Revisiting Retirement and Social Security Claiming Decisions

Neha Bairoliya* Neha.Bairoliya@marshall.usc.edu

Kathleen McKiernan† kathleen.mckiernan@vanderbilt.edu

March 26, 2026

Abstract

The structure of the Social Security program presents a trade-off between benefit duration and size; individuals who claim early receive lower monthly payments. Surprisingly, even wealthy Americans with high life expectancy frequently claim benefits at the earliest eligibility age, a behavior inconsistent with conventional life-cycle theory predictions. In fact, there is a striking similarity in claiming ages across groups with vastly different life expectancies, wealth levels, and incentives to delay. This paper resolves this puzzle by introducing policy details and subjective beliefs about longevity and program rules into a rich, forward-looking life-cycle model. These factors, particularly subjective beliefs, vary significantly across socioeconomic groups and incentivize early claiming among high-wealth individuals. When included in the model, subjective beliefs substantially narrow the claiming age gap between education (college vs. non-college) and marital status (married vs. single) groups. Specifically, differences in program beliefs alone account for 48.6% of the education claiming age gap and 75.1% of the marital status gap. Counterfactual policy simulations demonstrate that failing to account for these claiming mechanisms dramatically overstates individual responses to an increase in the normal retirement age and understates the resulting fiscal savings.

Keywords: Labor supply, Social Security, annuity, misbeliefs, life cycle, health, marriage, spousal benefits, survivors benefits, bequest motive, preference heterogeneity.

JEL Classification Numbers: J14, J26, E21, H55

*University of Southern California, Department of Finance and Business Economics, Los Angeles, CA 90089-1422, United States.

†Vanderbilt University, Department of Economics, Nashville, TN 37235, United States.

1 Introduction

An aging U.S. population is increasing fiscal pressure on the Social Security system and prompting urgent discussions around policy reform. The potential impact of any such policy, however, depends critically on how individuals respond to these changes. This, in turn, requires a deep understanding of the mechanisms behind current decisions—particularly when such decisions appear at odds with standard economic theory. In this paper, we present a comprehensive framework for understanding the drivers of Social Security claiming behavior and how the strength and nature of these drivers vary across different socioeconomic groups. Among other important insights, our framework demonstrates that models that do not account for the behavior of heterogeneous individuals overstate responses to an increase in the normal retirement age (NRA) and, thus, may underestimate the resulting fiscal savings of the reform.

Workers can claim Social Security benefits as early as age 62. However, claiming benefits prior to the normal retirement age is heavily penalized: early claimants receive a smaller fraction of benefits annually than if they had waited. As shown in prior work (Coile et al., 2002; Sun and Webb, 2011; Meyer and Reichenstein, 2010; Shoven and Slavov, 2014a,b), many Americans would benefit from delaying claims. Yet, among a cohort of men born 1931–1935, over two-thirds claimed before the NRA and the average claiming age was just 63.4 years.

Rust and Phelan (1997) and Gustman and Steinmeier (2005) document sharp spikes in retirement at ages 62 and 65, with the latter emphasizing that these patterns are difficult to reconcile within a standard life-cycle framework. This ignited a robust structural literature studying retirement and claiming decisions. Proposed explanations include heterogeneity in preferences (Gustman and Steinmeier, 2005, 2008; Maurer et al., 2021), stochastic returns on savings (Gustman and Steinmeier, 2015), reference dependence (Behaghel and Blau, 2012; Brown et al., 2016), bequest motives (Pashchenko and Porapakkarm, 2024), uncertainty about future Social Security solvency (Benitez-Silva et al., 2009; Gustman and Steinmeier, 2015), and specific policy rules, such as marital benefits (Gustman and Steinmeier, 2015) or the retirement earnings test (Pashchenko and Porapakkarm, 2024). Although these mechanisms explain why many individuals claim at or before the normal retirement age, they leave a crucial puzzle unresolved: why claiming decisions remain uniform across groups that face distinct economic incentives.

College-educated men live 4.3 years longer than non-college-educated men, and married men live 4.4 years longer than singles. Longer lifespans among higher socioeconomic status (SES) groups strengthen incentives to delay benefit claiming by extending the expected payout horizon. Standard economic reasoning would predict that these groups delay claiming relative to those with shorter expected lifespans. However, the average claiming age differs by only 0.6 years between education groups and by 0.1 years between marital status groups. While early claiming is more

prevalent among low-income individuals (Venti and Wise, 2014; Armour and Knapp, 2023), the magnitude of these differences is far smaller than variations in life expectancy would imply. This striking homogeneity in claiming behavior across groups with such divergent longevity presents a novel puzzle that standard explanations do not fully resolve.

To address this puzzle, we develop a rich life-cycle model that incorporates both economic and belief-based mechanisms driving claiming decisions. The model is designed not only to replicate overall claiming behavior but also to extend the previous literature and explain the claiming behavior of groups who differ vastly in life expectancies and financial resources. We find that while various channels contribute to claiming decisions, subjective beliefs about program rules and mortality play a central role in encouraging early claiming among high-SES individuals and compressing differences in claiming ages across education and marital status groups.

Our model includes household decisions regarding consumption, savings, labor supply, and Social Security claiming. It explicitly captures key features of the U.S. Social Security system, including the trade-off between claiming age and benefit size, the details of the retirement earnings test, and the presence of spousal and survivors benefits available to married households. Agents are heterogeneous with respect to education and marital status and face exogenous shocks to health, labor productivity, and employment. Additionally, they hold subjective beliefs about life expectancy and penalties for early claiming. Survival beliefs evolve over the life cycle based on age, education, marital status, and health shocks. Agents also recognize that their beliefs about early claiming penalties might be incorrect and thus assign strictly positive probabilities to alternative outcomes. This comprehensive quantitative framework captures the fundamental economic trade-offs shaping claiming decisions and detailed policy provisions, such as marital benefits, which may affect individuals differently. The breadth of these features is critical to our estimation strategy.

We estimate our model in two steps. In the first step, we directly estimate key components from data: health transitions, subjective survival beliefs, family structure, labor productivity, employment status, initial conditions, and policy parameters. In the second step, we structurally estimate preference parameters and the distribution of subjective beliefs about the Social Security penalty separately by education and marital status. Preference parameters target the life-cycle evolution of wealth and labor supply, while penalty uncertainty is inferred from observed claiming behavior.

This data-driven and theoretically grounded approach to estimating subjective program beliefs combines survey responses with structural estimation. We first capture stated beliefs and self-reported confidence directly from the data and then use a revealed preference framework to infer the unobserved probability weights individuals place on alternative penalty outcomes. We incorporate these survey-based belief types and confidence weights into our structural model to quantify the full distribution of subjective beliefs about Social Security penalties. We find that college-educated and married men exhibit higher confidence and more concentrated belief distributions.

We also find that belief uncertainty can have opposite effects: it incentivizes delaying claims for low-SES individuals but encourages early claims among high-SES groups.

The model successfully replicates overall Social Security claiming patterns by education and marital status for the 1931–1935 birth cohort. It also captures key variation in claiming behavior across wealth, health, and employment status. Importantly, it accounts for spousal claiming decisions in married households, distinguishing between claiming on one's own earnings record and claiming spousal benefits. With confidence in the structural model, we conduct counterfactual experiments to quantify the role of the mechanisms we study—particularly subjective beliefs about program rules and mortality—in explaining both overall claiming behavior and the homogeneity in claiming ages across education and marital status groups.

Our counterfactual experiments yield several interesting insights. First, in a simplified model where claiming is driven by empirically estimated heterogeneous objective survival probabilities and preferences consistent with observed labor supply and wealth profiles, the predicted average claiming age exceeds the data by over two years—with larger gaps for college-educated and married men. In this framework, low estimated impatience leads all workers to delay claiming. Moreover, longer life expectancy and relatively higher discount factors among high-SES groups generate sizable claiming age gaps: college-educated men claim 1.7 years later than non-college men, and married men claim 2.8 years later than singles. Adding standard mechanisms—health shocks, employment risk, and bequest motives—raises early claiming rates across all groups but leaves these large SES claiming age gaps largely intact.

Second, our results highlight that the heterogeneous policy treatment of married versus single men plays an important role in shaping claiming behavior—particularly among college-educated and married men. While spousal and survivors benefits exert opposing incentives, our estimated model indicates that the incentive to claim early due to spousal benefits outweighs the delaying effect of survivors benefits. These marital benefits, therefore, increases early claiming among married men. Notably, college-educated married men are more responsive to policy incentives than their less-educated counterparts, and the inclusion of marital benefits reduces their claiming age more substantially. Nonetheless, even with this greater responsiveness among high-SES groups, the model still predicts sizable claiming-age differences by education and marital status: college-educated men are expected to claim one year later than non-college-educated men, and married men 1.9 years later than single men—gaps that remain larger than those observed in the data.

Third, subjective beliefs about program rules have the largest impact on claiming behavior, primarily by shifting decisions among high-SES men. At the time of claiming average estimated subjective life expectancies fall below objective ones, especially among college-educated and married men. Reduced survival expectations function similar to a lower discount factor and lead to earlier claiming. Beliefs about the early claiming penalty and uncertainty around those beliefs also

contribute by altering the perceived costs of early claiming. Expectations of a small early claiming penalty decreases the perceived penalty among high SES men and amplifies early claiming in these groups. When both subjective program and survival beliefs are included, the model predicts widespread early claiming and small gaps in average claiming ages: 0.6 years across education groups and 0.3 years across marital statuses. These beliefs together explain all of the compression of education-related and marriage-related gap between a naïve baseline and our full quantitative specification. Differences in program beliefs alone account for 48.6% of the education claiming age gap and 75.1% of the marital status gap.

Finally, we use the model to simulate a policy reform that raises the NRA to age 70. We find that incorporating mechanisms such as subjective program beliefs significantly dampens responses to the reform. While simpler models predict substantial delays in claiming, our full model yields only modest shifts. As a result, models that omit these mechanisms overstate behavioral responsiveness and substantially understate the fiscal impact of the reform. Specifically, they underestimate the reduction in total lifetime Social Security benefits net taxes paid to the 1931–1935 male cohort by as much as $19.6 billion. These reform projections hold constant the subjective beliefs of the households. If, however, an increase in the NRA resets beliefs about the early claiming penalty or other policy features, behavioral adjustment could be distinct from our partial-information environment in which misperceptions about the annual early-claiming adjustment remain fixed.

Our paper makes important contributions to several strands of the literature. It is most closely related to prior studies such as Gustman and Steinmeier (2005, 2015); Benitez-Silva et al. (2009), and Pashchenko and Porapakkarm (2024) studying Social Security claiming behavior using structural models. While these studies have advanced our understanding of early claiming by incorporating heterogeneity in discount rates, liquidity constraints, bequest motives, and uncertainty about program solvency, they focus on explaining the average timing of claims. In contrast, our paper highlights and explains a novel empirical puzzle that has received little attention: the surprising similarity in claiming behavior across individuals with widely different lifespans and incentives.

Our model builds on prior work by incorporating many of the same mechanisms. However, it advances the literature by incorporating heterogeneous subjective beliefs about Social Security program rules and longevity estimated from micro data and demonstrating that these beliefs play an important role in compressing expected heterogeneity in claiming behavior. While earlier studies such as Benitez-Silva et al. (2009) and Gustman and Steinmeier (2015) model uncertainty about Social Security program solvency, they rely on inferred beliefs that are either not empirically grounded or not relevant to our cohort of interest, and—as Gustman and Steinmeier (2015) themselves note—lead to predictions (e.g., rising early claiming over time) that contradict observed

trends.[1] In contrast, our approach uses data on informed beliefs, varies them across demographic groups, and offers a more empirically disciplined explanation for observed claiming behavior. An additional, alternative behavioral explanation emphasizes that statutory ages may serve as reference points, generating bunching through loss aversion or framing effects (Behaghel and Blau, 2012; Brown et al., 2016). While these reference-point forces likely operate alongside standard incentives, a simple one-reference-age model cannot generate the dual-peaked claiming distribution in U.S. data without additional structure and identifying information. Our revealed preference approach to structurally estimate policy uncertainty, on the other hand, will reflect other behavioral forces, such as reference dependence, while also generating the claiming patterns observed.

We also contribute to a growing literature that uses rich structural life-cycle models of retirement, consumption, and savings to address questions of macroeconomics, labor economics, and public policy (Fan et al., 2022; Jones and Li, 2023; Bairoliya, 2019; Borella et al., 2023; De Nardi et al., 2010; Imrohoroğlu and Kitao, 2012; Jones and Li, 2018; Yu, 2024; French and Jones, 2011; Hubener et al., 2016; French, 2005; Hosseini et al., 2021). Our main contribution is to incorporate subjective beliefs into household decision-making. Prior structural work incorporating these beliefs has focused primarily on lifespan expectations (Van der Klaauw and Wolpin, 2008; Gan et al., 2015; Heimer et al., 2019; O'Dea and Sturrock, 2023; de Bresser, 2024). Van der Klaauw and Wolpin (2008) was among the first to integrate subjective expectations into a structural life-cycle model of retirement, studying a sample of relatively poor U.S. households while Gan et al. (2015) and Heimer et al. (2019) examine the role of subjective survival beliefs in wealth accumulation.

Two related papers also incorporate subjective survival expectations into life-cycle models to study pension reform (de Bresser, 2024) and Social Security claiming (Dai et al., 2024). de Bresser (2024) finds that subjective survival expectations are important to generate the small response of retirement ages to changes in the pension system in the Netherlands. Additionally, Dai et al. (2024) argue that the United States data does not feature large enough variation in subjective survival beliefs to explain the prevalence of age 62 claiming. Our contribution relative to this work is to highlight an additional source of subjective beliefs—subjective beliefs about the early-claiming adjustment—that dampens policy responsiveness and helps explain observed claiming patterns. While subjective survival expectations can be observationally equivalent to preferences within sample while implying different counterfactual responses to policy change (de Bresser, 2024), program-rule beliefs contain independent identifying content on the claiming margin: imposing full information about the early-claiming adjustment cannot be absorbed by re-optimizing preferences without generating deviations in the claiming-age distribution.

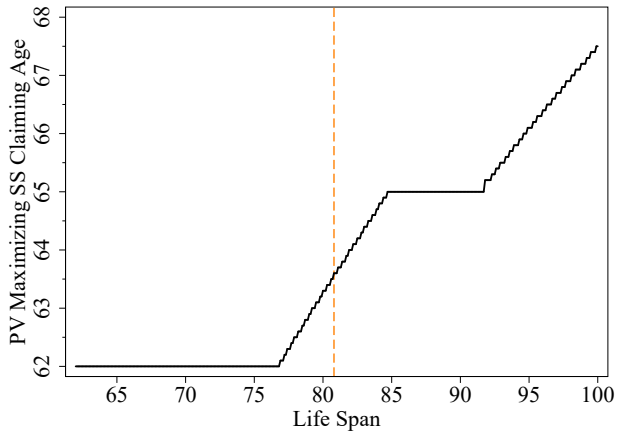
[1] Projections from the 1996 Social Security Trustees Report indicated that the system would remain fully solvent for over three more decades, with no legislated or expected benefit cuts in the near term. As shown in Table II.F18 of the report, cost and income rates remained closely aligned well into the future, leaving those retiring in the late 1990s with no compelling financial reason to claim early out of solvency concerns (SSA Trustees Report, 1996).

2 Motivation

The U.S. Social Security system pays retirement benefits from the claiming date until death. Benefits are a progressive function of average indexed monthly earnings (AIME), with replacement rates declining at higher income levels. Individuals can claim reduced benefits at age 62 or claim full benefits at age 65 (the normal retirement age for the 1931–1935 cohort). Claiming before age 65 permanently reduces monthly benefits by 6.7% per year early, while delaying up to age 70 increases them by 5.5% per year. The program therefore generates a fundamental trade-off between the duration and the level of monthly benefits.

2.1 Present discounted value analysis

Figure 1: PDV Maximizing Claiming Age



Notes: Figure shows the claiming age that maximizes the present discounted value (PDV) of total benefits received by lifespan. Calculations assume an individual with $30,000 average lifetime earnings (the average Social Security wealth for the 1931–1935 cohort) and a 3% annual discount rate. Early claims (before age 65) incur a 6.7% annual penalty; delayed claims (ages 66–70) earn a 5.5% credit. The dashed line indicates the average life expectancy at age 62: 80.8 years.

Because monthly benefits increase monotonically with the claiming age, for any given lifespan there is a unique claiming age between 62 and 70 that maximizes the PDV of lifetime benefits. Figure 1 shows this relationship. As expected, the PDV-maximizing claiming age is weakly increasing in lifespan: longer-lived individuals benefit more from delaying. For the cohort-average life expectancy, the PDV-maximizing claiming age is 63.4 years.[2]

[2]This PDV-maximizing age is sensitive to the rate at which future benefits are discounted. We use a 3% annual discount rate. If we discount at 2% annually, the PDV-maximizing claiming age rises to 64.2 years; if we discount at 4% annually, the PDV-maximizing claiming age falls to 62.8 years.

If life expectancy varies across individuals, PDV-maximizing claiming ages should vary as well. Empirical studies document substantial differences in life expectancy by education and marital status (Pijoan-Mas and Ríos-Rull, 2014). As shown in Table A.1, conditional on survival to age 62, college-educated men live, on average, 4.3 years longer than non-college-educated men, and married men live 4.4 years longer than singles. Based on the incentives in Figure 1, these longevity gaps would imply sizable differences in optimal claiming ages across groups. However, the data reveal remarkably small claiming-age gaps. Figure 2 shows the empirical distributions of claiming ages by education and marital status. The distributions are strikingly similar: the gap in average claiming age between college and non-college men is 0.6 years, and the gap between married men and singles is only 0.1 years. This homogeneity in claiming behavior has important consequences for the total benefits received. On average, individuals in the 1931–1935 birth cohort forewent over $7,000 of benefits by deviating from the PDV-maximizing claiming age, with married and college-educated workers foregoing nearly $8,000 and over $12,000 of benefits, respectively.[3]

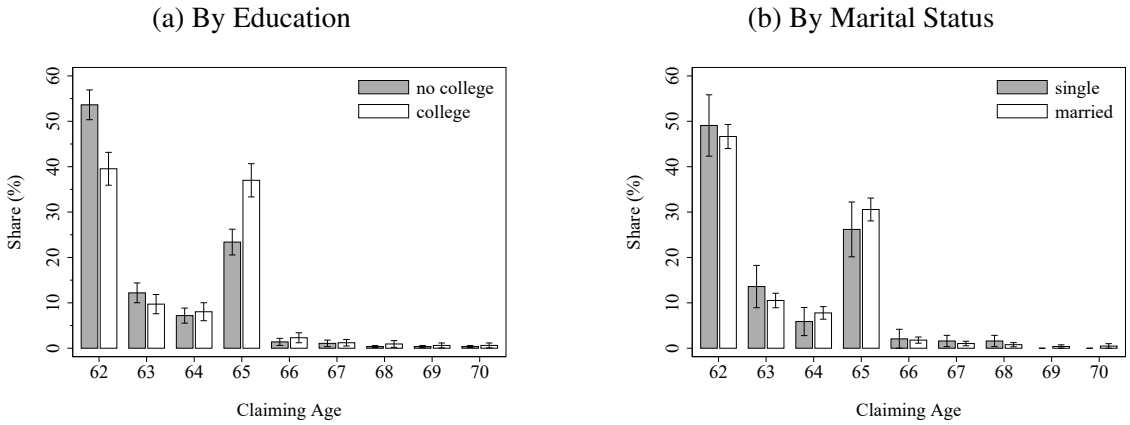
Table 1: Life Expectancy and Social Security Claiming Gaps

	Average	Education			Marital Status		
		Non-college	College	Gap	Single	Married	Gap
Panel A. Longevity							
Age 62 life expectancy	80.8	78.3	82.6	(4.3)	77.8	82.2	(4.4)
Panel B. Claiming Behavior							
PDV-maximizing age	63.6	62.6	64.2	(1.6)	62.4	64.1	(1.7)
Average age (data)	63.4	63.1	63.7	(0.6)	63.3	63.4	(0.1)
Panel C. Forgone Benefits (1998 $)							
PDV of forgone benefits	3,526	2,619	4,742		3,165	3,883	
Total forgone benefits	7,046	3,058	12,140		3,362	7,647	

Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records; authors' calculations.
Notes: PDV-maximizing claiming age assumes an individual with $30,000 average lifetime earnings and a 3% annual discount rate. Early claims (before age 65) incur a 6.7% annual penalty; delayed claims (ages 66–70) earn a 5.5% credit. Average forgone benefits are measured relative to the PDV-maximizing claiming age using the claiming shares shown in Figure 2.

[3]These calculations hold average earnings constant at $30,000 for all groups. Data show modest variation by education. For our cohort of interest, average indexed annual earnings at age 62 were roughly $25,000 for non-college educated workers and around $33,000 for college educated workers. Accounting for these differences implies total forgone benefits of $2,679 ($2,295 in PDV) for non-college workers and $13,043 ($5,094 in PDV) for college workers. These numbers are in line with Choukhmane et al. (2025) who find that couples forgo roughly $14,000 in wealth at retirement ($3,778 in forgone dollars) by making sub-optimal choices that do not maximize the match of employer-sponsored retirement plans. Refer to Appendix A for more details on forgone dollar calculations.

Figure 2: Social Security claiming behavior education and marital status



Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.
Notes: Bars show the share of men born 1926–1939 who claimed Social Security benefits at each age from 62 to 70, by education (Panel (a)) and marital status (Panel (b)). Grey bars show distributions for non-college or single men; white bars for college-educated or married men. 95% confidence intervals shown.

Married households face additional claiming incentives through spousal and survivor benefits.[4] Spousal benefits (providing up to 50% of the primary earner's benefit) can only be accessed after the primary beneficiary claims his own benefits and do not accrue delayed retirement credits. This structure weakens the incentive to delay claiming, particularly when spouses are of similar age. In contrast, survivors benefits continue after the death of the primary earner, effectively extending the payout horizon beyond the primary earner's lifetime and increasing the PDV-maximizing claiming age. The incentive to delay is particularly strong in couples where the spouse is expected to outlive her husband by many years (for example, if she is younger). These opposing forces have different effects across households: if the spouse qualifies for her own benefits, these incentives vanish entirely. We explore these dynamics in detail in our quantitative analysis.

