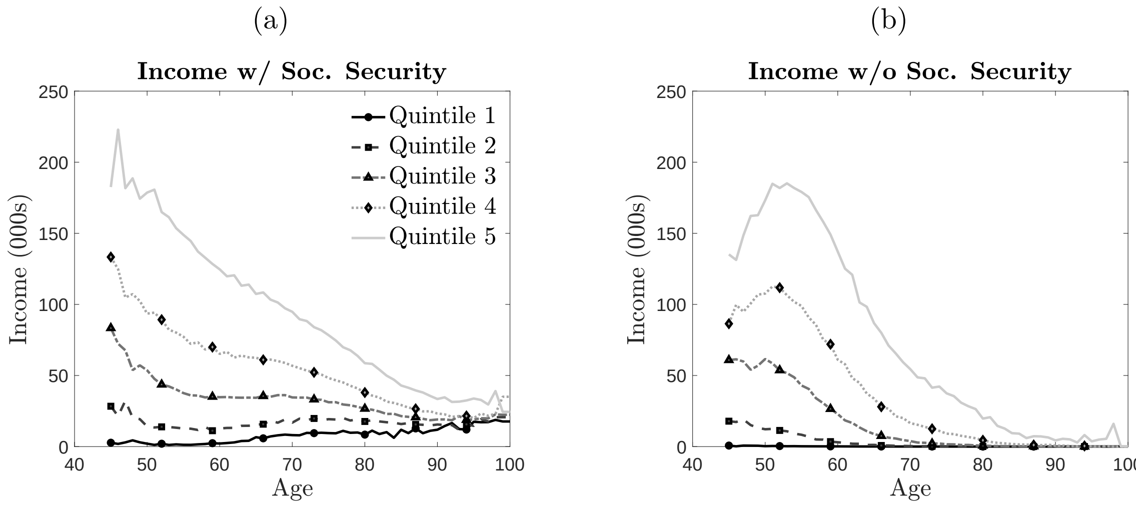

Note: Each line reflects the age profile of deterministic labor income for a given income quintile, estimated from the HRS panel data. The left panel includes social security payments, while the right panel excludes them.

C.5. As we can see, labor income profiles shift down, with there being a larger proportional effect on the lower income quintiles.

C.3 Assets

Financial wealth is defined as the sum of an IRA (individual retirement account) and employer-sponsored retirement, checking, saving, money market, mutual fund, certificate of deposit, brokerage, and education-related accounts. Total wealth is defined as financial wealth plus housing wealth. For insurance, 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.

D Simulation and Estimation of Model

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 first 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 $X_i = \{s_i, f, y_i, b_i, x_i^{\text{Ann}}, x_i^{\text{LI}}, x_i^{\text{LTC}}, p_i^{\text{LTC}}, p_i^{\text{LI}}\}$. The model is solved by conducting a grid search over the state variables. Once we have the policy functions defined over X_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, $\hat{X}_i$ for each individual. We simulate each individual 100 times, drawing relevant shocks $\hat{\Psi}$ from the distributions of full default, annual nonpayment, medical costs as well as the health transitions, and simulate the behavior implied by the optimal policies.[25] We then aggregate these individual behaviors to construct simulated moments from the model, $o(\hat{X}, \theta, \hat{\Psi})$, 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 nine 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 calibrating the model.

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

D.2 Calibration Strategy

We develop a calibration 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 $\Theta := \{\beta, \theta_{3,45}, \gamma, m, r\}$. 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. First, we target wealth percentiles across the life cycle, which are commonly used in life-cycle savings models. Second, we target average participation

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

rates for each of the three insurance products that we analyze. 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 behavioral 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 generated by the model, $o(\hat{X}, \Theta, \hat{\Psi})$, and empirical data, $s(Z)$ (which we define as $g(\hat{X}, \Theta, \hat{\Psi}) = o(\hat{X}, \Theta, \hat{\Psi}) - s(Z)$). We then solve the following problem:

$$\hat{\Theta} = \arg\min_{\Theta} g(\hat{X}, \Theta, \hat{\Psi})' W g(\hat{X}, \Theta, \hat{\Psi})$$

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 is equal. We now define the three categories of moments below:

D.2.1 Wealth Moments

The first set of moments consists of age-specific wealth percentiles. Specifically, we target the 25th, 50th, and 75th percentiles of the wealth distribution within 5-year age brackets (50–54, 55–59, . . . , 80–84). We stop at age 85 due to small sample sizes at older ages. This yields eight age groups and, with three percentiles per group, a total of 24 targeted wealth moments. Model-implied percentiles are directly compared to their empirical counterparts.