While the PDV framework clarifies the incentives embedded in the policy, claiming decisions are made in a richer setting. Labor market shocks, health and longevity risk, household insurance within marriage, and precautionary saving motives all interact to shape the joint retirement–saving–claiming problem (Chan and Huff Stevens, 1999, 2001, 2004; French, 2005; Jones and Li, 2023; Imrohoroğlu and Kitao, 2012; Brown et al., 2010; Coile et al., 2002; Sun and Webb, 2011; Shepard, 2011; Pashchenko and Porapakkarm, 2024). In the remainder of the paper, we study these forces jointly using a quantitative life-cycle model in which households choose labor supply, saving, and Social Security claiming in the presence of earnings, health, and mortality risk. Our central departure from previous literature is to incorporate two empirically supported belief channels—subjective survival expectations and subjective knowledge about program rules—and to quantify

[4] Appendix Figure A.1 in Appendix A shows the PDV-maximizing claiming age by lifespan for couples.

how these beliefs shape claiming behavior and contribute to the homogeneity in claiming.

2.2 Empirical analysis of subjective beliefs

Before turning to the model, we document the prevalence of these beliefs in the data and provides reduced-form evidence linking them to claiming behavior. We focus on two groups of subjective beliefs: subjective expectations about survival and subjective knowledge about how Social Security benefits adjusted when claiming prior to the normal retirement age.

Descriptive Statistics We first use the Health and Retirement Study (HRS) to document patterns in subjective survival beliefs, and we then document the prevalence of subjective beliefs about Social Security program rules using both the HRS and the Understanding America Study (UAS).

We define subjective survival beliefs in the HRS as the deviation between an individual's reported probability of surviving to age 75 (conditional on survival to age 60) and the corresponding objective probability for their education–marital status group.[5] We classify deviations of subjective from group-specific objective survival (in percentage points, pp) into five bins: (1) < -25 (subjective survival more than 25 pp below objective); (2) $[-25, -5)$ (subjective 5–25 pp below objective); (3) $[-5, 5]$ (subjective within ± 5 pp of objective); (4) $(5, 25]$ (subjective 5–25 pp above objective). Similar to the literature on subjective survival beliefs (Hurd and McGarry, 1995, 2002; Hurd et al., 2004; Dai et al., 2024; O'Dea and Sturrock, 2023; de Bresser, 2024), we find that although the average deviation is small, fewer than 30% of individuals report survival probabilities within 5 pp of the objective benchmark. Deviations are common in both directions: over 30% of respondents deviate by more than 25 pp (either positively or negatively).[6]

Next, we investigate subjective beliefs about program rules. First, the HRS asks respondents about the nominal benefit amounts they would expect to receive if they claimed at various ages. By combining responses for various ages, we construct an implied perceived penalty—a measure of how much the respondent believes benefits change when claiming is delayed by one year. While this measure allows for unconstrained answers on perceived penalty, it is not possible to determine whether reported benefit differences reflect knowledge of the statutory adjustment or other factors that may influence stated nominal benefit amounts, such as expectations about earnings, labor supply, or inflation (Rohwedder and Kleinjans, 2006). We therefore view this measure as a noisy proxy for penalty beliefs. Using this HRS-based measure, we classify beliefs into five categories: (1) $[-1.5, 1.5)$ (no penalty or a small credit for claiming early), (2) $[1.5, 4.7)$, (3) $[4.7, 8.7]$ accurate

[5] This measure captures the extent, though not the source, of belief deviation. For instance, these deviations could reflect private information, omitted determinants of mortality, or true misperceptions. We do not take a stand on the source of the deviations; instead, we study how their presence affects behavior.

[6] The distribution over survival beliefs is shown in Appendix Figure B.1.

(within ± 2 pp), (4) $(8.7, 12.0]$, and (5) $\geq 12.0\%$ (roughly 5 pp or more above the true penalty).

Second, we use the UAS, which asks respondents about the expected change in benefits for a hypothetical retired individual who delays benefit claiming from age 63 to age 64.[7] Because both ages are below the normal retirement age and the question explicitly states that the individual is already retired, this setting isolates expectations about the early-claiming adjustment more cleanly. This question therefore provides our preferred measure for the quantitative analysis. However, respondents choose among four discrete options (approximately 0%, 1%, 6%, or 15%), which constrains responses to this menu.

Figure 3 reports the distribution of penalty beliefs in the HRS (Panel 3a) and the UAS (Panel 3b). Despite differences in how the question is posed, the two sources paint a consistent picture: roughly 40% of respondents report beliefs close to the statutory adjustment, while about 60% report beliefs that deviate from it, with variation in both direction and magnitude. It is important to highlight that both datasets have limited availability: HRS began asking benefit expectation questions in 2002 while UAS covers 2017–2022. Since most respondents likely had access to the Social Security Statement (a document distributed since 1995 to inform individuals about benefit rules), our estimated share of mistaken beliefs about the penalty may represent a lower bound for older cohorts who did not have access to these statements.

An additional advantage of the UAS is that it includes questions designed to measure knowledge and confidence about Social Security rules. We use a question in which respondents report how knowledgeable they feel about the relationship between claiming age and benefits. The second bar in Figure 3b reports the share who describe themselves as very knowledgeable. Only about 25% report being very knowledgeable, with the highest subjective knowledge among those who are correct about the early-claiming penalty. This suggests that many respondents may be uncertain about their own answers.

Impact on Behavior Having established the prevalence of subjective beliefs, we next exploit the panel structure of the HRS to examine the empirical relationship between beliefs and Social Security claiming behavior. We focus on the HRS for this analysis because the UAS does not provide a long enough panel to link current beliefs to subsequent claiming. We estimate:

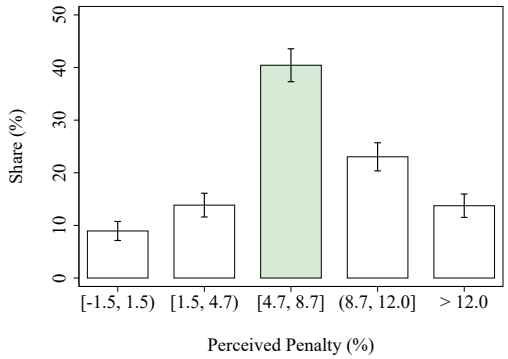
$$SS_i^{age} = \alpha + x_i'\beta_1 + \beta_2 S_i + \beta_3 K_i + \varepsilon_i \tag{1}$$

where the dependent variable SS_i^{age} is respondent i's Social Security claiming age. The regressors include cohort and demographic controls x_i', subjective survival beliefs S_i, and subjective beliefs about the early-claiming penalty K_i. We consider two specifications: (i) treating both beliefs as

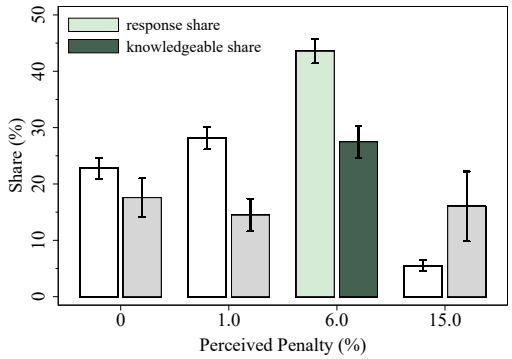
[7] Full survey text is in Appendix B.

Figure 3: Distribution of Penalty Beliefs

(a) Health and Retirement Study



(b) Understanding America Study

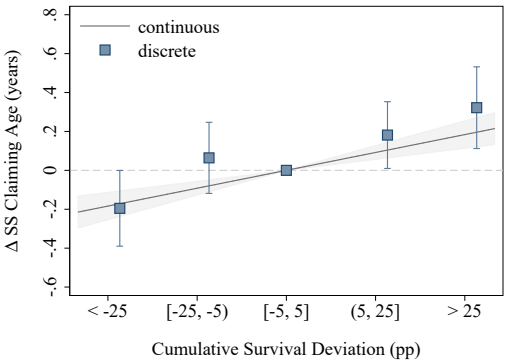


Source: Health and Retirement Study (2016, v1); Understanding America Study (waves 94, 241)
Notes: Figures show the distribution of early claim penalty beliefs in Health and Retirement Study (Panel (a)) and the Understanding America Study (Panel (b)). Panel (b) also includes the share of workers who respond they are very knowledgeable about the relationship between claiming age and benefits.

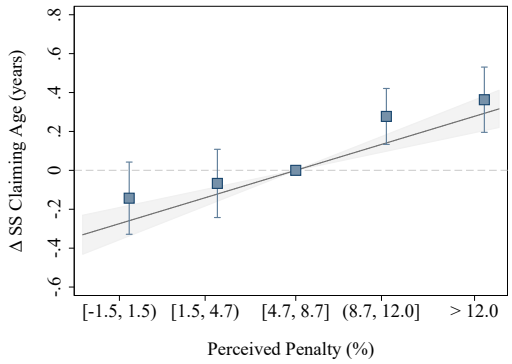
continuous variables and (ii) using the categorical belief measures defined above.

Figure 4: Impact of subjective beliefs on claiming age
regression coefficients

(a) Survival beliefs (β_2)



(b) Program knowledge beliefs (β_3)



Source: Health and Retirement Study (2016, v1), authors' calculations.
Notes: Figure shows regression coefficients from Equation 1 under two specifications. The first specifies beliefs as a continuous variable while the second defines beliefs as a categorical variable. Both specification includes cohort and demographic controls (gender, education, marital status, normal retirement age, cohort). Panel (a) reports coefficients for survival beliefs; Panel (b) for program knowledge beliefs. All values are reported relative to accurate beliefs. 90% confidence intervals shown.

Figure 4 reports the coefficient estimates (see Appendix B for details). In both specifications, coefficients are reported relative to individuals with accurate beliefs and capture how belief deviations correlate with claiming ages. The results align with the PDV framework. Underestimating

survival probability and expecting a low early-claiming penalty are both associated with earlier claiming, while overestimating survival and expecting a high early-claiming penalty are associated with later claiming. Specifically, respondents with the largest negative deviations in subjective survival beliefs, or who perceive the early-claiming penalty as minimal, claim about 0.2 years earlier than those with accurate beliefs. Conversely, respondents whose subjective survival exceeds their group objective benchmark by more than 25 pp, or who expect the largest penalties, claim 0.2–0.4 years later than those with accurate beliefs. In Appendix B we show that these results are robust to alternative categorizations of beliefs.[8]

While these regression results illustrate an association between subjective beliefs and claiming, they do not provide a full quantitative account of the role of beliefs and uncertainty in claiming. First, because the HRS question is in nominal benefit levels, the implied penalty may partly reflect expectations unrelated to policy knowledge; such measurement error would attenuate reduced-form relationships. Second, HRS respondents are on average 59.4 years old when answering these questions, so we observe beliefs immediately prior to retirement rather than over the full life cycle. Third, because beliefs may shape wealth and labor supply before claiming, specifications that control for these variables can be difficult to interpret; we therefore use a baseline specification without such controls. Finally, the UAS evidence indicates substantial uncertainty about program knowledge; if respondents are uncertain, the behavioral response to stated beliefs may be muted in reduced-form regressions. To quantify the role of subjective beliefs and uncertainty in explaining claiming-age gaps, we turn to a structural model that is disciplined by survival beliefs from the HRS and program-knowledge beliefs and confidence from the UAS.

3 Quantitative model

This section presents a structural life-cycle model of retirement and Social Security. Labor supply (h_t), consumption (c_t), savings (a_{t+1}), and Social Security benefit application (b_t^{ss}) of a male household head is modeled. Individuals make these decisions in every time period t and adjust their behavior in response to uncertainty pertaining to employment, wages, health status, subjective survival, and subjective program knowledge.

Individuals of ages $t = 25, 26, ..., 99$ differ along both fixed characteristics and state variables which vary over the life cycle. Fixed characteristics include education (e), marital status (q), and subjective SS penalty belief type (k). Marriage is summarized by a pair $q = (m, \iota)$ where m is an indicator variable taking a value of 1 if married and, if married, ι denotes the age gap

[8] Hurd et al. (2004) also uses the HRS to study the impact of subjective survival on claiming and finds that reporting a very low probability of surviving to age 85 is associated with higher rates of claiming at age 62. We use a related survival question (to age 75) and focus on the overall claiming age rather than only age-62 claiming.

between spouses. Time varying states include stochastic labor productivity (η_t); employment status (λ_t); health status (μ_t); liquid assets (a_t); Social Security wealth of the head of household (a_t^{ss}); if applicable, Social Security wealth of a spouse (a_{st}^{ss}); and household benefit application status (b_{t-1}^{ss}). Given this vector of states $\left(e, q, k, \eta_t, \lambda_t, \mu_t, a_t, a_t^{ss}, a_{st}^{ss}, b_{t-1}^{ss}\right)$, individuals choose optimal consumption, labor supply and make Social Security benefit application decisions (if eligible) to maximize the present discounted value of lifetime utility.

3.1 Preferences

In each period t, agents derive utility from consumption c_t and leisure l_t, with non-separable preferences over the two given by:

$$U^{e,m}(c_t, l_t) = \frac{1}{1-\rho^{e,m}} \left(\left(\frac{c_t}{\zeta_t^{e,m}} \right)^{\nu^{e,m}} l_t^{1-\nu^{e,m}} \right)^{1-\rho^{e,m}}$$

where $\rho^{e,m}$ is the coefficient of relative risk aversion, $\nu^{e,m}$ is the utility weight on consumption, and $\zeta_t^{e,m}$ is the equivalent scale in consumption. All these parameters are permitted to vary by both education (e) and marital status (m). The total amount of leisure in period t is given by:

$$l_t = \bar{l}^{e,m} - h_t - \phi_P^{e,m}(t)\mathbb{I}\{h_t > 0\} - \phi_H^{e,m}(\mu_t, t) \tag{2}$$

Here, $\bar{l}^{e,m}$ is the total time endowment, h_t is hours worked, $\phi_P^{e,m}(t)$ is the participation cost incurred when $h_t > 0$, and $\phi_H^{e,m}(\mu_t, t)$ is the time cost of poor health. The participation cost is:

$$\phi_P^{e,m} = \frac{\exp(\phi_0^{e,m} + \phi_1^{e,m} t + \phi_2^{e,m} t^2)}{1 + \exp(\phi_0^{e,m} + \phi_1^{e,m} t + \phi_2^{e,m} t^2)} \times \bar{l}^{e,m} \tag{3}$$

We adopt the functional form for health-related time cost from Jones and Li (2023):[9]

$$\phi_H^{e,m}(\mu_t, t) = \begin{cases} 0 & \text{if } \mu_t = 0 \\ \{\phi_a^e + \phi_b^e \log(67 - t)\} \times \bar{l}^{e,m} & \text{if } \mu_t = 1 \quad t \leq 65 \\ \phi_a^e \times \bar{l}^{e,m} & \text{if } \mu_t = 1 \quad t > 65 \end{cases} \tag{4}$$

Upon dying, individuals value bequests of any leftover bequeathable wealth, A_t^q, according to the utility function developed by De Nardi (2004):

[9]The "good health" state in Jones and Li (2023) corresponds to our good health category, while their fair/poor health state maps to our bad health category. We take their estimates of education-specific parameters ϕ_a^e and ϕ_b^e, which imply that entering bad health at age 25 reduces the effective time endowment by 14.2% for non-college men and 38.7% for college graduates.

$$beq^{e,q}(A_t^q) = \frac{\theta_{beq}^{e,m}}{1-\rho^{e,m}} \left(A_t^q + \kappa_{beq}^{e,m}\right)^{(1-\rho^{e,m})\nu^{e,m}}$$

Bequeathable wealth, A_t^q, includes remaining assets, a_t, and Social Security survivors benefits, if eligible. Eligibility depends on marital status and spousal age gap q (see Section 3.4.3). The strength of the bequest motive is governed by $\theta_{beq}^{e,m}$, while $\kappa_{beq}^{e,m}$ captures the curvature of the bequest function. Higher values of $\theta_{beq}^{e,m}$ increase the marginal utility of bequests and higher $\kappa_{beq}^{e,m}$ implies bequests are more luxury-like. Both parameters vary by education e and marital status m.

3.2 Health and mortality

Individuals face an exogenous health shock every period, μ_t, that varies by education and marital status. Health influences survival beliefs for the next period, $\pi_t^s(e, m, \mu_t)$ and the total time endowment. In Appendix I.2, we extend the model to allow for medical expenditures.

3.3 Employment and wages

Individuals face unemployment shocks with probability π^λ. These shocks reduce labor productivity and introduce wage-scarring effects. The hourly wage in each period depends on an education- marriage- and age-specific component $\omega(e, m, t)$, current unemployment status λ_t, and an auto-regressive idiosyncratic component η_t:

$$w_t = \xi(\lambda_t)\exp(\omega(e, m, t) + \eta_t);\ \ \eta_t = \rho^w \eta_{t-1} + \epsilon_t^w;\ \ \epsilon_t^w \sim N(0, \sigma_{\epsilon^w}^2) \tag{5}$$

Individual receiving an unemployment shock may re-enter the labor market immediately but incur wage penalty $\xi(\lambda_t)$. This allows the model to capture wage scarring effects of unemployment.

3.4 Social Security

Social Security benefits are computed in several steps. First, the highest thirty-five earning years are averaged to compute an annualized version of the worker's Average Indexed Monthly Earnings (AIME) or Social Security wealth. Let a_t^{ss} be Social Security wealth. The evolution of Social Security wealth is approximated in the model in the following simple way:

$$a_{t+1}^{ss} = \min\{[a_t^{ss} + \max\{0, (w_t h_t - a_t^{ss})/35\}], a^{\max}\} \tag{6}$$

where $a^{\max}$ is the threshold at which the Social Security wealth is capped and $w_t h_t$ denotes annual labor earnings for period t. Note that in Equation 6, we assume that the high earnings year only

replaces an average earnings year, as modeling the actual system would require keeping track of entire earnings history which is computationally infeasible.

Second, our Social Security wealth measure is converted into the Primary Insurance Amount (PIA) using the following piece-wise linear function:

$$\begin{aligned} pia(a_t^{ss}) = 0.90 \times \min\{a_t^{ss}, b_0\} + 0.32 \times \min\{\max\{a_t^{ss} - b_0, 0\}, b_1 - b_0\} \\ + 0.15 \times \max\{a_t^{ss} - b_1, 0\} \end{aligned} \tag{7}$$

where the bend points of this function, b_0 and b_1, are set based upon Social Security policy.

3.4.1 Adjustments

Social Security benefits, ssb_t, are a function of the primary insurance amount, $pia(a_t^{ss})$, and two possible adjustments: a penalty/credit for early/late benefit claims $\Gamma(t^{ss})$ and a temporary decrease in benefits for those workers who continue working while also claiming benefits (Υ_t).

$$ssb_t = pia(a_t^{ss}) * \Gamma\left(t^{ss}\right) - \Upsilon_t\left(pia(a_t^{ss}), t^{ss}, w_t h_t\right) \tag{8}$$

Early/late claiming penalty

SS benefits can be claimed without penalty at the normal retirement age (t_{NRA}), but individuals may claim earlier—from the early retirement age (t_{ERA}) of 62—at a permanent reduction in benefits. For each year claimed before t_{NRA}, the benefit amount is reduced by the early claiming penalty; for each year delayed beyond t_{NRA}, it is increased by the delayed claiming credit. These adjustments are captured by the percentage change γ_1^{ss} for claims before the NRA and γ_2^{ss} for claims after the NRA. The total penalty or credit, the share of the primary insurance amount individuals are eligible to receive as benefits, for a claiming age of t^{ss} is:

$$\Gamma\left(t^{ss}\right) = \begin{cases} 1 - \left(t_{NRA} - t^{ss}\right)\gamma_1^{ss} & \text{if } t^{ss} < t_{NRA} \\ 1 & \text{if } t^{ss} = t_{NRA} \\ 1 + \left(t^{ss} - t_{NRA}\right)\gamma_2^{ss} & \text{if } t^{ss} > t_{NRA} \end{cases} \tag{9}$$

Retirement earnings test

The Social Security retirement earnings test taxes the labor income (above a certain threshold y_t^{ss}) of the Social Security beneficiaries at a rate τ_t^{ss}, until the normal retirement age. Following the abolition of the earnings test in 2000 for beneficiaries above the NRA, we apply the earnings test only before the NRA. For each additional dollar earned above the threshold, SS benefits are

reduced by τ_t^{ss} until all benefits are taxed away:

$$\Upsilon_t\left(pia(a_t^{ss}), t^{ss}, w_t h_t\right) = \min\{pia(a_t^{ss}) * \Gamma\left(t^{ss}\right), \max\{0, w_t h_t - y_t^{et}\}\tau_t^{et}\}$$

Υ_t denotes benefits lost through the earnings test. However, the retirement earnings test is not a true tax as withheld benefits are credited back through permanent increases in future benefits. We implement this in the model through increases in the Social Security wealth.[10]

3.4.2 Subjective program beliefs

We model program misinformation by including subjective and uncertain beliefs about the annual adjustment factor for early claiming, $\widetilde{\gamma}_1^{ss}$. Individuals have a fixed, observed belief type, denoted by k, meant to represent an agent's best guess of the penalty. However, we allow uncertainty in subjective penalty beliefs; subjective penalties are distributed according to a belief distribution $B^{e,m}\left(k\right)$ that is permitted to vary by education, marital status, and belief type.

$$\widetilde{\gamma}_1^{ss} \sim B^{e,m}\left(k\right) \tag{10}$$

The subjective belief over the annual penalty, then, will map into a subjective total adjustment, $\widetilde{\Gamma}\left(t^{ss}\right)$ defined below.

$$\widetilde{\Gamma}\left(t^{ss}\right) = \begin{cases} 1 - \left(t_{NRA} - t^{ss}\right) \cdot \widetilde{\gamma}_1^{ss} & \text{if } t^{ss} < t_{NRA} \\ \Gamma\left(t^{ss}\right) & \text{if } t^{ss} \geq t_{NRA} \end{cases} \tag{11}$$

Equation 11 implies that this subjective adjustment deviates from the objective adjustment only prior to the normal retirement age. In other words, individuals have accurate information about delayed retirement credits.[11]

Perceived Social Security benefits (given by Equation 12) are a function of the primary insurance amount for an individual of a given SS wealth, $pia\left(a_t^{ss}\right)$, the uncertain total adjustment, $\widetilde{\Gamma}\left(t^{ss}\right)$, and the retirement earnings test, Υ_t.[12]

$$\widetilde{ssb}_t = pia\left(a_t^{ss}\right) \cdot \widetilde{\Gamma}\left(t^{ss}\right) - \Upsilon_t\left(pia(a_t^{ss}), t^{ss}, w_t h_t\right) \tag{12}$$

This way of modeling subjective beliefs is akin to agents operating under limited information with their actions resulting in irreversible mistakes in terms of claiming.

[10] Note that in practice, benefits withheld due to the earnings test are credited back at the normal retirement age. For computational tractability, we assume these are credited immediately.

[11] We make this assumption due to limitations on survey evidence on beliefs about these delayed retirement credits.

[12] In Appendix I.4, we consider an extension which relaxes the assumption of full information related to RET.

3.4.3 Marriage related benefits

Spousal benefits We assume that married couples claim Social Security benefits jointly. If both spouses are eligible at the time of the household head's claim, both the head's benefits and spousal benefits begin immediately. If the younger spouse is not yet eligible (i.e., under age 62) when the head claims, spousal benefits are received once she reaches age 62.[13] We also assume that subjective beliefs about the early claims penalty apply to spousal benefits.

Spousal Social Security wealth, a_{st}^{ss}, and primary insurance amount (PIA), $pia\left(a_{st}^{ss}\right)$, evolve based on exogenous spousal income, y_{st}, following the rules defined in Equations 6 and 7. The resulting spousal PIA is then used to determine Social Security benefits, as specified in Equation 13. If the benefits a spouse qualifies for based on her own earnings record exceed 50 percent of her husband's PIA, she receives those primary benefits. Otherwise, if her own earnings entitle her to no or relatively small benefits, she receives dual or auxiliary spousal benefits equal to up to 50 percent of the household head's PIA. Perceived spousal benefits are computed as:

$$\widetilde{ssb}_t^s = \max\left\{pia\left(a_{st}^{ss}\right)\cdot\widetilde{\Gamma}^s, 0.5\cdot pia_t\cdot\min\{\widetilde{\Gamma}^s, 1\}\right\} \tag{13}$$

We include the subjective adjustment for spousal claiming age in these calculations, $\widetilde{\Gamma}_{st}$. We assume that spouses retire by the point of claiming, and, thus, spousal benefits are not adjusted for the retirement earnings test. The adjustment for claiming age follows the formula shown in Equation 9 as a function of the spousal claiming age, $\max\{t^{ss}-\iota, 62\}$, with one adjustment. Policy rules impose that spousal benefits do not accumulate delayed retirement credits and the claiming age adjustment cannot be above 1.

Survivors benefits Married individuals may also leave their Social Security benefits—net of any spousal benefits already being received—to their surviving spouses upon death.[14] We model these survivors benefits as part of the bequeathable wealth, A_t^q, which is defined as follows:

$$A_t^q = \begin{cases} a_t + \sum_{j=t-\iota}^{T} \frac{\pi_j^s}{(1+r)^{j-(t-\iota)}} \max\left\{\widetilde{ssb}_t - \widetilde{ssb}_t^s, 0\right\} & \text{if } m = \text{married}, t-\iota \geq 62 \\ a_t & \text{otherwise} \end{cases} \tag{14}$$

In addition to any leftover assets, a_t, bequeathable wealth is a function of Social Security

[13] This assumption is supported by data for the 1931–1935 cohort: over 65% of eligible couples claimed within a year of each other, and over 70% of ineligible younger spouses claimed at age 62. Appendix I.3 explores an extension where spouses' claiming ages are modeled independently.

[14] In practice, after a beneficiary's death, the surviving spouse receives the larger of their own or the deceased spouse's benefits. Accordingly, we assume the value of survivors benefits generating warm-glow utility equals the difference between the two.

wealth if the individual is married and the spouse is over the age of 62.[15] These survivors benefits are calculated as the present value of the expected stream of additional benefits a spouse would receive from the time of the death of the household head until the end of her own life. This present value is a function of the household head's age t, the spousal age gap, ι, and the difference in perceived benefits between the head of household and spouse.[16] Once Social Security benefits are claimed, survivors benefits are a function of actual retirement benefits, ssb_t and ssb_t^s.

3.5 Budget constraint

Before claiming Social Security, individuals face the following budget constraint:

$$c_t + a_{t+1} = a_t + W\left(y_t, y_{st}, \bar{r}a_t, \tau\right) + \mathbb{I}\left\{b_t^{ss} = 1\right\} \times \left(\widetilde{ssb}_t + \mathbb{I}\left\{t - \iota \geq 62\right\} \times \widetilde{ssb}_t^{\,s}\right) + tr_t + pb_t \tag{15}$$

An individual's disposable household income, $W(\cdot)$, includes several components: labor income from hours worked, $y_t = w_t h_t$; exogenous spousal income, y_{st} (if married); and asset returns, ra_t. In addition, individuals may receive government transfers, tr_t (defined in Equation 17) and private pension benefits, pb_t, as well as pay taxes, τ. Importantly, before claiming, individuals base their decisions on perceived Social Security benefits, $\widetilde{ssb}_t$, and the perceived future spousal benefits (conditional on spousal eligibility), $\widetilde{ssb}_t^{\,s}$. Once Social Security benefits are claimed (i.e., when $b_t^{ss} = 1$), actual retirement benefits, ssb_t and ssb_t^s are revealed.

There is a standard borrowing constraint on assets and a consumption floor which guarantees a minimum level of consumption (Hubbard et al., 1995). given by:

$$a_{t+1} \geq 0; \; c_t \geq 0 \; \forall t \tag{16}$$

Government transfers, tr_t, bridge the gap between this minimum level of consumption and individual's liquid resources. This is a simple approximation to the federal safety net programs in the U.S. such as Supplemental Nutritional Assistance Program (SNAP), Supplemental Security Income (SSI), Temporary Assistance for Needy Families (TANF), and other programs.

$$tr_t = \min\{0, \underline{c} - (a_t + W_t + \mathbb{I}\left\{b_t^{ss} = 1\right\} \times (ssb_t + \mathbb{I}\left\{t - \iota \geq 62\right\} \times ssb_t^s) + pb_t)\} \tag{17}$$

where W_t is the total disposable household income as defined in Equation 15.

[15] Under Social Security's widow benefits provisions, surviving spouses may receive reduced benefits starting at age 60 and full benefits at 65. In our model, we choose an intermediate eligibility age of 62. We exclude child benefit provisions under OASDI, as early-life beneficiary death is sufficiently rare.

[16] Further details of spousal and survivors benefit calculations are provided in Appendix C.

3.6 Recursive formulation

Individuals solve a dynamic life-cycle optimization problem, choosing consumption, hours worked, and Social Security benefit application to maximize expected discounted lifetime utility subject to exogenous processes for wage, health, employment, subjective survival, and program beliefs, budget, borrowing, and time constraints, transfer rules, taxes, and Social Security policy. The life-cycle is separated into three phases: (1) the employment phase (ages 25–61) during which individuals choose consumption, savings, and labor supply.[17], (2) the retirement choice phase (ages 62–69) where individuals continue to make consumption, labor supply, and savings decisions but also decide whether to claim Social Security benefits, and (3) the retirement phase (ages 70-99) in which individuals make only consumption and saving decisions.

We present below a general value function for a household head with education level e, marital status m, and Social Security fixed belief type k. Time varying states include stochastic labor productivity (η_t); employment status (λ_t); health status (μ_t); liquid assets (a_t); Social Security wealth of the head of household (a_t^{ss}); if applicable, Social Security wealth of a spouse (a_{st}^{ss}); and household benefit application status (b_{t-1}^{ss}). Expectations are taken over productivity (η_t), employment (λ_t), Social Security program beliefs, and health uncertainty (μ_t). Phase-specific decision problem, which incorporate life-stage-dependent constraints and transitions, are provided in Appendix D.

$$
\begin{aligned}
V_{e,q,k}(a_t, a_t^{ss}, a_{st}^{ss}, \eta_t, \lambda_t, \mu_t, b_{t-1}^{ss}) = \max_{\{c_t, h_t, b_t^{ss}\}} \Bigg\{ & U^{e,m}(c_t, l_t) \\
& + \beta^{e,m} \pi_{t+1}^{s} \Big[EV_{e,q,k}(a_{t+1}, a_{t+1}^{ss}, a_{st+1}^{ss}, \eta_{t+1}, \lambda_{t+1}, \mu_{t+1}, b_t^{ss}) \Big] \\
& + \beta^{e,m} (1 - \pi_{t+1}^{s}) beq^{e,m}(A_{t+1}^{q}) \Bigg\}
\end{aligned}
$$

$$
\text{s.t.} \quad (2), (6\text{-}13), (15), \text{and } (16).
$$

4 Estimation

We estimate the model for male household heads born between 1931 and 1935 using a two-step strategy following Gourinchas and Parker (2002).[18] In the first step, we estimate processes that can be identified independently of the structural model using data from the Panel Study of Income Dynamics (PSID), Health and Retirement Study (HRS), the Household Component of the Medical

[17] We abstract from disability insurance (DI) and focus on those who claim benefits through the non-DI pathway.

[18] We focus on this cohort for two reasons: (1) we observe most of their life cycle in the data, and (2) they retire before major policy changes to the normal retirement age and the retirement earnings test took place.

Expenditure Panel Survey (MEPS), and the Understanding America Study (UAS).[19] In the second step, we combine initial conditions for this cohort with the structural model and first-step estimates to identify preference parameters and uncertainty about subjective program beliefs.

4.1 First step

4.1.1 Health and subjective survival beliefs

Health takes two values, $\mu_t \in \{\text{good}, \text{bad}\}$, and is identified using self-reported health status from MEPS and HRS.[20] Health transitions are estimated from the MEPS Household Component using data on men ages 20–90. We estimate an ordered probit regression of self-reported health status on lagged health, education, marital status, and a cubic function of age. Estimation details are provided in Appendix F.1.

We estimate subjective survival beliefs using a three-step procedure. First, we construct objective survival probabilities by combining MEPS and HRS data: MEPS for individuals under age 50, and HRS for those 50 and older. We rely on HRS beyond age 50 as MEPS excludes the institutionalized population, potentially biasing survival estimates at older ages. We estimate objective survival probabilities using an ordered probit model of the death indicator regressed on lagged health status interacted with a third-order polynomial in age. We also include interactions of education (college), marital status, and gender with the same age polynomial as well as birth-year cohort controls. Because these objective probabilities are central to identifying subjective survival beliefs, we scale the estimated annual survival probabilities so that life expectancies at ages 50, 60, and 70 align with Social Security Administration life tables for men in our cohort of interest.[21] The resulting objective survival probabilities by age, education, marital status, and self-reported health are depicted as dashed lines in Appendix Figure F.3.

Second, we assume that prior to age 50 objective and subjective survival probabilities are identical and estimate cumulative subjective survival curves for those age 50 and above using a method similar to O'Dea and Sturrock (2023). This method fits a Weibull distribution—a two-parameter distribution commonly used to model aging—to subjective reports of the probability of survival to given ages. While O'Dea and Sturrock (2023) estimates individual-level survival curves in the context of the United Kingdom, this level of granularity is infeasible in our structural model. Rather,

[19] Details on data source, samples are in Appendix E; details on the first stage estimation are in Appendix F.

[20] Both surveys ask respondents to rate their health on a 1–5 scale, from "Excellent" (1) to "Poor" (5). For computational simplicity, we collapse this into a binary measure: responses of "Excellent," "Very Good," or "Good" are classified as good health; "Fair" or "Poor" as poor health. Alternatively we could model an objective measure, such as frailty index in Hosseini et al. (2022). However, as Miller and Bairoliya (2021) discusses, self-rated health remains predictive of mortality even after controlling for objective measures.

[21] This scaling is modest: survival probabilities are multiplied by 1.002 between ages 50–59, 1.0035 between 60–69, and 1.006 for ages 70 and older. These comparisons are reported in Table F.2.

we extend their methodology to estimate education, marital status, and self-reported health specific curves using median subjective survival probability reports for each group from the HRS.

In the final step, we use the estimated subjective cumulative survival curves and objective survival probabilities to compute an age-dependent scaling factor for each education–marital status–health group. This factor is such that when it is multiplied by the objective annual survival probabilities the resulting cumulative objective survival curve exactly matches the estimated subjective survival curve. The scaled subjective survival probabilities used in our quantitative model are presented as colored, marked lines in Appendix Figure F.3. Additional methodological details are provided in Appendix F.2.

Table 2 shows how incorporating subjective survival probabilities alters implied remaining life expectancy at age 60 for individuals in both bad and good health.[22] On average, singles report subjective life expectancy that exceeds the objective levels, whereas married individuals report subjective life expectancy below the objective benchmark. A similar pattern emerges by education: non-college men are close to objective longevity on average, while college-educated men systematically underestimate their survival prospects. These average patterns are driven largely by individuals in good health, who substantially underestimate their objective probability of survival.

Table 2: Age 60 Life Expectancy: Objective and Subjective

	All		Bad Health		Good health	
	Objective	Subjective	Objective	Subjective	Objective	Subjective
Non-college singles	15.4	17.5	11.3	16.6	23.2	19.3
College singles	19.1	19.8	12.0	17.2	24.3	20.9
Non-college married	19.2	17.9	14.2	16.4	26.4	19.4
College married	23.3	20.0	14.6	16.2	27.2	21.2
By Education:						
Non-college	17.6	17.7	13.0	16.5	25.1	19.4
College	22.3	20.0	14.0	16.4	26.5	21.1
By Marital Status:						
Single	17.1	18.5	11.6	16.9	23.7	20.0
Married	21.8	19.2	14.5	16.3	26.9	20.6
Overall:	20.3	19.0	13.5	16.5	25.9	20.4

Notes: objective and subjective life expectancy at age 60 from the HRS are reported for men born between 1931 and 1935, both for the overall population in this birth cohort and conditional on health status at the given ages. "Good health" refers to the top three categories of self-reported health, while "bad health" corresponds to the bottom two categories.

[22] Comparable results at ages 50 and 70 are reported in Appendix Table F.3.

4.1.2 Family structure

Family structure determines three inputs for married men: the consumption equivalence scale, $\zeta_t^{e,m}$, the spousal age gap, ι, and spousal income, y_{st}.

The equivalence scale which adjusts household consumption is based on family size and varies by education and marital status. For single households, this parameter is fixed at 1. For married households, the scale reflects the presence of a spouse and the average number of children living in the household. We estimate the number of children age 0-17 by age and education using data from the PSID. We then map these family sizes to values of $\zeta_t^{e,m}$ using the OECD equivalence scale.[23]

The spousal age gap for married men is modeled using the empirical distribution of age differences between husbands and wives observed in the PSID for the relevant cohort. We discretize this distribution using four representative values: 0, 1, 4, and 8 years. The mass at each point is assigned based on observed frequencies: 8.7% of couples have no age gap, 26.2% have a one-year gap, 46.1% have a four-year gap, and 19% have an eight-year gap. Additional estimation details related to children and spousal age gaps are provided in Appendix F.3

Spousal income, y_{st}, is estimated using the restricted Social Security Administration Summary Earnings file linked to the HRS. It is modeled as a function of the household head's state variables: age, education level, spousal age gap, and labor income. To allow for zero spousal earnings, we adopt a two-step approach. First, we estimate spousal labor force participation as a function of the household head's characteristics. Then, conditional on participation, spousal earnings are predicted using the same set of state variables. For both of these steps an indicator of participation or real spousal labor earnings in thousands of 1998 dollars are regressed on a fourth-order age polynomial, the labor income of the head interacted with this fourth-order age polynomial, and education and spousal age gap interacted with age. Estimation details are provided in Appendix F.4.

4.1.3 Labor productivity, employment shocks, and wage scarring

Wages are modeled as the sum of two components: (1) a deterministic component varying by age, education, and marital status, and (2) a persistent stochastic shock. These components are estimated from the PSID in two steps. First, we estimate the deterministic wage profiles by age, education, and marital status using a Heckman two-step procedure to account for selection into employment. Second, we estimate the persistence and standard deviation of the wage shocks by fitting an AR(1) process to the residuals from the first step using a minimum distance estimator.

Individuals face an exogenous employment shock characterized by involuntary labor market separation and a wage penalty. This shock, λ, is assumed constant across education and marital groups and calibrated to match an annual separation rate of $\lambda = 0.056$. The wage penalty, ξ,

[23] The OECD scale assigns a weight of 1 to the household head, 0.5 to the spouse, and 0.3 to each child.

represents the proportional wage loss one year after displacement and is estimated from the PSID by regressing log hourly wages on indicators for years since displacement, along with standard controls, including quadratic terms in age and experience. We estimate a 14% wage loss one year after displacement ($\xi = 0.86$).[24]

While we address the selection issue from unobserved wages among the non-employed using a Heckman selection model, the resulting profiles remain biased due to wage scarring from unemployment spells. Observed wages reflect the cumulative impact of past unemployment on current and future earnings. Since our structural model explicitly incorporates unemployment shocks and their wage-scarring effects, we correct the exogenously estimated wage profiles to avoid double-counting. We calibrate adjusted profiles so that wages simulated from the process in Equation (5)—which includes the unemployment term $\xi(\lambda)$—match those from the original unadjusted profiles using a modified process that excludes unemployment effects (i.e., setting $\xi(\lambda) = 1$). Appendix Figure F.7 shows the estimated age-wage profiles $\omega(e, m, t)$, both unadjusted and adjusted, for each education-marital group.

4.1.4 Social Security and private pensions

We study a cohort of men born between 1931 and 1935. For this group, we set the early claiming age to 62 and the normal retirement age to 65. Benefits are reduced by 6.7 percent per year for early claiming and increased by 5.5 percent per year for delays beyond the NRA, up to age 70.[25] As this cohort reaches retirement in the mid-1990s, we calibrate Social Security parameters using the 1998 rules (see Section 3.4). Maximum Social Security wealth, $a^{\max}$, is set at \$68,400. The bend points in the benefit formula, b_0 and b_1 in Equation (7), are \$5,724 and \$34,500, respectively, and do not vary with claiming age. Earnings above \$9,120 before the NRA are subject to a 50 percent tax rate; these withheld benefits are returned through future increase in benefits.

For computational tractability, and given the low participation of the 1930s birth cohort in defined contribution plans, we model private pensions to reflect the core features of defined benefit (DB) plans: benefits are illiquid before a fixed eligibility age (assumed to be 62 in the model), and non-linear pension accrual profiles generate sharp work disincentives at certain ages. To capture these features, pb_t is imputed as a function of an individual's primary insurance amount and modeled analogously to French (2005):

[24] This estimate aligns with the wage-scarring literature (Jacobson et al., 1993; Huff Stevens, 1997; Huckfeldt, 2016). Further wage estimation details and separation rate and wage penalty estimation are in Appendix F.5.

[25] The delayed retirement credit is cohort-specific and increased for later birth cohorts, implying that the return to delaying claiming past the FRA varies substantially across cohorts. Our quantitative analysis focuses on men born 1931–1935, for whom the NRA is 65 and the delayed retirement credit is approximately 5.5% per year between ages 66–70.

$$pb_t = pb(pia_t) + \epsilon_t^p$$

where $pb(pia_t)$ is a regressive function of lifetime earnings, and ϵ_t^p reflects age-specific pension accrual. To reflect the significant heterogeneity in pension wealth across education and marital groups, we adjust the scale parameter in the $pb(pia_t)$ function from French (2005) separately for each group. These adjustments ensure that predicted average pension benefits at older ages match observed values in the HRS data (see Appendix F.6 for details).

4.1.5 Taxes

Individuals in the model pay a proportional payroll tax, τ_t^{ss}, and labor income taxes, $\tau^{e,m}$. The proportional payroll tax includes both the Social Security payroll tax of 6.2 percent on income until the taxable maximum, $a^{\max}$, and the Medicare tax of 1.45 percent on total labor income.

Following the literature, we adopt a smooth functional form for the labor income tax that allows for negative tax rates in order to incorporate Earned Income Tax Credit (EITC). Total taxes paid, $T(y_t)$, on income y_t is given by:

$$T(y_t) = y_t - \lambda_t y_t^{1-\xi_t}.$$

We allow parameters of this tax function to vary by age, education, and marital status and estimate the function using PSID data. Age-dependent parameters capture changes in tax policy over our cohort's life cycle while education-specific parameters reflect differences in the number of children (refer to Appendix F.7 for estimation details).

4.2 Second step estimation

In the second step, we use initial conditions for the 1931 - 1935 cohort alongside our structural model and the parameters from the first step to estimate the preference parameter vector and belief-specific uncertainty over SS penalty $\{\beta, \rho, \nu, \theta_{beq}, \kappa_{beq}, \phi_P(t), \bar{l}, B(k)\}$ by education and marital status. We employ the Method of Simulated Moments (MSM) to find the parameter values that allow model simulated moments to best match the data profiles for this cohort of male household heads according to a GMM criterion function.[26] Additional details on the second-step estimation are discussed in Appendix G.

[26] We use the identity matrix for weighting (see Altonji and Segal, 1996).

Preferences

Table 3 reports the second-step MSM estimates of the preference parameters by education and marital status. The estimates are disciplined by the joint life-cycle profiles of assets and labor supply. Broadly, β is primarily pinned down by the rate of wealth accumulation earlier in the life cycle, whereas late-life features of the wealth profile discipline the bequest parameters $(\theta_{beq}, \kappa_{beq})$.[27] The curvature parameter ρ (CRRA) governs intertemporal smoothing and the strength of precautionary saving and is informed by how strongly households build buffer-stock wealth and smooth consumption in response to risk. Labor-supply moments identify the intratemporal trade-off between consumption and leisure through the utility weight (ν) and the overall time endowment and participation costs ($\bar{l}, \phi_P(t)$)) jointly govern the level and age profile of participation and hours.

Our preference parameter estimation uses 31 moments for each education–marital group: labor-force participation and hours worked conditional on participation in 5-year age bins between ages 25–69, and median household assets in 5-year age bins between ages 25–99.[28]

Table 3: Preference parameters

Parameter	Description	Singles		Married	
		Non-college	College	Non-college	College
β	discount factor	0.963	0.975	1.030	1.071
ρ	relative risk aversion	3.711	2.459	1.837	2.519
ν	consumption weight	0.578	0.579	0.528	0.662
θ_{beq}	bequest intensity	0.782	0.988	1.245	1.659
κ_{beq} (in 000s)	bequest curvature	217.7	106.7	47.6	67.4
$\bar{l}$	time endowment	5053	4504	4574	4354

* The estimated β's do not include discounting due to survival risk, which enters separately. The effective discount factor is $\tilde{\beta}_t = \beta\, \pi_t^s$, where π_t^s is the subjective probability of survival to $t+1$.