To construct the wealth moments in the model, we treat each simulated individual-age observation as a cross-sectional data point. That is, every time one of the 108,500 simulated individuals is observed at a given age, it enters the corresponding age bracket. For example, the percentiles for ages 50–54 are computed using all simulated observations that fall within that age range.

Empirical wealth moments are constructed using the HRS. We use the HRS rather than the UAS because it provides a larger and therefore less noisy estimate of the wealth distribution. This approach relies on the UAS being broadly representative of the HRS population. As shown in Table 1, however, the UAS sample differs somewhat from the HRS, particularly in marital and retirement 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 individuals 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

We compute insurance participation rates similar to how we compute wealth moments. The key difference is that we only target the aggregate participation rates of life insurance, annuities and LTCI, leaving untargeted the life-cycle profile of insurance ownership. Our empirical moments are the aggregate participation rates from the 2014 HRS, using the adjusted sample weights as described above. This generates empirical targets of 61% for life insurance, 12% for LTCI, and 4% for annuities.

D.2.3 SSQ Moments

This final moment set is constructed entirely from SSQ responses of the sample of individuals in Ameriks et al. (2020). The key difference here is that we fix the parameter values of $(\theta_2, \bar{\theta}_3, \kappa_2, \kappa_3, \bar{c}_2)$ to those estimated by the paper as their posterior means, leaving us only with γ as the estimated variable. The intuition for this is that γ is the parameter that, locally, most affects insurance ownership moments, which were not studied in Ameriks et al. (2020).

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)**

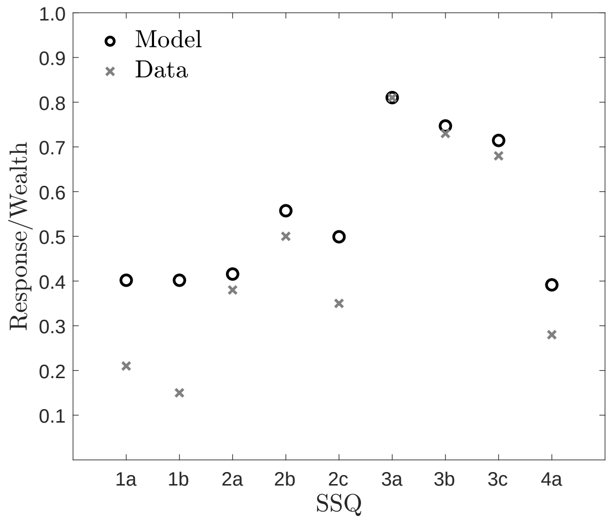

Note: Shows mean responses to the nine SSQs from Ameriks et al. (2020), in the data via the cross and the model via the circles. Compared to Ameriks et al. (2020), in order to match low insurance demand, our baseline calibration has a smaller coefficient of relative risk aversion coefficient. This worsens our fit to questions 1a and 1b in the SSQs.

D.3 Computing Consumption-Equivalent Welfare

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

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

where n indexes each simulation of health states, medical 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

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

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 individuals 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 - γ

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

In our calibration, we need γ to be relatively low in order to be close in matching LTCI participation rates. Otherwise there is too much demand for LTCI. 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 individuals. The main goal here is to ensure that the bequest motive for young individuals is strong enough such that their life insurance participation is high. Doing this requires raising $\theta_{3,45}^{-\gamma}$. There is a balancing act, however, because this also raises LTCI 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.2. As a result, LTCI acts similarly to life insurance by providing incidental bequests when young.

D.4.4 Parameter ρ

This parameter ρ controls the shape of the bequest motive over the life cycle, governed by equation (2). For ρ that is very close to 0, it approximates a linear function, but as ρ increases it tilts the bequest motive when young higher and creates a steeper declining slope for later ages. Our calibration is for ρ of 1.5 is 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., $\rho \to 0$), life insurance ownership would fall too quickly with age, hence why we need a higher ρ.

The value of ρ can't be set too high or it would generate too high life insurance ownership rates at young ages. 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. 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.

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 nonpayment risk are non-rational. In these experiments, insurance is safe, meaning that payouts are always made. Subjective beliefs, however, are the same as in the main text, calibrated to those elicited in the survey.