How beliefs interact with preference identification. The two novel belief objects in our framework—subjective survival beliefs and subjective program-rule beliefs—affect the mapping from observed behavior to estimated preferences in distinct ways. Consistent with the identification insight emphasized by de Bresser (2024), replacing subjective survival beliefs with objective survival leaves

[27] In Appendix I.1, we test the robustness of our key findings to fixing bequest motives across groups. However, substantial variation in pension wealth as well as in survival beliefs makes it infeasible to match key data moments, particularly wealth profiles, using a common discount factor.

[28] Details on the construction of these life-cycle profiles are provided in Appendix G.1.

the model approximately observationally equivalent in-sample once preferences are re-estimated: the shift in perceived longevity after age 50 primarily affects late-life saving incentives and is largely absorbed by parameters that govern the marginal value of wealth in old age, especially $(\theta_{beq}, \kappa_{beq})$. By contrast, shutting down subjective program-rule beliefs changes perceived claiming incentives and the timing of benefit flows in the household budget which in turn affects both labor supply exits as well as wealth evolution. While re-estimating preferences can preserve the fit to wealth and labor-supply profiles through adjustments in ν, $\bar{l}$, participation costs, and bequest motives, it does *not* preserve the within-sample fit of the claiming-age distribution.[29] This difference highlights that program-rule beliefs contain independent identifying content for the claiming margin that cannot be absorbed by preferences alone.

Parameter estimates. We fix the real interest rate at $r = 0.03$ for all groups. Table 3 shows that the estimated discount factor ranges from 0.963 for non-college singles to 1.071 for married college men.[30] Note that β is reported net of survival discounting: because married groups are, on average, more pessimistic about survival in our estimates, their effective discounting $\beta\pi_t^s$ is lower, and the model correspondingly estimates a higher β to match the same life-cycle asset accumulation.

The bequest parameters $\theta_{beq}^{e,m}$ and $\kappa_{beq}^{e,m}$ govern the strength and curvature of the warm-glow bequest motive. Because the implied propensity to bequeath is a nonlinear function of all preference parameters, we summarize these implications in Appendix Figure G.1a, which plots the share of resources households would leave as bequests if the probability of death next period were one, and Appendix Figure G.1b, which benchmarks our implied bequest motives against the literature. Our estimates imply bequests that are less "luxury-good" than in, e.g., Lockwood (2012); De Nardi et al. (2010, 2021), but closer to French (2005) and Ameriks et al. (2020).

Finally, the estimated time endowment $\bar{l}^{e,m}$ ranges from 4,354 hours annually (84 hours/week) for married college men to 5,053 hours (97 hours/week) for non-college singles. These estimates reflect two patterns in the data: married households face greater nonmarket time demands (e.g., child care and home production), and hours worked conditional on participation vary little across education groups despite differences in wages. A higher time endowment for less-educated men helps reconcile similar hours with greater leisure consumption among the less educated.[31] We also estimate the time cost of working via the function in equation (3); Appendix Figure G.2 plots the implied participation costs by group.

[29] Refer to Appendix Table G.3 for parameter estimates for these restricted belief models.

[30] A positive association between education and patience is consistent with prior evidence; see, e.g., Becker and Mulligan (1997) and Doepke and Zilibotti (2008).

[31] Aguiar and Hurst (2008) document that less-educated individuals consume more leisure than higher-educated individuals, even conditional on working.

Social Security program beliefs

Our quantitative model includes uncertain beliefs about the size of the early-claiming penalty for Social Security benefits.[32] This section discusses the how we estimate the subjective belief distribution over the annual early-claiming penalty, $B^{e,m}(k)$, that individuals use when evaluating the claiming decision.

Data inputs: belief types and confidence. The Understanding America Study (UAS) asks respondents about how Social Security benefits change with claiming age. The question presents respondents with four discrete options for the perceived annual early-claiming penalty,

$$\tilde{\gamma}_1^k \in \{0.0,\ 0.01,\ 0.067,\ 0.15\}.$$

We define an individual's *belief type* $k \in \{1, 2, 3, 4\}$ as the option they select. Table 4 reports (i) the empirical distribution across these belief types and for each education and marital group and (ii) a second UAS measure that asks how knowledgeable respondents feel about the relationship between claiming age and benefits. We map this self-assessed knowledge into a numerical confidence scale and use the group-by-type average, denoted by ACS, as a directly observed measure of the weight that type k places on its penalty belief.[33]

Modeling belief uncertainty. We allow individuals to be uncertain about the true penalty and to assign positive probability to each of the three alternative penalty values from the UAS question. Formally, we represent program-belief uncertainty by a probability vector

$$\boldsymbol{\varepsilon}_k^{ss} \equiv \left(\varepsilon_{k,1}^{ss}, \varepsilon_{k,2}^{ss}, \varepsilon_{k,3}^{ss}, \varepsilon_{k,4}^{ss}\right), \qquad \varepsilon_{k,j}^{ss} \geq 0,\ \sum_{j=1}^{4} \varepsilon_{k,j}^{ss} = 1.$$

We set the diagonal elements using the survey-based confidence measure, $\varepsilon_{k,k}^{ss} = \text{ACS}_k$, and interpret the remaining mass as the probability that respondents attach to other penalty values.

Because choices depend on continuation values (not only the mean penalty), individuals form expectations over penalty schedules when evaluating the claiming decision. For exposition, we

[32] Unlike MacCuish (2024), we do not allow individuals to learn about the statutory penalty over the life cycle. Using panel evidence in the HRS, we find no systematic updating in average belief errors over age; see Appendix F.8 and Appendix Figure F.10.

[33] The confidence scale is: very knowledgeable (1.0), somewhat knowledgeable (0.75), not too knowledgeable (0.25), and not at all knowledgeable (0.0).

Table 4: Program misinformation shares and confidence

		Belief Type, k			
		$\tilde{\gamma}_1^k = 0.0$	$\tilde{\gamma}_1^k = 0.01$	$\tilde{\gamma}_1^k = 0.067$	$\tilde{\gamma}_1^k = 0.15$
Non-college, single	Share	33.8	27.3	30.4	8.5
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.464	0.495	0.556	0.507
College, single	Share	21.9	25.9	47.6	4.6
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.651	0.535	0.710	0.847
Non-college, married	Share	29.2	32.3	34.1	4.4
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.604	0.592	0.658	0.546
College, married	Share	17.5	27.7	49.2	5.6
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.658	0.671	0.756	0.719

Source: Understanding America Study (UAS), authors' calculations.

Notes: For each education-marital status group, the columns report stated annual early claiming penalty,$\tilde{\gamma}_1^k$, for belief type k. The top line (Share) reports the distribution of responses across the four belief survey options. The second line (ACS) shows the average confidence scale for the group. The confidence scale is based on self-assessed knowledge: very knowledgeable (1.0), somewhat knowledgeable (0.75), not too knowledgeable (0.25), or not at all knowledgeable (0.0). The ACS is the average of these scores within the group.

summarize the implied distribution by the type-specific uncertainty weighted penalty:

$$\mathbb{E}[\tilde{\gamma}_1^{ss} \mid k] = \sum_{j=1}^{4} \varepsilon_{k,j}^{ss}\, \tilde{\gamma}_1^j = \varepsilon_{k,k}^{ss} \tilde{\gamma}_1^k + \sum_{j \neq k} \varepsilon_{k,j}^{ss} \tilde{\gamma}_1^j. \tag{18}$$

Estimation and identification. The UAS identifies $\tilde{\gamma}_1^k$, the distribution of these belief types across education and marital groups, and $\varepsilon_{k,k}^{ss}$, the weight placed on the individual's own belief, but it does not reveal how the remaining probability mass is allocated across $j \neq k$. We estimate the off-diagonal elements $\{\varepsilon_{k,j}^{ss}\}_{j \neq k}$ within the structural model using claiming behavior. Specifically, for each education and marital group, we match observed cumulative claiming rates between ages 62 and 69. Because each row k of $\boldsymbol{\varepsilon}_k^{ss}$ has three off-diagonal elements that sum to $1 - \text{ACS}_k$, each row contributes two free parameters; with four belief types this yields $4 \times 2 = 8$ free parameters per group. We match eight claiming moments per group, so the claiming distribution tightly disciplines these uncertainty weights.[34]

This approach assumes that the model captures all major drivers of claiming behavior and any

[34] We test the robustness of our key results to an alternate distance-based multinomial-logit parameterization of the off-diagonal elements; details are in Appendix J.3.

residual mismatch between the model and observed claiming distributions reflects uncertainty in beliefs about the penalty. This assumption is reasonable within our rich quantitative framework, as uncertainty about program rules can also serve as a proxy for broader social influences, such as peer effects which are known to shape both retirement savings and timing decisions (Duflo and Saez, 2002; Brown and Laschever, 2012); uncertainty about the solvency of the Social Security system, which has been shown to encourage individuals to move claims earlier (Gustman and Steinmeier, 2015; Benitez-Silva et al., 2009); or other behavioral mechanisms including reference dependent behavior (Behaghel and Blau, 2012).[35]

Figure 5 describes the identification of the off-diagonal weights. Panel (a) plots the deviation of the subjective penalty from the true penalty for each education and marital group, $\mathbb{E}[\tilde{\gamma}_1^{ss}] - \gamma_1^{ss}$, where $\gamma_1^{ss} = 0.067$, under three models: (i) No subjective beliefs (all agents understand the true penalty), (ii) Fixed stated beliefs (a special case corresponding to $\varepsilon_{k,k}^{ss} = 1$ and $\varepsilon_{k,j\neq k}^{ss} = 0$.), and (iii) Stated beliefs + estimated uncertainty (our full model program-belief specification). Panel (b) shows the corresponding deviation in average claiming ages (model minus data).

In a model without subjective beliefs (i), the subjective penalty is mechanically identical to the true, statutory penalty. However, the average claiming ages differ markedly from what is observed in data. Moving from (i) to (ii), the group-level expected penalty falls because the empirical distribution of stated beliefs is skewed toward low penalties (Table 4); as a result, the model shifts claiming earlier. Moving from (ii) to (iii), the estimated off-diagonal weights reallocate part of the residual probability mass toward higher-penalty schedules for groups that would otherwise claim too early relative to the data, thereby increasing the perceived cost of early claiming and delaying claims. Consistent with this logic, the estimated uncertainty-weighted penalty can exceed the statutory penalty for singles (for whom the fixed-belief model over-predicts early claiming), while remaining below the statutory penalty for married households (for whom fixed beliefs already bring claiming closer to the data). Finally, although the off-diagonal weights are estimated flexibly, the resulting belief matrices are well-behaved: in every group and belief type, the stated belief receives the largest probability mass. Therefore, the implied uncertainty weighted penalty remains close to the stated penalty for high-confidence types.

5 Model fit

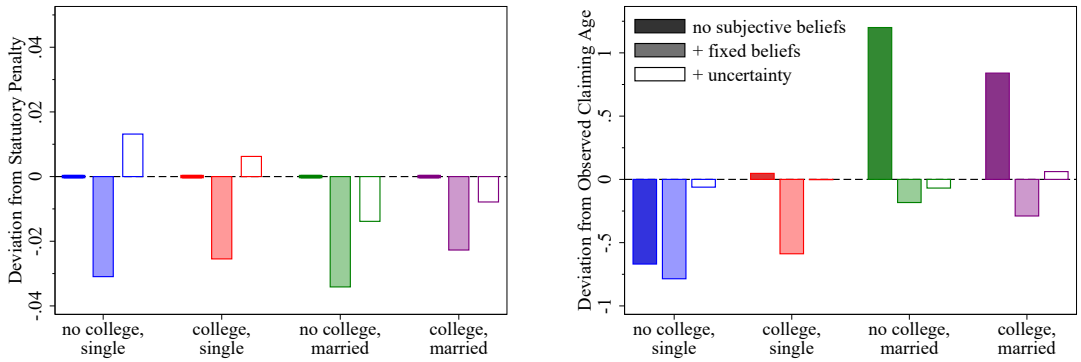
5.1 Targeted moments

Figures G.3-G.6 in Appendix G.2 show the benchmark model fit for average labor force participation, hours worked (conditional on participation), median wealth, and cumulative Social Se-

[35] We include a discussion of reference dependence in Appendix J.1

Figure 5: Deviations from Observed Claiming and Statutory Penalty

(a) Weighted Penalty Belief - Statutory Penalty (b) Model Implied - Observed Claiming Age



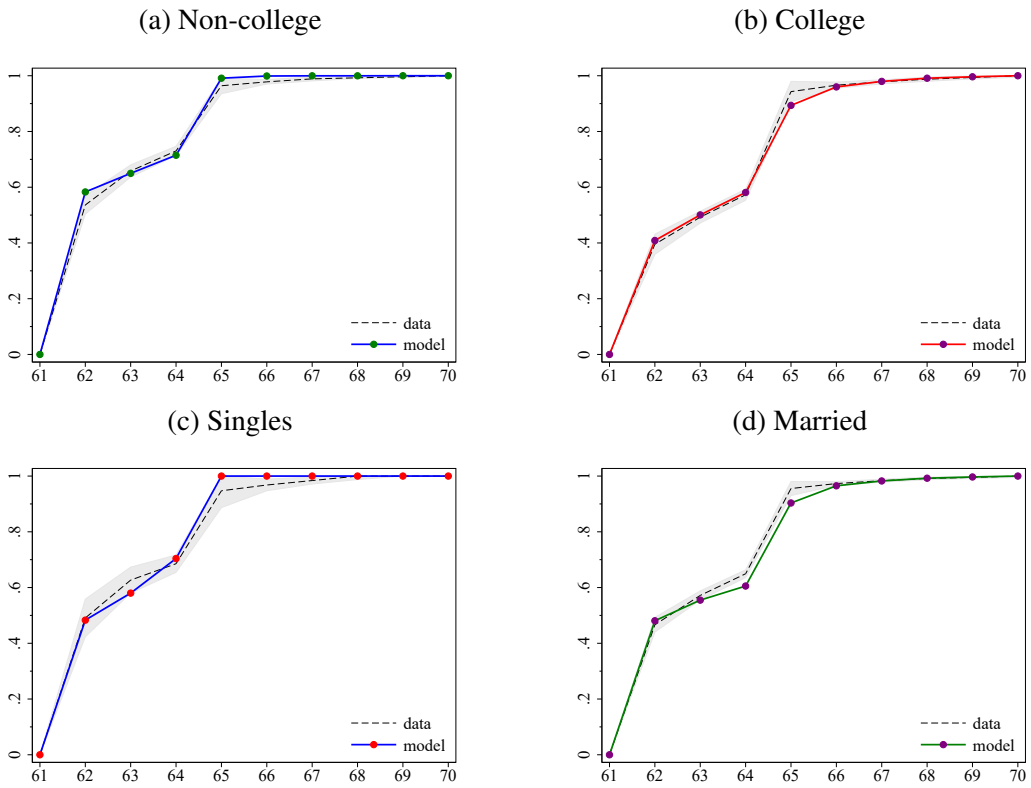
Notes: figure shows the deviation between the average perceived penalty and the statutory penalty of 6.67% for a given education-marital status group (Panel (a)) and the deviation between the model implied average claiming age and that observed in the Restricted HRS data for a given education-marital status group (Panel (b)) for three models: (1) a model with no subjective program beliefs in which all agents know the true penalty, (2) a model in which agents have a fixed subjective belief about the penalty, and (3) a model which includes our estimated uncertainty parameters.

curity benefit claims over the life cycle for male household heads by education and marital status groups. The model matches most empirical patterns well, though it somewhat underpredicts the decline in hours among non-college single men later in life. This is consistent with prior evidence on occupational downgrading and limited access to flexible or part-time employment among less-educated workers (Rutledge et al., 2017), and may be particularly relevant for non-college singles, the most financially constrained group in our model. Notably, the model matches the decline in hours for other education-marital groups more closely. While the wage process accounts for selection into employment, fully capturing this decline likely requires an explicit modeling of occupational transitions—an extension we view as beyond the scope of the current work.

Figure 6 shows cumulative Social Security claiming rates separately by education and marital status. As the model is calibrated to match claiming behavior within each education-marital status group, the aggregate claiming profile naturally aligns closely with the data. While this is not surprising, it is reassuring: the model replicates average claiming behavior for the 1931–1935 cohort with near-perfect accuracy (see panel (a) of Appendix Figure G.6), leaving no concern that the average claiming behavior might be misrepresented due to composition effects.

Apart from matching the overall Social Security claiming behavior extremely well, the model is also able to predict several important characteristics of claiming behavior. Most notable among these facts is the lack of heterogeneity in claiming ages for married vs. singles (claiming ages 63.4 and 63.5 respectively; gap +0.1) and college vs. non-college groups (claiming ages 63.7 and 63.2 respectively; gap +0.6), despite there being significant heterogeneity in their budget sets

Figure 6: Benchmark: Social Security claiming behavior



Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.
Notes: shows the share of non-college educated (Panel (a)), college educated (Panel (b)), single (Panel (c)), and married (Panel(d)) men born 1926-1939 who have claimed before or at each age 61-70.

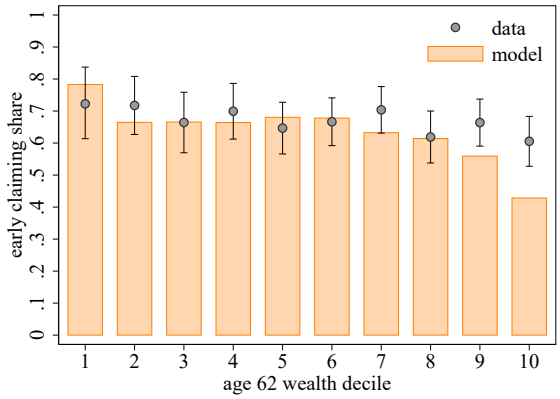
and life expectancies. The model closely mimics this lack of gap in claiming ages across diverse socioeconomic groups where the model counterparts for the claiming age gaps across marital status and education groups are +0.3 and +0.6 respectively.

5.2 Non-targeted moments

Next, we check to see if the model is able to generate a reasonable prediction for moments that were not targeted explicitly in the structural estimation. Wealth is an important indicator of an individual's ability to delay claiming to avoid costly penalties. Figure 7 shows the percent of those claiming Social Security benefits before the normal retirement age in each wealth decile (at age 62) for both the data and the model. The model lines up with the data well, despite predicting somewhat lower claiming among the top decile of wealth. Appendix G.2 presents additional model fit results by health status and labor force participation prior to claiming.

It is important that the model captures some aspects of heterogeneity in claiming within married households as Social Security policy varies by marital status. The model allows for the flexibility

Figure 7: Early claiming by wealth deciles



Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.
Notes: Scattered points show mean estimates and capped lines indicate 95% confidence intervals for the share of men (born 1926–1939) in each wealth decile at age 62. Orange bars show corresponding model values.

for spouses to receive benefits either on their own earnings record or on that of a spouse and successfully replicates the observed patterns in the data as shown in Table 5. Average spousal claiming ages both in the model and data roughly hover around 62 (except for college educated married men where the model under predicts spousal claiming age by nearly a year). Approximately 60 percent of model simulated married couples have a spouse claiming auxiliary spousal benefits. This is broadly consistent with the 52 percent observed in the data. The model also captures the qualitative pattern that auxiliary benefit claiming is more prevalent among spouses of non-college-educated heads than among those married to college-educated men. While the model is not explicitly designed to explain spousal claiming behavior, it nevertheless generates these empirical patterns and closely replicates spousal AIME at age 60 across education groups—even though spousal earnings are modeled using an exogenously estimated life-cycle process.

Note that these comparisons of model and data statistics puts the model to a considerably strong test. Our benchmark model is only explicitly designed to fit the life-cycle labor supply, wealth and cumulative social security claims of these different groups of individuals. The fact that it is in addition able fit several dimensions of heterogeneity in early claims is a testament to its strength as an approximation of the true data generating process. These tests demonstrate that the mechanisms incorporated in the structural model, along with the detailed policy features, are potentially driving claiming decisions in important ways.

Table 5: Non-targeted moments for married men

	Model	Data	
		Mean	[95% confidence interval]
Spouse claiming age	62.2	62.9	[62.9, 62.9]
Non-college	62.1	62.8	[62.8, 62.9]
College	62.2	63.1	[63.1, 63.2]
Spouse claiming on husband's record (%)	59.8	51.7	[44.3, 59.0]
Non-college	65.2	54.6	[44.3, 64.9]
College	56.8	49.0	[38.5, 59.5]
Spousal Average Age-60 AIME (1998$)			
Non-college	6,747.20	6,893.80	[6,225.67, 7,451.99]
College	10,709.40	9,932.90	[9,124.27, 10,741.55]

Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.
Notes: a spouse is defined as claiming benefits on her husband's record if her initial benefit claims are dual retirement benefits or auxiliary spousal benefits, and she does not receive survivors benefits before age 70. Spousal claiming age is completely determined by head's claiming age and the age-gap between the two.

6 Counterfactual experiments

Equipped with a model that replicates key features of Social Security claiming behavior, we conduct a series of experiments to understand the role of different mechanisms in explaining the behavior of heterogeneous groups. We summarize the claiming age distribution using average claiming ages as reported in Table 6. We also report early claiming rates in Appendix table H.1. We adopt an additive approach—introducing mechanisms incrementally to a simpler baseline—rather than removing them from the full model presented in Section 3. This avoids non-linear interactions that could mask or attenuate the effects of individual channels.

We begin with a simple model (Baseline A) in which claiming decisions are primarily driven by the trade-off between the number of years benefits are received and the size of annuity payments, as detailed in the PDV framework in Section 2. Variation in behavior under Baseline A stems from differences in estimated objective survival probabilities and preferences across education and marital groups. While this baseline omits key mechanisms from the full model and is not expected to match observed behavior, it serves as a clean starting point for subsequent extensions.

Next, we construct Baseline B by adding mechanisms identified in prior literature as important for labor market exits and benefit claims: employment shocks, health status, and bequest motives. These additions improve the model's fit to overall claiming behavior but still fail to reproduce the limited heterogeneity observed across socioeconomic groups. Baseline B serves as our reference

when evaluating the impact of the key mechanisms emphasized in this paper: marital benefits and subjective beliefs about survival and program rules.

6.1 Baseline A: no mechanisms

The second row of Table 6 reports average claiming ages, as well as gaps in claiming ages, by education and marital status under the no-mechanisms baseline.[36] This simple specification predicts an average claiming age of 65.8 years, about 2.4 years higher than observed, and an early claiming rate of only 21.3 percent, over 40 percentage points below the data (see rows 1 and 2 of Table 6 and Appendix Table H.1). In this setup, claiming decisions are largely driven by life expectancy and time preferences. Differences in estimated discount factors and survival probabilities generate counterfactual variation in predicted claiming behavior across groups. For example, while observed average claiming ages cluster around 63, the no mechanisms baseline predicts a range from 64 (singles) to nearly 67 (married and college-educated men).

Claiming age gaps: This simple model generates exaggerated differences in claiming ages across groups. The education gap is 1.7 years, compared to 0.6 years in the data; the marital status gap is 2.8 years versus just 0.1. This mirrors our present value calculations in Section 2, which also predict larger gaps based solely on life expectancy than those observed empirically. Notably, the model incorporates heterogeneity in preferences, income, and wealth—further amplifying these gaps relative to simple PDV-based predictions.

6.2 Baseline B: standard mechanisms

We expand Baseline A by incorporating three mechanisms emphasized in the prior literature—health risk, employment shocks, and a bequest motive—and refer to this enriched specification as Baseline B. The magnitude of each mechanism is allowed to vary by education and marital status, consistent with the heterogeneity observed in the data. The effects of these mechanisms, both individually and jointly, are reported in rows 3–6 of Table 6 and discussed below.

Health and employment shocks Health and employment shocks affect claiming behavior primarily by reducing lifetime earnings and thus accumulated Social Security wealth. Relative to Baseline A (where we shut down the labor-market consequences of these shocks), introducing the health-related work limitation lowers mean Social Security wealth at age 60 by about 3%, while introducing unemployment risk with wage scarring lowers it by about 9%. Additionally, age-60

[36] Appendix H outlines how this baseline is derived from the full model in Section 3.

Table 6: Social Security claiming behavior: role of different channels

Experiment	All	Education status			Marital status		
		Non-College	College	Gap	Singles	Married	Gap
Data	63.4	63.1	63.7	+ 0.6	63.3	63.4	+ 0.1
Baseline A (No-Mechanisms)	65.8	64.7	66.5	+ 1.7	63.9	66.7	+ 2.8
Health Risk	65.8	64.7	66.6	+ 1.8	63.8	66.7	+ 2.9
Employment Shocks	65.8	64.5	66.7	+ 2.2	63.7	66.8	+ 3.1
Bequest Motive	64.6	63.9	65.2	+ 1.3	63.1	65.4	+ 2.2
Baseline B (Health+Employment+Bequest)	64.6	63.7	65.2	+ 1.5	62.9	65.4	+ 2.6
(1) Baseline B + Marital Benefits							
Spousal	64.2	63.5	64.6	+ 1.1	62.9	64.8	+ 1.9
Survivors	65.0	63.9	65.8	+ 1.9	62.9	66.0	+ 3.1
Spousal & Survivor	64.4	63.7	64.9	+ 1.3	62.9	65.1	+ 2.3
(2) Baseline B + Program Misbeliefs							
Fixed Beliefs	63.4	62.8	63.8	+ 1.0	62.6	63.7	+ 1.2
Beliefs & Uncertainty	63.6	63.0	64.1	+ 1.1	63.1	63.9	+ 0.9
(3) Baseline B + Mortality Misbeliefs	64.1	63.4	64.6	+ 1.2	63.0	64.6	+ 1.7
(4) Baseline B + Subjective Beliefs (2) + (3)	63.4	63.1	63.7	+ 0.6	63.2	63.5	+ 0.3
Full Model: Baseline B + (1) + (2) + (3)	63.4	63.1	63.7	+ 0.6	63.2	63.5	+ 0.3

Notes: *Average SS benefit claiming age and claiming age gap across groups is reported for each experiment. Each experiment listed under (1)-(3) entails including given channel to the standard mechanisms of Baseline B to understand its impact on claiming. No mechanisms (Baseline A) case refers to a scenario in which all mechanisms are omitted.

consumption is approximately 12–17 percent lower when we consider a model with employment risk (see Appendix K and Appendix Figure K.1).

Lower wealth increases the marginal utility of current consumption and weakens the incentive to delay claiming in order to obtain higher future benefits. As a result, lower-SES groups—who are more exposed to these shocks—respond by claiming earlier. In particular, employment shocks increase early claiming rates by 6.9 and 6.6 percentage points for non-college-educated and single individuals, respectively. In contrast, higher-wealth groups are better able to smooth consumption and therefore continue to delay claiming to avoid early-claiming penalties (see Appendix Figure K.2).

Although these mechanisms raise early claiming rates overall, they do not affect average claiming age (rows 3 and 4 of Table 6). Importantly, because non-college-educated and single individuals respond more strongly, these shocks increase heterogeneity in claiming behavior: employment shocks widen the education gap in claiming ages from 1.7 years in Baseline A to 2.2 years, and the marital-status gap from 2.8 to 3.1 years.

Bequest motive An individual with a bequest motive and an uncertain lifespan may claim earlier to preserve the wealth that can be passed to heirs as Social Security benefits stop at death. Bequest motives, therefore, encourage early claiming and the use of benefit flows to finance consumption and preserve private wealth for bequests.

Adding bequest motives to the no-mechanisms baseline lowers the average claiming age by 1.2 years (see row 5 of Table 6). In fact, when all standard mechanisms are combined (row 6), the entire reduction in the average claiming age is attributable to bequest motives. While these motives affect all groups, higher bequest intensities among college-educated and married individuals make the effect somewhat stronger: their average claiming ages decline by 1.3 years, compared to 0.8 years for non-college-educated and single individuals, respectively. As a result, bequest motives slightly compress claiming age differences by education and marital status.

Claiming age gaps: Addition of bequest motives narrows the education gap from 1.7 years in Baseline A to 1.3 years, and the marital gap declines from 2.8 to 2.2 years. However, when bequest motives are combined with health and employment shocks—which tend to somewhat widen these gaps—their equalizing effect is largely offset. Under all standard mechanisms together (Baseline B), the marital claiming gap drops to only 2.6 years, while the education gap narrows only slightly from 1.7 to 1.5 years. These findings suggest that, while standard mechanisms do shape overall claiming behavior, they explain little of the observed compression in claiming age gaps across education and marital groups. We next examine whether marital benefits and subjective beliefs offer a more compelling explanation for these patterns.

6.3 Marital benefits

As Social Security benefits differ by marital status, we examine the role of marital benefits in shaping claiming behavior. As outlined in Section 2, we focus on two key provisions available to married individuals: spousal benefits and survivors benefits. The value of these depends on the spouse's earnings. Spousal benefits encourage earlier claiming when they exceed the benefits the spouse would receive on her own record. In contrast, survivors benefits are valuable to households where the spouse's benefit (own or spousal) is lower than the primary earner's as delaying claiming raises the eventual survivors payout.

Section (1) of Table 6 isolates these forces by adding marital benefits, one at a time, to Baseline B. Among married men, introducing spousal benefits alone reduces the average claiming age by 0.6 years (and raises early claiming by 32.3 percentage points). Introducing survivor benefits alone has the opposite effect, increasing the average claiming age by 0.6 years (and reducing early claiming by 10.3 percentage points). These impacts are mechanical: spousal benefits are tied to the primary earner's claiming decision and do not generate delayed retirement credits for the spouse, so they weaken the incentive to delaySurvivor benefits, by contrast, extend the payoff to delaying beyond the primary earner's lifetime and therefore strengthen the incentive to delay.

When both spousal and survivor benefits are included, the net effect is a 0.3-year decline in the average claiming age for married men (and a 25.5 percentage point increase in early claiming), indicating that the spousal-benefit channel dominates survivors channel. Singles are mechanically unaffected by these provisions, so this asymmetric response reduces the difference in claiming behavior between married and single men.

Heterogeneity in the response across education groups arises primarily from baseline behavior. Under Baseline B, non-college married men already claim relatively early (average claiming age 64.6), leaving limited scope for spousal benefits to further accelerate claiming. College-educated married men claim later (65.8) and therefore have more scope to respond; accordingly, their claiming age falls by 0.4 years once marital benefits are introduced (Appendix Table H.2).

Claiming age gaps: Marital benefits affect education-related claiming gaps through both composition and differential behavioral responses. 75 percent of college-educated men are married as opposed to 58 percent of non-college men. Therefore, they are more exposed to the impact of spousal and survivor benefits. Combined with the stronger behavioral response described above, this composition effect lowers the average claiming age of college-educated men, while leaving the non-college average roughly unchanged.

Quantitatively, adding marital benefits to Baseline B reduces the education claiming-age gap from 1.5 years to 1.3 years and the marital-status gap from 2.6 years to 2.3 years. As summarized in Table 7, marital benefits therefore explain 29.2 percent of the compression in the education gap

and 12.8 percent of the compression in the marital gap between Baseline B and the full model.

6.4 Subjective beliefs

To quantify how subjective beliefs about Social Security program rules and one's own survival prospects affect claiming decisions, we introduce each belief object into Baseline B holding all other model components fixed. The resulting changes in claiming behavior are reported in Sections (2) and (3) of Table 6.

Program beliefs The UAS provides two key inputs for program beliefs (discussed in Section 4.2): individuals' stated belief about the annual early-claiming penalty (0%, 1%, the true 6.7%, or 15%) and the confidence that each type places on its stated belief. To isolate the roles of (a) the *level* of stated beliefs and (b) *uncertainty* around those beliefs, we introduce program beliefs into Baseline B in two stages.

Role of fixed beliefs: We first consider a fixed-beliefs version of the model in which agents differ in their stated penalty belief type k according to Table 4, but hold full certainty in that belief (i.e., the weight place on the belief $\varepsilon^{ss}_{k,k}$, is equal to 1). This exercise isolates the mechanical implications of replacing the statutory penalty with the (mis)perceived penalty schedules present in the data.

Section (2) of Table 6 shows that introducing only fixed stated beliefs lowers the overall average claiming age by 1.2 years relative to Baseline B. The effect is concentrated among higher-SES groups: claiming ages fall by 1.4 years for college-educated individuals and by 1.7 years for married household heads. By contrast, non-college and single individuals—who already claim near age 62 in Baseline B—experience smaller changes (0.9 and 0.3 years, respectively).

The mechanism is transparent: the stated-belief distribution in Table 4 is skewed toward low penalties (0% and 1%) in every group. Therefore, replacing the statutory schedule with these perceived schedules lowers the perceived cost of early claiming and shifts claiming earlier. This is also visible in Figure 5a): moving from a model without subjective beliefs to one with fixed stated beliefs generates a similar negative deviation between the average perceived penalty and the statutory penalty in all groups. Because college and married heads delay more in Baseline B, there is more scope for these low-penalty beliefs to pull their claims toward earlier ages.

Role of uncertainty: We next incorporate uncertainty around the stated belief types. The diagonal weight on the stated belief, $\varepsilon^{ss}_{k,k}$, is disciplined by the UAS confidence measure, while the off-diagonal weights $\varepsilon^{ss}_{k,j\neq k}$ are estimated to match the group-specific cumulative claiming distribution between ages 62 and 69.

Introducing uncertainty attenuates the fixed-belief effects. The overall average claiming age falls by 1.0 year (instead of 1.2 years). The largest declines remain among higher-SES groups (1.1 years for college and 1.5 years for married). By contrast, singles are largely unaffected: their average claiming age *rises* slightly (by 0.2 years), while non-college heads experience a smaller decline (0.7 years).

Panel (a) of Figure 5 clarifies why uncertainty pushes in different directions across groups. For singles, fixed beliefs push claiming *too early* relative to the data. Therfore, the estimated uncertainty weights shift residual probability mass toward higher-penalty schedules (notably the 15% option), implying an uncertainty-weighted penalty that can exceed the statutory value and modestly delaying claims. For married households, fixed beliefs already bring claiming close to the data, so less upward re-weighting is required and the uncertainty-weighted penalty remains below the statutory penalty. These opposite adjustments imply that composition matters for the education gap: because non-college men are more likely to be single, the uncertainty correction that raises perceived penalties receives greater weight in the non-college average; because college men are more likely to be married, the persistent underestimation among married household heads receives greater weight in the college average. Consequently, program beliefs compress the education-based claiming-age gap, even though the primary source of heterogeneity in estimated belief uncertainty operates along the marital-status margin.

Table 7: Social Security claiming age gaps: role of different channels

% gap explained	Claiming age-gap	
	Education	Marriage
Marital benefits	29.2	12.8
Program beliefs	48.6	75.1
Mortality beliefs	38.9	38.7
Subjective beliefs (Program + Mortality)	99.7	99.0

Notes: Gap reduction explained is the difference between the claiming age gap under Baseline B and after introducing a given channel. Total gap is the difference between Baseline B and the full model (education: $0.9 = 1.5 - 0.6$; marital: $2.3 = 2.6 - 0.3$). Percent reduction explained = gap reduction / total gap (e.g., for program beliefs and education: $(1.5 - 1.1)/0.9 \approx 49\%$).

Survival beliefs Table 2 shows both objective and subjective life expectancy at age 60. Two empirical patterns emerge. First, there are systematic differences in average subjective longevity across groups. On average, singles report subjective life expectancy that exceeds the objective benchmark, while married individuals report subjective life expectancy below objective levels. A

similar pattern emerges by education: non-college men are close to objective longevity on average, whereas college-educated men systematically underestimate their survival prospects. Second, these differences are driven by those in good health who substantially underestimate objective life expectancy. This interacts with the fact that health composition differs sharply by socioeconomic status: high-SES groups are much less likely to be in bad health at age 60 (Appendix Table K.1). As a result, pessimistic subjective survival beliefs are especially relevant for the groups that otherwise have the strongest incentives to delay claiming.

Section (3) of Table 6 shows that introducing subjective survival beliefs lowers the average claiming age by 0.5 years relative to Baseline B. As with program beliefs, the largest effects occur among higher-SES groups: married individuals' claiming age falls by 0.8 years, and college-educated individuals' claiming age declines by 0.6 years. Singles are essentially unchanged (a 0.1-year *increase*), consistent with the fact that their average subjective life expectancy is slightly higher than objective life expectancy (Table 2).

Finally, beyond explaining heterogeneity in claiming, subjective beliefs also help account for key empirical features such as the spike in claiming at age 65, a focus of recent literature (Pashchenko and Porapakkarm, 2024; Behaghel and Blau, 2012). In our model, the low delayed retirement credit contribute to this spike, consistent with previous studies. However, absent subjective beliefs, our higher estimated discount factors would generate an exaggerated peak at age 65—except for singles—as shown in Baseline B (Appendix Figure K.3).

Claiming age gap Subjective beliefs are the central force compressing claiming-age differences across socioeconomic groups. Introducing program beliefs with uncertainty reduces the education gap from 1.5 to 1.1 years and the marital gap from 2.6 to 0.9 years. Introducing subjective survival beliefs reduces these gaps to 1.2 and 1.7 years, respectively. As summarized in Table 7, program beliefs account for 48.6% of the education-gap reduction and 75.1% of the marital-gap reduction, while survival beliefs account for an additional 38.9% and 38.7%, respectively.

These effects operate through both *behavioral* responses within groups and *composition* across groups. In particular, education groups differ in their marital composition and in the distribution of belief types and confidence. Thus, education-level averages weight married- and single-specific belief distortions differently. Because program beliefs affect married and single households in opposite directions once uncertainty is accounted for, these composition differences contribute to the compression of the education claiming-age gap.

When both belief objects are included jointly, they explain nearly all of the compression in claiming ages between Baseline B and the full model (99.7% of the education gap and 99.0% or the marital status gap).

7 Policy experiment

The counterfactual exercises in the Section 6 quantify how each mechanism helps reconcile the observed homogeneity in claiming ages across groups. We now ask how these mechanisms shape policy responsiveness and the fiscal projections from a major Social Security reform. We consider an increase in the Normal Retirement Age (NRA) to age 70 (henceforth,"NRA 70"). Under this reform, delayed retirement credits are eliminated and all claiming prior to age 70 is subject to the 6.7% actuarial reduction.

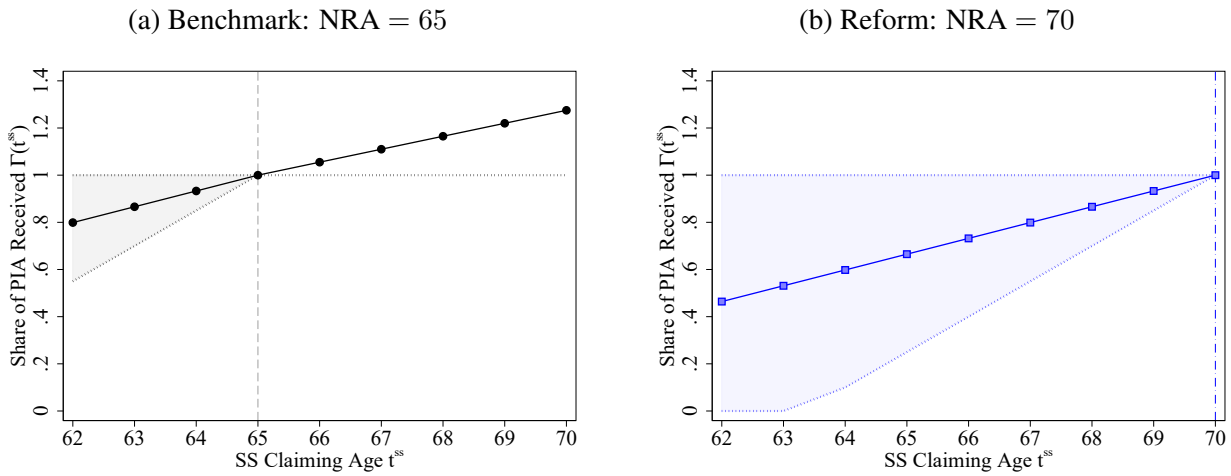
Policy design and beliefs under the reform In our benchmark model (NRA of age 65 with delayed retirement credits from ages 66–70), $\Gamma\left(t^{ss}\right)$ represents the share of the the Primary Insurance Amount (PIA) paid as benefits when claiming at age t^{ss} and $\Gamma^{P}\left(t^{ss}\right)$ denote the analogous schedule under NRA 70 policy.

The solid lines in Figure 8 plots these schedules for the baseline policy (Figure 8a) and the reform (Figure 8b). These schedules represent the statutory change in benefits due to the claiming age. Policy knowledge, however, may cause the perceived change to deviate from this schedule. The perceived schedule may fall anywhere in the shaded areas of each graph with the top of the area representing someone who believes there is no adjustment for early claims and the bottom of the shaded are representing someone who believes benefits decrease 15% for each year prior to the NRA that benefits are claimed. We assume that subjective beliefs on the early claiming penalty stays constant after the reform, i.e. the slope of the perceived benefit schedule remains the same in both policy scenarios.

Three mechanical features of the reform are useful to keep in mind. First, moving the NRA to age 70 represents a benefit cut at any given claiming age. For example, claiming at age 65 delivers 100% of PIA under benchmark rules but only about 67% of PIA under NRA 70 policy. Second, the marginal return to delaying claiming between ages 65 and 70 becomes larger because delaying avoids an additional early-claiming reduction under NRA 70 (rather than earning a smaller delayed retirement credit of 5.5% under the benchmark). Finally, as we model program knowledge and uncertainty as augmenting the perceived annual adjustment for claims prior to the NRA, the shaded region in which the perceived adjustment schedule falls grows larger as the NRA moves later. For example, if an individual believes there to be a 15% prior to the reform, he would believe claiming at age 64 provides 85% of PIA as benefits. Under the reform, this 15% compounds over a larger number of years and the same individual would expect claiming at 64 to provide only 10% of PIA.

Claiming Response For each model specification, we compute (i) the change in the average claiming age and (ii) the percent change in lifetime value of *net* Social Security payments for the

Figure 8: Benefit adjustment schedules and program-belief heterogeneity



Notes: The figure plots the share of PIA paid as a function of the claiming age. The solid line corresponds to the schedule implied by the accurate annual adjustment (6.7%), while the shaded region spans the schedules implied by the extreme belief types (0% and 15%).

household head under NRA 70 relative to the corresponding baseline. Net payments are defined as lifetime Social Security benefits received minus Social Security portion of the FICA taxes paid.[37]

Table 8 shows that NRA 70 generates substantial increases in claiming ages in models without the key mechanisms emphasized in this paper, but markedly smaller increases in our baseline model. In Baseline A (No-Mechanisms) and Baseline B (Standard-Mechanisms), the average claiming age rises by 3.1 years and 2.7 years, respectively. By contrast, in our benchmark model the average claiming age rises by only 1.5 years. The reform creates a strong incentive to delay claiming, but the full model contains mechanisms that reduce the elasticity of claiming behavior.

Simpler baselines also imply substantially more heterogeneity in policy responsiveness than the full model. Under Baseline B, married households are predicted to respond very strongly (+3.6 years) relative to singles (+0.7 years), and college-educated individuals respond much more than non-college (+3.3 vs. +1.7 years). In the full model, these gaps shrink markedly: the response is +1.8 years for married versus +0.9 for singles, and +1.8 for college versus +1.1 for non-college.

To isolate the source of this muted responsiveness, we start from Baseline B, adding each mechanism at a time as in Table 7 above. Adding *program beliefs* (and holding their distribution and uncertainty fixed under the reform) reduces the average response from +2.7 to +1.7 years—nearly matching the full model's +1.5 years. Notably, program beliefs are also the main force compressing the predicted heterogeneity: relative to Baseline B, the response of married individuals falls sharply from +3.6 to +2.0 years and the college response falls from +3.3 to +2.0 years. Adding *mortality beliefs* alone yields a substantially larger response (+2.2 years) and pre-

[37] This experiment retains the benchmark earnings test parameters.