E.1 Insurance Ownership and Savings: Non-Rational Expectations

Table E.4 repeats the same exercise as under Table 9 but under non-rational expectations. Results are very similar in the two cases. Ultimately, portfolio choice is driven by beliefs, whether or not they are accurate. The slight differences arise because the evolution of wealth in the simulations differs—individuals in this safe-insurance setting receive more payments than if their beliefs about risky insurance were accurate.

Table E.4: **Insurance Ownership and Savings: Impact of Risky Insurance and Loads (Non-Rational Expectations)**

	Data	Baseline	Safe Insurance	Zero Loads	Zero Loads & Safe Insurance
A. *Ownership Rate*					
LI	61%	56%	73%	72%	84%
LTCI	12%	15%	43%	30%	55%
Ann	4%	2%	8%	3%	10%
B. Δ *Personal Savings*					
25th percentile	–	−38%	−48%	−47%	−55%
50th percentile	–	−17%	−27%	−23%	−33%
75th percentile	–	−6%	−16%	−14%	−22%

Note: Panel A reports aggregate ownership rates for three insurance products under alternative scenarios. Panel B reports percentage changes in wealth across the distribution, relative to values from a no-insurance benchmark, averaged across age with population weights. Baseline is the calibrated model; Safe Insurance sets default probabilities to zero (i.e., no nonpayment risk); Zero Loads sets loads to zero; Zero Loads & Safe Insurance combines both adjustments under non-rational expectations.

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

Table E.5 repeats the same exercise as in Table 11 but under non-rational expectations.

Table E.5: **Welfare Gains from Individual Insurance Products (Non-Rational)**

	Baseline	Safe Insurance	Zero Loads	Zero Loads & Safe Insurance
LI	0.7%	0.1%	1.0%	1.1%
LTCI	0.9%	0.3%	0.4%	0.6%
Ann	0.0%	0.2%	0.4%	0.7%
All	1.6%	0.7%	1.7%	2.4%

Note: Consumption-equivalent welfare gains relative to a no-insurance benchmark (non-rational beliefs case). Each row isolates the introduction of only the specified product. The final row introduces all three products and reproduces Panel A from Table E.4. Baseline is the calibrated model; Safe Insurance sets all default probabilities to zero; Zero Loads sets loads to zero; Zero Loads & Safe Insurance sets all nonpayment probabilities and loads to zero.

E.3 Role of Public Long-Term Care: Non-Rational Expectations

Table E.6 runs the same experiment as in Table 12 but under the assumption of non-rational expectations.

Table E.6: **The Effects of Government LTC Provision (Medicaid): Non-Rational**

	Baseline $\bar{c}_2 = 77.43$	$\bar{c}_2 = 60$	$\bar{c}_2 = 40$
A. *Ownership Rate*			
LI	56%	55%	52%
LTCI	15%	26%	45%
Ann	2%	2%	2%
B. *LTCI Ownership Rate*			
0–25th Wealth Percentile	10%	13%	43%
25–50th Wealth Percentile	29%	45%	74%
50–75th Wealth Percentile	16%	31%	57%
75–100th Wealth Percentile	13%	37%	47%
C. *Welfare Effects*			
Aggregate	-	-1.0%	-2.8%
0–25th Wealth Percentile	-	-1.2%	-4.6%
25–50th Wealth Percentile	-	-1.4%	-4.7%
50–75th Wealth Percentile	-	-1.1%	-1.6%
75–100th Wealth Percentile	-	-0.3%	-0.3%

Note: Columns correspond to alternative levels of the public long-term care consumption guarantee ($\bar{c}_2$). Panel A reports model-implied insurance participation rates. Panel B reports private long-term care insurance ownership rates across the wealth distribution. Panel C reports consumption-equivalent welfare changes relative to the Baseline, in aggregate and by wealth percentile, using population weights.

E.4 Role of Social Security: Non-Rational Expectations

This runs the same set of experiments as under Table 13 but under non-rational expectations.

Table E.7: **The Effects of Social Security: Non-Rational Expectations**

	Data	**No Insurance**	**Baseline**	**No Social Security**
A. *Ownership Rate*				
LI	61%	0%	56%	20%
LTC	12%	0%	15%	14%
Ann	4%	0%	2%	4%
B. *Welfare Gains*				
Δ **Welfare**	-	-	-	1.6%

*Note:*Panel A reports insurance ownership rates in the data and across model scenarios. Panel B reports consumption-equivalent welfare changes relative to Baseline. No Social Security removes Social Security payments and rebates their present value to individuals as a lump-sum transfer.