Table 8: NRA 70: Changes in Social Security claiming behavior

Experiment	All	Education status		Marital status	
		Non-college	College	Singles	Married
Baseline A (no-mechanisms)	65.8	64.7	66.5	63.9	66.7
Δ *claiming age*	+3.1	+3.3	+2.9	+2.8	+3.2
$\%\Delta$ *in lifetime SS benefits paid*	-21.8	-23.5	-21.0	-24.9	-20.8
Baseline B (standard-mechanisms)	64.6	63.7	65.2	62.9	65.4
Δ *claiming age*	+2.7	+1.7	+3.3	+0.7	+3.6
$\%\Delta$ *in lifetime SS benefits paid*	-23.4	-26.5	-21.7	-28.4	-21.7
Baseline B + marital benefits	64.4	63.7	64.9	62.9	65.1
Δ *claiming age*	+2.1	+1.9	+2.3	+0.7	+2.8
$\%\Delta$ *in lifetime SS benefits paid*	-23.5	-26.3	-22.0	-28.4	-21.8
Baseline B + program beliefs	63.6	63.0	64.1	63.1	63.9
Δ *claiming age*	+1.7	+1.3	+2.0	+1.0	+2.0
$\%\Delta$ *in lifetime SS benefits paid*	-24.3	-27.3	-22.7	-28.1	-23.0
Baseline B + mortality beliefs	64.1	63.4	64.6	63.0	64.6
Δ *claiming age*	+2.2	+1.1	+3.0	+0.8	+2.9
$\%\Delta$ *in lifetime SS benefits paid*	-23.9	-27.2	-22.1	-28.3	-22.4
Baseline B + subjective beliefs	63.4	63.1	63.7	63.2	63.5
Δ *claiming age*	+1.6	+0.9	+2.0	+0.9	+1.9
$\%\Delta$ *in lifetime SS benefits paid*	-24.5	-27.7	-22.9	-28.3	-23.3
Full model	63.4	63.1	63.7	63.2	63.5
Δ *claiming age*	+1.5	+1.1	+1.8	+0.9	+1.8
$\%\Delta$ *in lifetime SS benefits paid*	-24.4	-27.4	-22.9	-28.3	-23.2

Notes: For each experiment, the first line reports the average claiming age under the status quo baseline for that model, the second line reports the change in the average claiming age under NRA 70, and the third line reports the percent change in the present-discounted value of lifetime Social Security *net payments* for the household head (benefits received minus FICA payroll taxes paid) under NRA 70 relative to that experiment's baseline.

serves larger cross-group differences (e.g., $+3.0$ for college and $+2.9$ for married), indicating that survival-belief heterogeneity is not the primary driver of dampened policy responsiveness in this reform. Adding *marital benefits* also reduces the response (to $+2.1$ years), but the comparison between "Baseline B + program beliefs" and the full model shows that program-rule beliefs are the dominant channel behind the muted aggregate response. This shows that the dampening effect is largely attributable to subjective program-rule expectations.

This reform experiment assumes that the reform itself does not impact beliefs about the policy. If, rather, the reform leads individuals to reset their own expectations about policy rules, we would expect behavioral responses to be altered. However, whether these altered beliefs mitigate or amplify behavioral responses would depend on how beliefs are updated.[38]

Fiscal Implications. Differences in behavioral responses to the same NRA 70 policy across model specifications have first-order implications for evaluating its fiscal impact. To connect behavior to fiscal projections, Table 8 also reports the implied percent change in the lifetime value of lifetime Social Security benefits net taxes paid for the household head.[39] A more negative value therefore corresponds to larger fiscal savings due to the reform.

Consistent with our benefit claiming response, the net-payments fall -21.8% due to the reform in Baseline A and -23.4% in Baseline B; this drop is -24.4% in the full model. The difference reflects that when households hold subjective beliefs about the program adjustments embedded in the reform or their lifespan, a larger mass continues to claim relatively early. As the behavioral adjustment is muted and a larger portion of the mechanical benefit cut translates into lower net payments. Figure 8 provides a transparent mapping from the reform to the program belief channel: under NRA 70, perceived incentives can remain weak for households who place substantial probability on low (or zero) early-claiming penalties. Aggregated to the population of men born 1931–1935, Baseline A understates the fiscal savings from raising the NRA by $19.6 billion (5.8 percent), and Baseline B understates them by $13.1 billion (3.9 percent).[40]

8 Conclusion

This paper studies a novel puzzle: the surprising similarity in Social Security claiming behavior across socioeconomic groups that differ substantially in life expectancy and financial resources. Standard life-cycle models predict that high-SES individuals should delay claiming more than their lower-SES counterparts, yet the data show only minimal differences in claiming ages. To explain this homogeneity, we develop and estimate a rich structural life-cycle model of consumption, labor supply, savings, and Social Security claiming that incorporates heterogeneity in preferences, marital benefits, and, critically, subjective beliefs about survival and program rules.

[38] For example, if the policy adjusts a reference point to age 70, we may expect individuals to respond more strongly to policy. If, rather, these updated beliefs incorporate increased uncertainty of the future of Social Security, we may expect behavior to be even less responsive to the policy change.

[39] We net out Social Security taxes paid by individuals.

[40] Baseline A predicts that NRA 70 reduces lifetime Social Security net payments by $72,532 per person on average; Baseline B predicts a reduction of $74,015 per person; and the full model predicts a reduction of $77,010 per person. Results are aggregated using a population of 4,382,482 men born 1931–1935 who would have survived to age 25 and be covered by Social Security (assumed to be 90% of workers).

Our counterfactual analysis shows that previously studied mechanisms—such as preference heterogeneity, health and employment shocks, and bequest motives—help explain overall claiming behavior but cannot account for the convergence in claiming ages across education and marital status. Instead, we find that subjective survival expectations and uncertainty about early claiming penalties play a central role, explaining most of both the education and marital gaps. Moreover, models that omit these mechanisms understate the fiscal savings from raising the Normal Retirement Age by up to $19.6 billion.

Our paper incorporates empirically grounded beliefs into a rich framework, offering new insights into how subjective perceptions shape claiming decisions and helping reconcile puzzling patterns in the data. Our findings have important implications for the design and evaluation of policy reforms and motivate further research into the formation of these beliefs. Finally, while our analysis focuses on a cohort facing a common set of policy rules, exploiting the substantial cross-cohort variation policy rules is a promising direction for future work.

References

Aguiar, M. and Hurst, E. (2008). The increase in leisure inequality. *NBER Working Paper No. w13837*.

Altonji, J. G. and Segal, L. M. (1996). Small-sample bias in gmm estimation of covariance structures. *Journal of Business & Economic Statistics*, 14(3):353–366.

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., and Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, 128(6):2375–2451.

Armour, P. and Knapp, D. (2023). The consequences of claiming social security benefits at age 62. *Journal of Pension Economics and Finance*, 22(2):211–237.

Bairoliya, N. (2019). Pension plan heterogeneity and retirement behavior. *European Economic Review*, 116:28–59.

Becker, G. S. and Mulligan, C. B. (1997). The endogenous determination of time preference. *Quarterly Journal of Economics*, 112(3):729–758.

Behaghel, L. and Blau, D. M. (2012). Framing social security reform: Behavioral responses to changes in the full retirement age. *American Economic Journal: Economic Policy*, 4(4):41–67.

Benitez-Silva, H., Dwyer, D. S., and Sanderson, W. C. (2009). Retirement and social security reform expectations: A solution to the new early retirement puzzle. *mimeo*.

Borella, M., De Nardi, M., and Yang, F. (2023). Are marriage-related taxes and social security benefits holding back female labour supply? *The Review of Economic Studies*, 90(1):102–131.

Brown, J. R., Coile, C. C., and Weisbenner, S. J. (2010). The effect of inheritance receipt on retirement. *The Review of Economics and Statistics*, 92(2):425–434.

Brown, J. R., Kapteyn, A., and Mitchell, O. S. (2016). Framing and claiming: How information-framing affects expected social security claiming behavior. *Journal of Risk and Insurance*, 83(1):139–162.

Brown, K. M. and Laschever, R. A. (2012). When they're sixty-four: Peer effects and the timing of retirement. *American Economic Journal: Applied Economics*, 4(3):90–115.

Chan, S. and Huff Stevens, A. (1999). Employment and retirement following a late-career job loss. *American Economic Review*, 89(2).

Chan, S. and Huff Stevens, A. (2001). Job loss and the employment patterns of older workers. *Journal of Labor Economics*, 19(2):484–521.

Chan, S. and Huff Stevens, A. (2004). How does job loss affect the timing of retirement. *Contributions to Economic Analysis and Policy*, 3(1).

Choukhmane, T., Goodman, L., and O'Dea, C. (2025). Efficiency in household decision-making: Evidence from the retirement savings of us couples. *American Economic Review*, 115(5):1485–1519.

Coile, C., Diamond, P., Gruber, J., and Jousten, A. (2002). Delays in claiming social security benefits. *Journal of Public Economics*, 84:357–85.

Dai, T., Sun, W., and Webb, A. (2024). How do subjective mortality beliefs affect the value of social security and the optimal claiming age? *International Studies of Economics*, 19:92–116.

de Bresser, J. (2024). Evaluating the accuracy of counterfactuals: heterogeneous survival expectations in a life cycle model. *Review of Economic Studies*, 91(5):2717–2743.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3):743–768.

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2021). Why do couples and singles save during retirement? Technical report, National Bureau of Economic Research.

De Nardi, M., Jones, J., and French, E. (2010). Why do the elderly save? *Journal of Political Economy*, 118(1):39–75.

Doepke, M. and Zilibotti, F. (2008). Occupational choice and the spirit of capitalism. *The Quarterly Journal of Economics*, 123(2):747–793.

Duflo, E. and Saez, E. (2002). Participation and investment decisions in a retirement plan: The influence of colleagues' choices. *Journal of public Economics*, 85(1):121–148.

Fan, X., Seshadri, A., and Taber, C. R. (2022). Estimation of a life-cycle model with human capital, labor supply and retirement. Technical report, National Bureau of Economic Research.

French, E. (2005). The effects of health, wealth, and wages on labor supply and retirement behavior. *The Review of Economic Studies*, 72(2):395–427.

French, E. and Jones, J. B. (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79(3):693–732.

Gan, L., Gong, G., Hurd, M., and McFadden, D. (2015). Subjective mortality risk and bequests. *Journal of Econometrics*, 188(2):514–525.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Gustman, A. and Steinmeier, T. L. (2015). Effects of social security policies on benefit claiming, retirement and saving. *Journal of Public Economics*, 129(C):51–62.

Gustman, A. L. and Steinmeier, T. L. (2005). The social security early entitlement age in a structural model of retirement and wealth. *Journal of public Economics*, 89(2):441–463.

Gustman, A. L. and Steinmeier, T. L. (2008). Projecting behavioral responses to the next generation of retirement policies. *Research in Labor Economics*, 28:141 – 195.

Health and Retirement Study (2024). Health and retirement study, restricted use dataset. https://hrs.isr.umich.edu/data-products. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740).

Heimer, R. Z., Myrseth, K. O. R., and Schoenle, R. S. (2019). Yolo: Mortality beliefs and household finance puzzles. *The Journal of Finance*, 74(6):2957–2996.

Hosseini, R., Kopecky, K. A., and Zhao, K. (2021). How important is health inequality for lifetime earnings inequality?

Hosseini, R., Kopecky, K. A., and Zhao, K. (2022). The evolution of health over the life cycle. *Review of Economic Dynamics*, 45:237–263.

Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995). Precautionary saving and social insurance. Technical report, National Bureau of Economic Research.

Hubener, A., Maurer, R., and Mitchell, O. S. (2016). How family status and social security claiming options shape optimal life cycle portfolios. *The Review of Financial Studies*, 29(4):937–978.

Huckfeldt, C. (2016). Understanding the scarring effect of recessions. *Working Paper*.

Huff Stevens, A. (1997). Persistent effects of job displacement: The importance of multiple job losses. *Journal of Labor Economics*, 15:165–188.

Hurd, M. D. and McGarry, K. (1995). Evaluation of the subjective probabilities of survival in the health and retirement study. *The Journal of Human Resources*, 30:S268–S292.

Hurd, M. D. and McGarry, K. (2002). The predictive validity of subjective probabilities of survival. *The Economic Journal*, 112:966–985.

Hurd, M. D., Smith, J. P., and Zissimopoulos, J. M. (2004). The effects of subjective survival on retirement and social security claiming. *Journal of Applied Econometrics*, 19:761–775.

Imrohoroğlu, S. and Kitao, S. (2012). Social security reforms: Benefit claiming, labor force participation, and long-run sustainability. *American Economic Journal: Macroeconomics*, 4(3):96–127.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *The American Economic Review*, 83:685–709.

Jones, J. B. and Li, Y. (2018). The effects of collecting income taxes on social security benefits. *Journal of Public Economics*, 159:128–145.

Jones, J. B. and Li, Y. (2023). Social security reform with heterogeneous mortality. *Review of Economic Dynamics*, 48:320–344.

Lockwood, L. M. (2012). Bequest motives and the annuity puzzle. *Review of economic dynamics*, 15(2):226–243.

MacCuish, J. H. (2024). Costly attention and retirement. *Working Paper*.

Maurer, R., Mitchell, O. S., Rogalla, R., and Schimetschek, T. (2021). Optimal social security claiming behavior under lump sum incentives: Theory and evidence. *Journal of Risk and Insurance*, 88(1):5–27.

Medical Expenditure Panel Survey (2022). Medical expenditure panel survey, public use data. https://www.meps.ahrq.gov/mepsweb/. Rockville, MD: U.S. Department of Health and Human Services, Agency for Healthcare Research and Quality.

Meyer, W. and Reichenstein, W. (2010). Social security: When to start benefits and how to minimize longevity risk. *Journal of Financial Planning*, 23(3):49–59.

Miller, R. and Bairoliya, N. (2021). Health, longevity, and welfare inequality of older americans. *The Review of Economics and Statistics*, pages 1–45.

O'Dea, C. and Sturrock, D. (2023). Survival pessimism and the demand for annuities. *The Review of Economics and Statistics*, 105(2):442–457.

Panel Study of Income Dynamics (2022). Panel study of income dynamics, public use dataset. https://doi.org/10.3886/ICPSR07451.v53. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI.

Pashchenko, S. and Porapakkarm, P. (2024). Accounting for social security claiming behavior. *International Economic Review*, 65(1):505–545.

Pijoan-Mas, J. and Ríos-Rull, J.-V. (2014). Heterogeneity in expected longevities. *Demography*, 51:2075–2102.

Rohwedder, S. and Kleinjans, K. J. (2006). Dynamics of individual information about social security. *Unpublished manuscript, RAND*.

Rust, J. and Phelan, C. (1997). How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica: Journal of the Econometric Society*, pages 781–831.

Rutledge, M. S., Sass, S. A., and Ramos-Mercado, J. D. (2017). How does occupational access for older workers differ by education? *Journal of Labor Research*, 38(3):283–305.

Shepard, M. (2011). Social security claiming and the annuity puzzle. *Working Paper*.

Shoven, J. B. and Slavov, S. N. (2014a). Does if pay to delay social security? *Journal of Pension Economics and Finance*, 13(2):121–144.

Shoven, J. B. and Slavov, S. N. (2014b). Recent changes in the gains from delaying social security. *Journal of Pension Economics and Finance*, 27(3):32–41.

SSA Trustees Report (1996). 1996 annual report of the board of trustees of the federal old-age and survivors insurance and disability insurance trust funds.

Sun, W. and Webb, A. (2011). Valuing the longevity insurance acquired by delayed claiming of socialsecurity. *Journal of Risk and Insurance*, 78:907–930.

Understanding America Study (2020). Understanding america study, public use dataset. https://uasdata.usc.edu. Center for Economic and Social Research, University of Southern California.

Van der Klaauw, W. and Wolpin, K. I. (2008). Social security and the retirement and savings behavior of low-income households. *Journal of Econometrics*, 145(1):21–42.

Venti, S. and Wise, D. A. (2014). The long reach of education: Early retirement. *NBER Working Paper No. 20740*.

Yu, Z. (2024). Why are older men working more? the role of social security. *Journal of Public Economics*, 231:105071.

Online appendix

A Present-discounted value analysis

We extend the present-discounted value (PDV) analysis from Section 2 by examining how the PDV-maximizing Social Security claiming age varies along two dimensions: (1) the age gap between spouses in married couples who claim simultaneously (or at age 62 if the spouse is not yet eligible), and (2) subjective beliefs about the early claiming penalty.

Figure A.1 illustrates how the PDV-maximizing claiming age changes with lifespan for couples with age gaps of 0, 1, 4, and 8 years (Panels (a)–(d), respectively), where a positive gap indicates the husband is older. The results show that age gaps substantially influence the benefit maximizing claiming age due to the asymmetric treatment of spousal and survivors benefits.

Spousal benefits—up to 50% of the head's primary insurance amount (PIA)—can only be claimed after the head has claimed, are penalized for early claiming, and do not receive delayed retirement credits. As a result, larger age gaps reduce the relevance of spousal benefits for early claiming. For instance, with an 8-year gap (Panel (d)), the spouse isn't eligible before the head turns 70, making the PDV-maximizing age similar to that of a single individual. With no age gap (Panel (a)), both spouses are penalized equally for early claims, but only the head benefits from delayed credits, making present value maximizing age flat at 65 for all lifespans. For intermediate gaps (Panels (b) and (c)), the trade-off is more nuanced: the younger spouse's ineligibility alters the penalty structure and lowers the PDV-maximizing age for longer lifespans compared to singles, as delayed spousal benefits are not rewarded.

While survivors benefits do not directly alter the PDV for a given lifespan, they extend the effective duration over which benefits are received. For example, a household head living to the average married male life expectancy of 82.2 years has a spouse who, on average, lives 8.2 years longer. This extended horizon (marked by the green dashed line in Figure A.1) raises the benefit maximizing claiming age, since PDV increases with lifespan.

Figure A.2 turns to how beliefs about the size of the early claiming penalty (e.g., 0–1%, 6.7%, or 15%) affect the PDV-maximizing age. Across all beliefs, the PDV maximizing claiming age increases with lifespan due to the fundamental trade-off between larger benefits and earlier receipt. However, the level of the penalty shifts the curve: lower perceived penalties reduce the cost of early claiming, lowering the PDV maximizing age, while higher perceived penalties raise the cost of early claims, pushing the PDV maximizing age upward.

We also compute the value of benefits (total and PDV) lost by our 1931-1935 cohort by following the claiming behavior observed in the data rather than claiming at the PDV-maximizing claiming age. Table A.1 shows the size of benefits forgone at each claiming age; a value of 0 is forgone for claiming at the PDV-maximizing claiming age. The average forgone benefits is computed by taking the weighted average of these numbers. The claiming shares in Figure 2 are used as weights.

B Empirical analysis of subjective beliefs

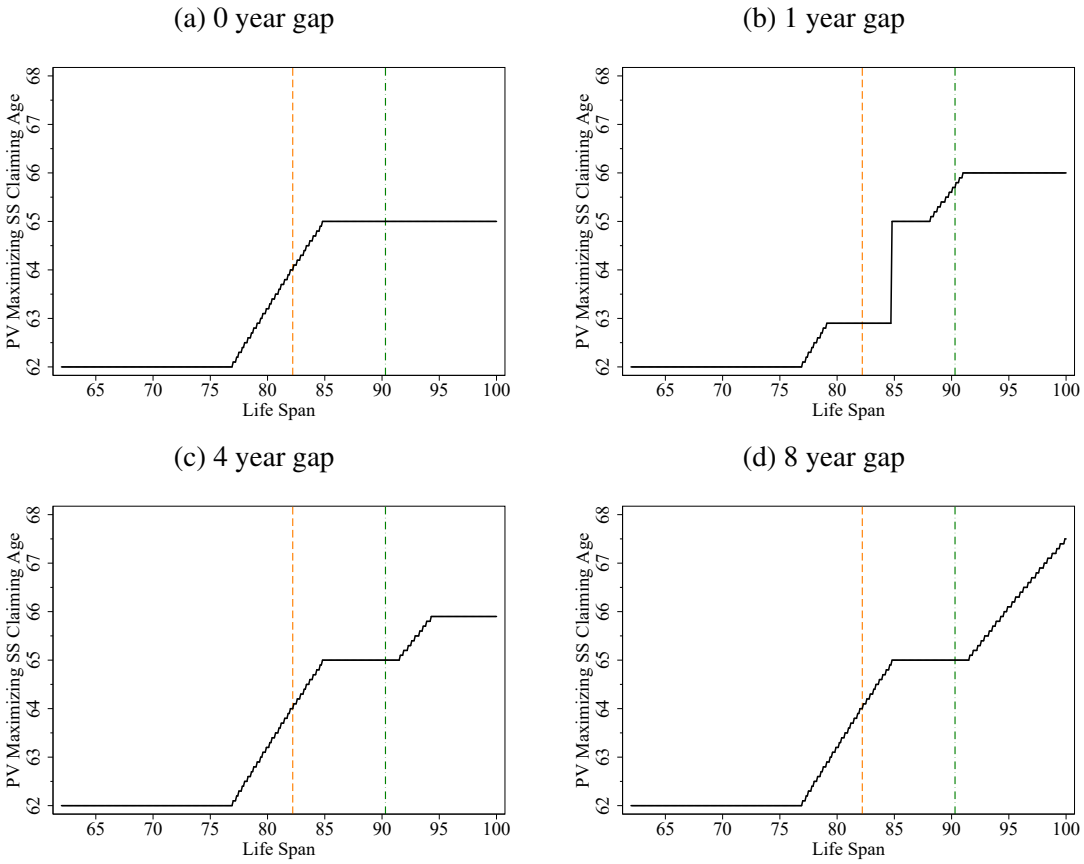
Section 2.2 contains empirical evidence studying the prevalence of subjective beliefs and linking these beliefs to Social Security claiming age by estimating regression Equation 1.

Table A.1: Forgone Benefits by Claiming Age

	AVERAGE	NON-COLLEGE	COLLEGE	SINGLE	MARRIED
Age 62 life expectancy	81	78	83	78	82
PDV-maximizing age	64	63	64	62	64
PDV of Forgone Benefits (1998 $)					
Age 62	2,580	474	4,664	474	3,522
Age 63	605	0	1,762	0	1,076
Age 64	0	896	229	896	0
Age 65	698	3,095	0	3,095	227
Age 66	22,343	23,240	22,573	23,240	22,343
Age 67	42,102	41,497	43,259	41,497	42,573
Age 68	60,055	57,948	62,139	57,948	60,996
Age 69	76,276	72,668	79,288	72,668	77,689
Age 70	90,841	85,732	94,780	85,732	92,724
Average	3,526	2,619	4,742	3,165	3,883
Total Forgone Benefits (1998 $)					
Age 62	10,517	3,527	18,372	3,527	12,248
Age 63	4,393	0	10,517	0	5,258
Age 64	0	-1,796	4,393	-1,796	0
Age 65	-2,662	-1,860	0	-1,860	-3,527
Age 66	24,109	22,313	28,501	22,313	24,109
Age 67	49,147	44,755	55,271	44,755	50,013
Age 68	72,455	65,465	80,310	65,465	74,186
Age 69	94,031	84,445	103,618	84,445	96,628
Age 70	113,876	101,693	125,194	101,693	117,339
Average	7,046	3,058	12,140	3,362	7,647

Notes: Age 62 life expectancy and PDV-maximizing claiming age are annual versions of information in Table A.1. Forgone benefits (both PDV and Total) are calculated relative to the benefits received for claiming at the PDV-maximizing age. Negative values in total foregone benefits reflect examples of calculations in which the claiming age that maximizes the PDV of benefits is different from the claiming age that maximizes total benefits received.

Figure A.1: Variation in the PDV maximizing claiming age by spousal age gap



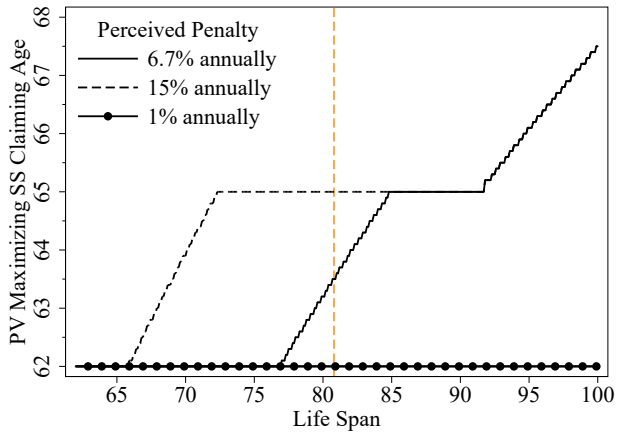
Notes: Figures show the claiming age that maximizes the present discounted value (PDV) of total benefits by lifespan, for spousal age gaps of 0 (Panel a), 1 (b), 4 (c), and 8 years (d). PDVs are calculated for a worker earning $50,000 annually, discounted at 3% per year. The dashed orange line marks the average life expectancy of a married man at age 62 (82.2 years); the dashed green line marks the average life expectancy of his wife at that age (90.3 years).

For this analysis we use two data sets: the Health and Retirement Study and the Understanding America Study. When we consider the HRS, we use a sample of 1,160 individuals born between 1930 and 1960 for whom all co-variates are non-missing. Table B.2 shows summary statistics for the HRS sample of those born 1930-1960 and our regression sample which includes all other co-variates. When we consider UAS, we use a sample of 3,648 observations for 2,339 individuals who are not already receiving Social Security benefits (details of sample are in Table E.5).

We first establish descriptive statistics related to subjective beliefs. Subjective survival is measured from the HRS while we use both HRS and the Understanding American Study to demonstrate subjective program knowledge.

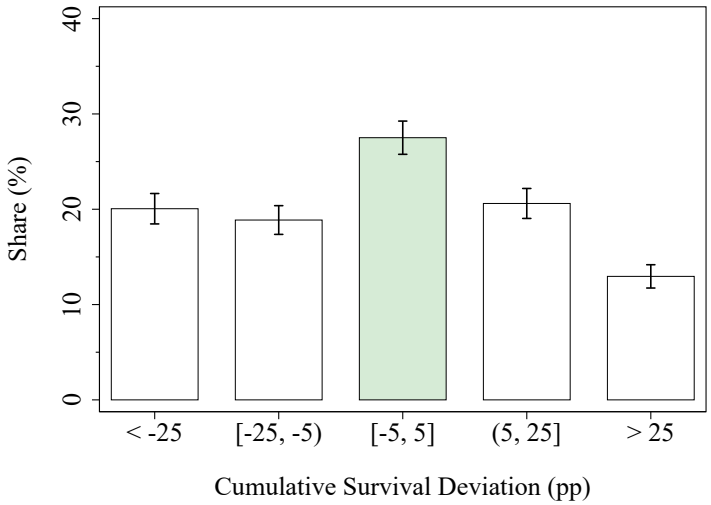
Subjective mortality expectations are measured from the Health and Retirement Study (HRS). One module of the survey asks respondents what they believe is the probability they will live to

Figure A.2: Variation in PDV maximizing claiming age by perceived penalty



Notes: Figures show the PDV-maximizing claiming age by lifespan under three perceived early claiming penalties: 6.7% (accurate), 15%, and 1%. PDVs are calculated for an individual earning $50,000 annually, discounted at 3%. Dashed orange line indicates average male life expectancy at age 62: 80.8 years.

Figure B.1: Distribution of Survival Beliefs



Source: Health and Retirement Study (2016, v1), authors' calculations.

Notes: figure displays the share of the sample who reported deviations in cumulative subjective survival to age 75 in each bin. Deviations are defined the difference between the subjective cumulative probability of survival and the average objective cumulative probability for a man of the same education and marital status.

certain ages: *(What is the percent chance) that you will live to be [X] or more?*[41]

We use a question in which workers are asked at age 60 about their probability of survival

[41] At the outset, respondents are prompted:*Next we would like to ask your opinion about how likely you think various events might be. When I ask a question, please give a number from 0 to 100, where "0" means you think there is absolutely no chance, and "100" means you think the event is absolutely sure to happen.*

to age 75 and compute the gap between this report and the objective probability of survival from age 60 to age 75 for a man of the same education and martial status. We cateogrize the reported cumulative probability deviations in five categories: (1) <-25pp, (2) [-25,5) pp, (3) [-5,5] pp, (4) (5,25] pp, and (5) >25 pp. The histogram showing the distribution of responses over these groups is shown in Figure B.1.

We document the presence of subjective beliefs about the early claim penalty in both the HRS and the UAS. These datasets ask different questions that each have pros and cons.

The Health and Retirement Study asks:
If you were to retire at age XX, and collect benefits, what do you think you Social Security benefits would be in today's dollars?

Using the response to this question at ages 62 and 65, we are able to construct the repondent's perceived annual adjustment. This question has the advantage of eliciting a continuous value for the response. Respondents are not confined to limited choices. However, as the question asks about the level of benefits, we cannot be sure that responses reflect knowledge of the adjustment or other factors that may impact nominal benefits.

The Understanding America Study, on the other hand, asks:
Imagine an individual, Mr. Spencer Wills, who is retired from work and turning 63 today. Spencer expects his monthly Social Security retirement benefit to be about $2,000 if he claims now. How much should he expect his monthly benefit to be if he decides to wait another year and claim when he is 64?

(a) it would stay the same

(b) approximately $2,120 (6% higher)

(c) approximately $2,020 (1% higher)

(d) approximately $2,300 (15% higher)

This question provides a clean setting. It is clear that the response to the question identifies how the individual believes benefits adjust due to the claiming age. However, this question does confine responses to the fours choices given. Additional information on the subjective program knowledge is discussed in Appendix F.8.

The pros and cons of these datasets are contained in Table B.1.

Finally, using the Understanding America Study, we also measure uncertainty of response. We use the following question:

Below is a list of aspects of the Social Security system. For each one do you feel you are very, somewhat, not too, or not at all knowledgeable about that aspect?
How Social Security retirement benefits change based on how old you are when you claim.

Here we detail the construction of variables used in this regression. The main variables of interest are the categorical variables defining subjective beliefs:

Subjective survival beliefs are defined as the deviation between an individual's reported probability of surviving to age 75 (conditional on survival to age 60) and the corresponding objective

For example, no one can ever be sure about tomorrow's weather, but if you think rain is very unlikely tomorrow, you might say there is a 10 percent chance. If you think there is a very good chance it will rain, you might say there is an 80 percent chance.

Table B.1: Health and Retirement Study vs. Understanding America Study

Health and Retirement Study (HRS)	
Pros	**Cons**
Includes measure of subjective survival expectations	Asks nominal values of expected benefits rather than expectations of SS rules
Panel structure linking beliefs to future claiming behavior	Only covers those age 50+
	Benefit expectation questions begin in 2002
Understanding America Study (UAS)	
Pros	**Cons**
Ask a clean question about a hypothetical retired individual delaying claims from age 63 to age 64	Lacks the long panel structure (only 2 waves)
Includes additional questions on program knowledge, confidence	Does not include subjective survival
Covers those early in the life-cycle	Covers years 2017-2022

probability for their education-marital status group in the HRS. These objective cumulative probabilities of survival are estimated by the authors using annual group specific survival probabilities estimated from HRS (additional details on the estimation of these probabilities is included in Section 4.1.1 and in Appendix F.1); these probabilities are estimated to be 0.429 for non-college educated singles, 0.621 for college educated singles, 0.620 for non-college educated married men, and 0.772 for college educated married men. This measure captures the extent, though not the source, of belief deviation. We classify deviations from group specific objective survival into five categories: (1) large negative (> 25 percentage points (pp) below objective), (2) small negative (5–25 pp below), (3) accurate (within ± 5 pp), (4) small positive (5–25 pp above), and (5) large positive (> 25 pp above). the distribution of beliefs across these five groups is shown in Figure B.1.

Beliefs about the penalty for early claiming are constructed from subjective reports of the expected size of future benefits would be for various claiming ages. Beginning in wave 6, the HRS began asking questions about how large a respondent expected benefits to be at various claiming ages. Using reports for ages 62 and 65, we calculate the percentage difference between these values and divide by three. This gives an estimate of how much the respondent believes their benefits change per year between these two dates. As both dates are prior to the normal retirement age we interpret this as the subjective early claiming penalty. For example, if a respondent reported

Table B.2: Sample Summary Statistics

	HRS Sample (born 1930-1960)		Estimation Sample (born 1930-1960) + non-missing covariates	
	mean	sd	mean	sd
SS claiming age	63.525	1.767	63.808	1.762
subjective survival (normalized)	2.084	25.444	0.559	24.899
subjective penalty belief (normalized)	-0.010	0.047	-0.012	0.047
male share	0.498	0.500	0.633	0.482
non-college, single share	0.114	0.318	0.121	0.326
college, single share	0.084	0.277	0.194	0.395
non-college, married share	0.458	0.498	0.287	0.453
college, married share	0.345	0.475	0.398	0.490
share spousal age gap 0 years	0.119	0.324	0.115	0.319
share age gap 1-2 years	0.214	0.410	0.234	0.423
share age gap 3-6 years	0.266	0.442	0.254	0.435
share age gap 7-10 years	0.106	0.308	0.083	0.276
share working at age 62	0.574	0.494	0.727	0.446
share in good health at age 62	0.879	0.326	0.922	0.268
log of real wealth at age 62	2.736	1.637	2.685	1.696
share holding a private pension	0.453	0.498	0.539	0.499
normal retirement age	65.329	0.424	65.882	0.198
Observations	12309		1142	

Source: Public-Use Health and Retirement Study (2016, v1), authors' calculations.

expected benefits of $100 monthly if benefits are claiming at 62 and $118 monthly if claimed at 65, our method would predict an implied subjective annual penalty of 6% per year. We then use this implied annual penalty to classify beliefs into 5 categories:(1) -1.5–1.5% (no penalty or a small credit for claiming early), (2) 1.5-4.7%, (3) accurate (within $\pm$ 2 pp), (4) 8.7–12.0%, (5) $\geq$ 12.0% (roughly 5 pp or more above the true penalty).

We estimate various regression specifications which add additional control variables that will appear in our quantitative analysis. We will perform these regressions both the case of a continuous measure of beliefs as well as the two categorical measures of beliefs: the baseline specification and a second in which we categorize beliefs into quintiles. The results of the regressions are included in Tables B.3 -B.6 and Figures B.2 and B.3.

First, we include cohort controls variables including gender, the normal retirement age, and cohort dummies. Gender is a dummy variable taking a value of 1 if the respondent is female. Normal retirement age and cohort controls are meant to control for policy as well as other cohort effects that may impact retirement age. The variable for the NRA is a continuous variable ranging from 65 to 67 and captures the normal retirement age for the respondent based upon year of birth. As the normal retirement age increased in 2 month increments, this variable also takes fractional values (e.g. 65.5 would signify that the NRA for this individuals was 65 year and 6 months of age).

This variable controls for policy changes over time. We also include dummy variables for 10-year cohort bins (birth years 1931-1940, 1941-1950, and 1951-1960). These additional cohort dummies are meant to capture any additional time trends (such as social norms or changes in employment policies) which may have impacted the claiming decision of households.

In our second specification, we include additional controls for demographic groups. Respondents are grouped into four categories based upon education and marital status: non-college educated singles, college educated singles, non-college educated married, and college educated married. College is defined as any schooling beyond high school (not restricted to have earned a degree). An individual is classified as married if the respondent is in a legal marriage (spouse present or absent) or reports being partnered. This will be our chosen specification for the empirical analysis.

Our final specification controls for other factors which have been studied as determinants of retirement and claiming behavior including labor force participation, self-reported health status, wealth, the age gap between spouses, and private pension access. Labor force participation, self-reported health status, and wealth are measured between ages 60-63 to capture the impact of these variables in the years prior to claiming. Labor force participation is a dummy variable taking a value of 1 if the respondent reports currently having a full-time or part-time job. Self-reported health takes two values. The respondent is classified as having good health if they report their own health as "Excellent", "Very Good" or "Good" and classified as having bad health if they report their own health as "Bad" or "Very Bad". Wealth measures include all financial wealth plus housing and is converted into 1998 dollars. The age gap between spouses is computed as the differences in ages between members of a married couple and discretized into four categories: (1) no age gap, (2) age gap of 1-2 years, (3) age gap of 3-6 years, or (4) age gap of 7-10 years. Negative age gaps (wife is older) and those more than 10 years are excluded.[42] Finally, the variable for private pension access is a dummy variable taking a value of 1 if the respondent reports having either a defined benefit or a defined contribution pension plan. Because early life-cycle subjective beliefs likely also impact the value of these controls as agents near retirement, these will not be our baseline specifications. However, these will be useful for robustness.

As a final robustness, we categorize between into quintiles rather than our chosen groups. This equally splits respondents over these groups. The average belief within each of these quintiles for both types of beliefs are shown in Table B.5.

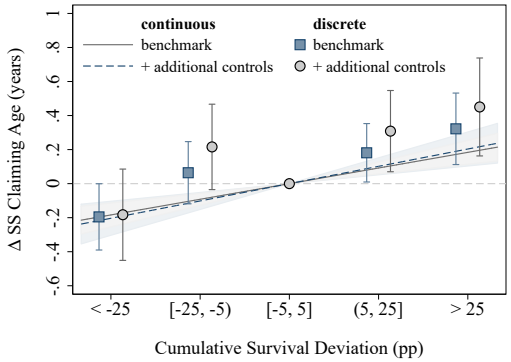
C Spousal and survivors benefits

Spousal benefits The early application penalty applies to both the old-age benefits of the head of household and the spousal benefits while delayed retirement credits are only awards to the head's benefits. Therefore, in practice, the benefits received at a household level (old age benefits + spousal benefits) are a function of the claiming age of both the head and the spouse. In order to reduce the computational burden of our model, we abstract from a separate claiming decision of the spouse and assume that all couples claim together. In the event that the spouse is not benefit eligible at the time of the head's claiming, we assume that spousal benefits are received as soon as the spouse is age 62. Therefore, spousal claiming age in the model—and thus any early claiming

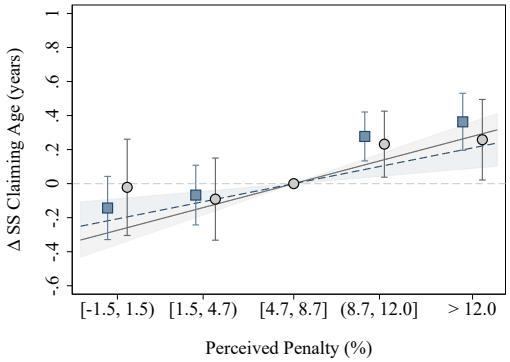
[42] This restriction lines up with how age gaps are discretized in the quantitative model. The criterion is much more restrictive for women in the sample. This is apparent in Table B.2 as the male share increases from 0.484 to 0.634

Figure B.2: Impact of subjective beliefs on claiming age
regression coefficients (additional controls)

(a) Survival beliefs (β_2)



(b) Program knowledge beliefs (β_3)

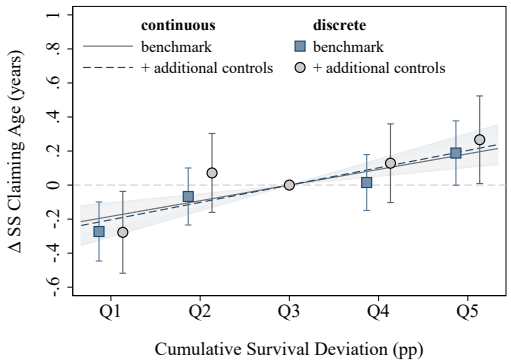


Source: Health and Retirement Study (2016, v1), authors' calculations.

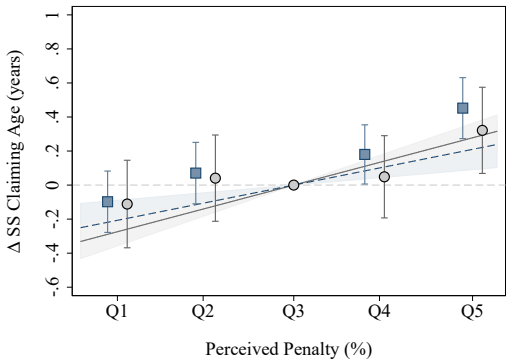
Notes: Figure shows regression coefficients from Equation 1 under two specifications. The first includes cohort and demographic controls (gender, education, marital status, normal retirement age, cohort); the second adds labor force participation, self-reported health, log real wealth prior to claiming, spousal age gap (if married), and private pension status. Panel (a) reports coefficients for survival belief dummies; Panel (b) for program knowledge belief dummies. The omitted (accurate belief) category is centered at 0. 90% confidence intervals shown.

Figure B.3: Impact of subjective beliefs on claiming age
regression coefficients (quintiles)

(a) Survival beliefs (β_2)



(b) Program knowledge beliefs (β_3)



Source: Health and Retirement Study (2016, v1), authors' calculations.

Notes: Figure shows regression coefficients from Equation 1 under two specifications. The first includes cohort and demographic controls (gender, education, marital status, normal retirement age, cohort); the second adds labor force participation, self-reported health, log real wealth prior to claiming, spousal age gap (if married), and private pension status. Panel (a) reports coefficients for survival belief dummies; Panel (b) for program knowledge belief dummies. The omitted (accurate belief) category is centered at 0. 90% confidence intervals shown.

penalty incurred on spousal benefits received—is completely determined by head's claiming age and the age gap between the spouse and the head of household.

Survivors benefits There are two important rules pertaining to survivors benefits that we try to replicate in the model. First, individuals can leave behind their Social Security benefits for their surviving spouses (even if they have not claimed their benefits at the time of the death). We allow this feature in the model as well. In the scenario where the head dies without claiming their benefits, survivors benefits in the model are calculated using the Primary Insurance Amount at the time of death. However, if the head has already claimed their benefits at the time of death, then survivors benefits are calculated using any adjustments that were applied to the head's benefits. Second, in practice, surviving spouses can receive reduced benefits as early as age 60 (age 50 in the case of disability and even earlier if they have not remarried and are caring for the deceased husband's child who is under the age of 16). For computational simplicity we abstract from some of these details and allow full survivors benefits to be bequeathed as long as the spouse is age 62.

D Recursive formulation

The first is the employment phase (ages 25–61), during which individuals choose consumption, savings, and labor supply. The second phase is the retirement choice phase (ages 62–69), where individuals continue to make consumption, labor supply and savings decisions and also decide whether to claim Social Security benefits. The third phase, beginning at age 70, is the retirement phase, in which individuals no longer work or claim benefits and make only consumption and saving decisions. The decision problem of a household head with education level e, marital status m, and Social Security belief type k for each phase is given below:

Employment phase

$$V_{e,q,k}(a_t, a_t^{ss}, a_{st}^{ss}, \eta_t, \lambda_t, \mu_t) = \max_{\{c_t, h_t\}} \Bigg\{ U^{e,m}(c_t, l_t) + \beta^{e,m} \pi_{t+1}^{s} \Big[EV_{e,q,k}(a_{t+1}, a_{t+1}^{ss}, a_{st+1}^{ss}, \eta_{t+1}, \lambda_{t+1}, \mu_{t+1}) \Big] + \beta^{e,m}(1 - \pi_{t+1}^{s}) beq^{e,m}(A_{t+1}^{q}) \Bigg\} \quad s.t.$$

$$a_{t+1} = a_t + W(y_t, y_{st}, \bar{r}a_t, \tau) + tr_t - c_t,$$
(2), (6-13), and (16).

where $y_t + y_{st} + \bar{r}a_t$ is the total pre-tax income and $W(., \tau)$ gives the level of post-tax income with the tax rate τ. Note that the expectation is taken with respect to wage, employment, SS program belief, and health uncertainty.

Retirement choice phase Starting at age 62, individuals also make benefit-claiming decisions, which are one-time choices. Benefits are determined by the age at which individuals first claim. When making these decisions, individuals consider the perceived benefits implied by their belief type:

$$V_{e,q,k}(a_t, a_t^{ss}, a_{st}^{ss}, \eta_t, \lambda_t, \mu_t, b_{t-1}^{ss}) = \max_{\{c_t, h_t, b_t^{ss}\}} \Bigg\{ U^{e,m}(c_t, l_t) + \beta^{e,m} \pi_{t+1}^{s} \Big[EV_{e,q,k}(a_{t+1}, a_{t+1}^{ss}, a_{st+1}^{ss}, \eta_{t+1}, \lambda_{t+1}, \mu_{t+1}, b_t^{ss}) \Big] + \beta^{e,m}(1 - \pi_{t+1}^{s}) beq^{e,m}(A_{t+1}^{q}) \Bigg\} \quad s.t.$$

$$a_{t+1} = a_t + W\left(y_t, y_{st}, \bar{r}a_t, \tau\right) + \mathbb{I}\left\{b_t^{ss} = 1\right\} \times \left(\overline{ssb}_t^{k} + \mathbb{I}\left\{t - \iota \geq 62\right\} \times \overline{ssb}_t^{sk}\right) + tr_t + pb_t - c_t,$$

(2), (6-13), and (16).

Note that the expectation is taken with respect to wage, employment, SS program belief and health uncertainty.

Retired phase At age 70, if an individual has not yet claimed benefits, they automatically begin receiving both their own benefits and any applicable spousal benefits (if applicable, as some younger spouses may not yet be SS eligible).

$$V_{e,q}(a_t, a_t^{ss}, a_{st}^{ss}, \mu_t) = \max_{c_t} \Bigg\{ U^{e,m}(c_t, l_t) + \beta^{e,m} \pi_{t+1}^{s} EV_{e,q}(a_{t+1}, a_{t+1}^{ss}, a_{st+1}^{ss}, \mu_{t+1}) + \beta^{e,m}(1 - \pi_{t+1}^{s}) beq^{e,m}(A_{t+1}^{q}) \Bigg\} \quad s.t.$$

$$a_{t+1} = a_t + W(y_{st}, \bar{r}a_t, \tau) + ssb_t + \mathbb{I}\left\{t - \iota \geq 62\right\} \times ssb_{st} + tr_t + pb_t - c_t,$$

(2), (7), and (16).

Note that in this phase of life, the expectation is taken with respect to only health uncertainty.

E Data: PSID, HRS, MEPS, and UAS

Panel Study of Income Dynamics (PSID) The description of the PSID data and collection methodology in this paper is based on material made publicly available by the study here: https://psidonline.isr.umich.edu/default.aspx.[43] The PSID is a nationally representative longitudinal survey in the United States that we use to estimate life-cycle profiles of labor force participation, hours, wealth, the wage process, and initial conditions. The original PSID sample was drawn from the nationally representative SRC sample and an oversample of the low-income SEO sample. We use a subsample of individuals from the SRC sample who were interviewed at least twice between

[43]The collection of data used in this study was partly supported by the National Institutes of Health under grant number R01 HD069609 and R01 AG040213, and the National Science Foundation under award numbers SES 1157698 and 1623684.

1968 and 2017. Our analytic sample consists of male household heads between ages 22 and 99, born between 1902 and 1965. The final sample includes 95,534 observations for 5,593 individuals. Table E.1 provides details on how we derive our sample.

Health and Retirement Study linked with SSA administrative records The Health and Retirement Study (HRS) is a longitudinal survey of Americans over age 50.[44]

In this work, we use HRS data (2016, v1) linked to SSA administrative records on benefits and earnings. This dataset is used to compute claiming age distributions and related statistics, spousal benefits within couples, and the evolution of spousal earnings and Average Indexed Monthly Earnings (AIME) of household heads and spouses over the life cycle.

We use three different samples: two derived from the Respondent Cross-Year Benefits file (male and spouse samples) and one from the Summary Earnings file (household earnings sample). The male sample, used to construct claiming behavior, consists of 2,296 men born between 1926 and 1939 who first claimed retirement benefits between ages 62 and 70 (refer to Table E.2 for details). We focus on this cohort to ensure a sufficiently large sample while including only individuals for whom Social Security was in place throughout their working lives and the normal retirement age was 65. The spouse sample includes 254 married couples in which the male household head was born between 1921 and 1939 and (i) claimed Social Security retirement benefits on his own earnings record between ages 62 and 70, and (ii) the spouse first claimed retirement benefits either solely on her own record, jointly on her own and her spouse's record, or exclusively as auxiliary benefits on her spouse's record (refer to Table E.3 for details).

The household earnings sample contains 72,545 observations for 1,270 households in which the male head was born between 1926 and 1940 and spousal earnings are non-missing (including zeros). Refer to Table E.4 for sample selection details.

Medical Expenditure Panel Survey (MEPS) We use the Household Component of Medical Expenditure Panel Survey (MEPS-HC) to identify health and mortality related parameters. MEPS-HC is a nationally representative survey of the U.S. civilian noninstitutionalized population. The sampling frame is drawn from respondents to the National Health Interview Survey (NHIS), which is conducted by the National Center for Health Statistics. We use a pooled sample (between the years 1996 and 2014) of males between the ages of 20 and 90 which gives us a total sample size of 89,870 observations.

Understanding America Study (UAS) Our project relies on data from surveys administered by the Understanding America Study (UAS), 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 UAS is a panel dataset of roughly 9,000 respondents per wave, designed to be nationally representative of the United States and administered in various modules. The panel is conducted online, allowing respondents to participate digitally at their convenience. We use this dataset to measure how well individuals understand the Social Security system. Specifically, we use a nationally representative sample of men ages 25–69 who are not currently receiving Social Security

[44]The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

retirement benefits, drawn from two waves of the study: UAS94 and UAS231. This sample includes 3,663 observations for 2,347 individuals collected between 2017 and 2022. After further restricting the sample to respondents who answered the subjective program knowledge questions, our final sample contains 3,648 observations for 2,339 individuals. Details on the sample selection process are provided in Table E.5.

F First-step estimation

F.1 Health and objective survival

We estimate health transitions using the Household Component of Medical Expenditure Panel Survey on a population between the ages of 20 and 90. The coefficient estimates for these health transitions are provided in Table F.1.

To estimate objective survival probabilities, we use the HRS tracker files to obtain individuals' age at death, based either on the reported year of death or, when unavailable, the last known date alive. Because the HRS is conducted biennially, we expand the dataset to account for intervening years. Marital and health status between waves are imputed using information from the preceding wave, while age is imputed by incrementing the prior value by one year. Education and birth cohort are treated as time-invariant. Using this expanded dataset and reported ages at death, we construct an indicator capturing the precise timing of death. Columns 2 and 3 of Table F.2 presents a comparison of life expectancies at ages 50, 60, and 70 between those constructed from HRS estimated objective survival probabilities and Social Security Administration life tables.

F.2 Subjective survival

Subjective mortality expectations are measured from the Health and Retirement Study (HRS). One module of the survey asks respondents what they believe is the probability they will live to certain ages: *(What is the percent chance) that you will live to be [X] or more?*[45]

The target age [X] varies by respondents' current age. Respondents under age 65 are asked about survival to age 75. If they report a positive probability, they are then asked about survival to age 80. Respondents aged 65–69 are asked about age 80, those 70–74 about age 85, 75–79 about age 90, 80–84 about age 95, and 85–89 about age 100. Respondents 90 and older are not asked these questions.

These questions yield subjective survival expectations to ages 75–100. To maximize usable reports, we impose four assumptions: (1) If a respondent reports a 0% probability of survival to a given age, we assume they believe they have a 0% chance of survival to all later ages; (2) If they report a 100% chance of survival to a given age, they also believe they have a 100% chance of survival to all prior ages; (3) If the same probability is reported for two non-consecutive ages,

[45] At the outset, respondents are prompted:*Next we would like to ask your opinion about how likely you think various events might be. When I ask a question, please give a number from 0 to 100, where "0" means you think there is absolutely no chance, and "100" means you think the event is absolutely sure to happen.*

For example, no one can ever be sure about tomorrow's weather, but if you think rain is very unlikely tomorrow, you might say there is a 10 percent chance. If you think there is a very good chance it will rain, you might say there is an 80 percent chance.

the same probability is assumed for all ages in between. (4) All respondents believe there is a 0% chance of survival to age 110. We drop respondents with inconsistent answers (e.g., higher survival chances to older ages) or who report 0% or 100% for all ages.[46]

The final sample includes men ages 45–89 born between 1926 and 1965.[47] This yields 60,308 observations for 12,851 individuals: 7,265 ages 45–54, 27,624 ages 55–64, 17,007 ages 65–74, and 8,412 ages 75–89. We compute the median subjective survival probability to each target age by reference (current) age, education, marital status, and health using HRS cross-sectional weights.

Using this sample, we estimate survival curves that vary by education, marital status, and self-reported health. Because many survival curves can rationalize the reported cumulative probabilities, we follow O'Dea and Sturrock (2023) and assume subjective survival probabilities follow a Weibull distribution:

$$S_{\mu}^{e,m}(\alpha; X) = exp\left[-\left(\frac{X-\alpha}{\kappa_{\mu}^{e,m}}\right)^{\lambda_{\mu}^{e,m}}\right] \tag{19}$$

where $\kappa_{\mu}^{e,m}$ and $\lambda_{\mu}^{e,m}$ are the estimated parameters. This yields survival curves from age α to 110 for each education-marital-health combination.

To incorporate evolving health, we construct survival curves in two steps. First, for each reference age $\alpha = 50, ..., 70$, we estimate cumulative survival curves in rolling 10-year bins. For example, the curve for reference age 60 uses data from ages 55–65. This produces 21 curves per demographic-health group. Second, we compute average cumulative survival probabilities to each age, $\bar{S}_{\mu}^{e,m}(X = \alpha, ...; 110)$, by averaging across all curves that include that age:

$$\bar{S}_{\mu}^{e,m}(X = \alpha, ...; 110) = \begin{bmatrix} 1 \\ S_{\mu}^{e,m}(50; 51) \\ \vdots \\ \sum_{\alpha=50}^{X-1} \frac{1}{X-\alpha} S_{\mu}^{e,m}(\alpha; X) \end{bmatrix}$$

This method allows the cumulative probability of survival to vary by not only fixed education and marital status but also by healths shocks after age 50.[48]

Figures F.1 show estimated objective, $\bar{O}_{\mu}^{e,m}(X)$, and subjective survival curves, $\bar{S}_{\mu}^{e,m}(X)$. The solid lines on this graph represent the subjective cumulative probability of survival to various ages, conditional on survival to age 50, for individuals who vary based upon education, marital status, and self-reported health. The cumulative probability of survival to age X for someone in bad/good health represents the cumulative probability of survival for someone who received a bad/good health shock in each period from age 50 to age X. The dashed lines represent the objective

[46] Papers such as Bissonnette et al. (2017) treat 50% responses as "focal answers" reflecting rounding or guessing. We retain them because they remain informative for deviations from objective responses, especially in affecting the mean or median responses we use in estimation.

[47] We expand the cohort to ensure enough observations to estimate subjective survival curves at younger ages. Using only the 1931-1935 cohort would preclude estimation of subjective survival at reference ages 50 and 55. Additionally, using only this cohort would leave a small sample for which estimation of parameters as a function of demographics and health would not be feasible.

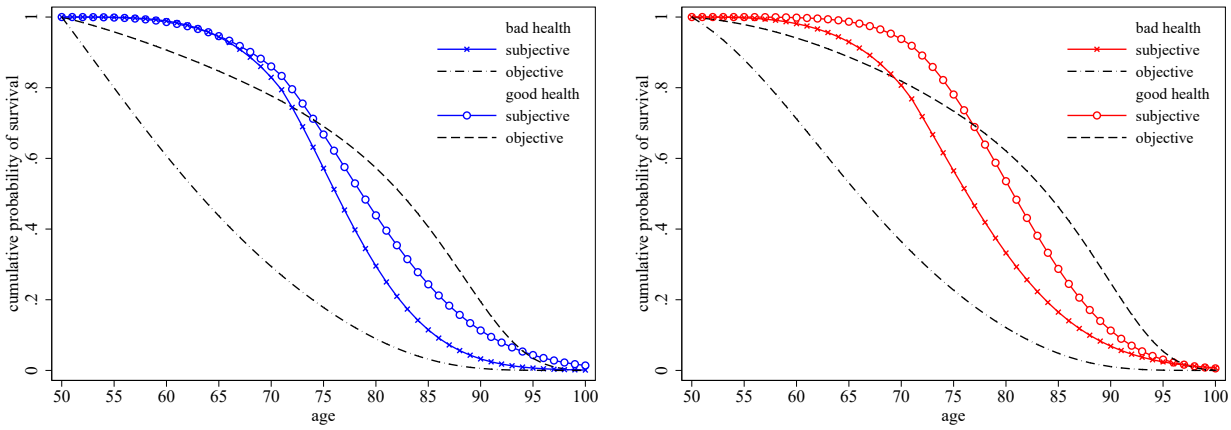
[48] Due to data limitations, we only estimate the age specific curves up until age 70 (using data on individuals up to age 75). The assumption we make is that shocks after age 70 do not impact your subjective perceptions of survival to various ages differently than the health shocks realized up until age 70.

cumulative probability of survival (as estimated from HRS using the process discussed in Section F.1) for these same groups. The gap between these lines captures deviations in subjective mortality from the mortality predicted by average objective survival probabilities.

Figure F.1: Objective $\left(\bar{O}_{\mu}^{e,m}\right)$ and subjective $\left(\bar{S}_{\mu}^{e,m}\right)$ cumulative survival curves by education, marital status, and self-reported health

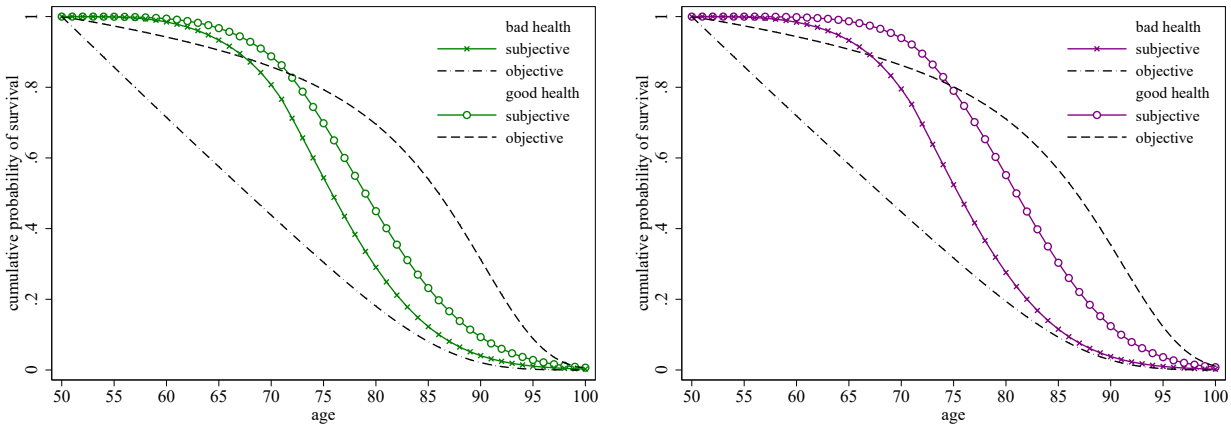
(a) Non-college, single

(b) College, single



(c) Non-college, married

(d) College, married



Source: Health and Retirement Study, authors' calculations

Notes: solid lines show the subjective survival curves of an individual of a given education-marital status group for bad (circle marked) and good (x marked) health. Dashed lines are the objective survival curves measured from HRS for the same education, marital status, and health groups.

We use these survival curves to compute the annual survival probabilities that will be used as inputs into our model. One option would be to back out annual survival probabilities directly from these curves as the ratio of the cumulative probability of survival to age $t+1$ to the cumulative probability of survival to age t. Specifically, $\pi_{t+1}^{s}\left(e,m,\mu\right) = \bar{S}_{\mu}^{e,m}\left(t+1\right)/\bar{S}_{\mu}^{e,m}\left(t\right)$. However, there are many options of annual survival probabilities that can deliver the same cumulative survival curve. Thus, we choose to use the structure of the estimated objective survival probabilities to guide the shape of the probabilities. Specifically, we will construct a scale factor which will

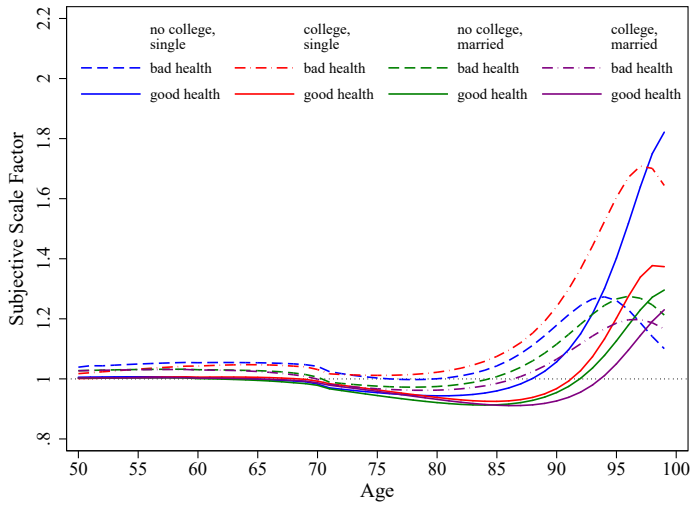
augment objective survival probabilities to make the dashed objective survival curves by health exactly match the estimated subjective survival curves.[49]

The objective cumulative probability of survival to age t, $\bar{O}_{\mu}^{e,m}(t)$, is equal to the product of the objective cumulative probability of survival to age $t-1$, $\bar{O}_{\mu}^{e,m}(t-1)$, and the probability of survival between ages $t-1$ and t. We compute a education, marital status, and health specific scale factor, $\gamma_t^{e,m,\mu}$, such that the cumulative objective probability of survival is equal to the cumulative subjective probability of survival estimated from HRS:

$$\bar{O}_{\mu}^{e,m}(t) = \left[\left(1+\gamma_j^{e,m,\mu}\right)\pi_t(e,m,\mu)\right] * \bar{O}_{\mu}^{e,m}(t-1) \quad \text{s.t. } \bar{O}_{\mu}^{e,m}(t) = \bar{S}_{\mu}^{e,m}(t)$$

This scale factor is shown in Figure F.2.

Figure F.2: Subjective scale factor for annual survival probabilities



Source: Health and Retirement Study, authors' calculations

Notes: shows the computed scale factor which equates objective and subjective cumulative probabilities of survival.

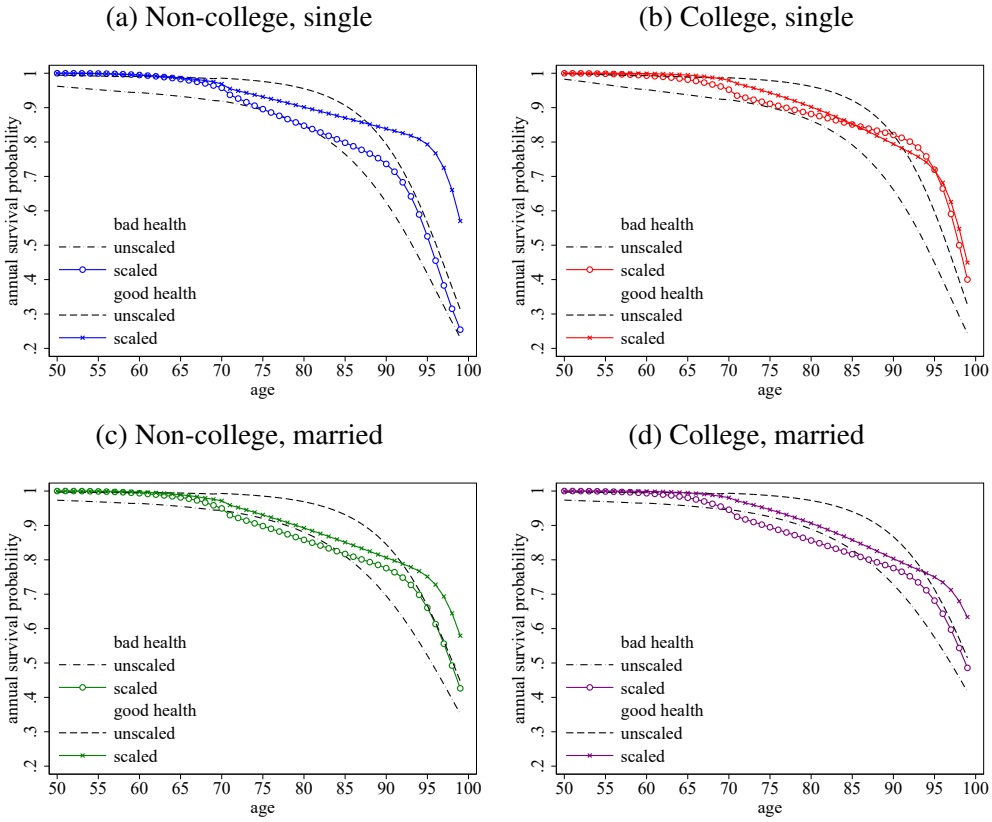
Comparison between subjective and objective survival We define the life expectancy gap as the difference between the subjective life expectancy (implied by the estimated subjective survival probabilities) and the measured life expectancy derived from education- and marital-status-specific survival probabilities (and the SSA life tables):

$$\text{life expectancy gap} = \text{subjective life expectancy} - \text{measured life expectancy}$$

As Table F.3 and Figure F.3 show, both subjective and objective annual survival probabilities and, thus, life expectancy vary across education-marital groups. Data on observed deaths reveal higher objective life expectancy among college-educated and married men, while subjective expectations are also higher for the college educated. Importantly, variation in these measures contributes

[49] Because annual probabilities of survival must be bounded above by 1, in some cases it may be impossible to find a scale factor that exactly matches objective and subjective probabilities.

Figure F.3: Subjective and objective survival probabilities
by education, marital status, and health



Source: Health and Retirement Study (2016, v1), authors' calculations.

Notes: solid, colored lines in each figure show the annual survival probabilities multiplied by the computed scale factor for bad (circled marked) and good (x marked) health for each education-marital status group. Dashed lines are the unscaled, objective survival probabilities estimated from HRS.

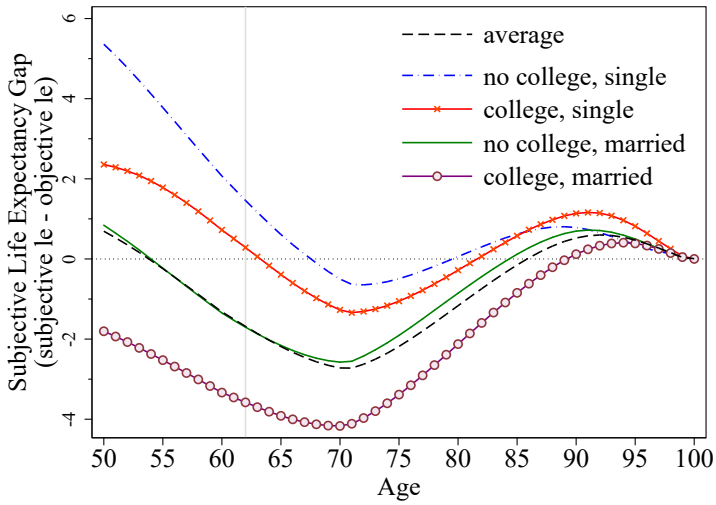
to deviations between how long a worker expects to live and how long the average individual of his demographic group does.

This gap is shown in Figure F.4. There is substantial variation across demographic groups and over the life cycle. At age 50, the average gap is about $+1$ year, meaning perceived life expectancy exceeds measured life expectancy by one year on average. However, across groups this ranges from nearly $+6$ years (for single men without a college degree) to about -2 years (for married college-educated men). By age 70, all demographic groups exhibit a negative gap, while by age 90 all have a positive gap.

Note that when we do not allow subjective survival to vary by health, the average life expectancy gap by education-marital status remains similar to what is shown in Figure F.4. However, in that case there is no within-group variation in the gap due to realized health outcomes. The life cycle and demographic heterogeneity are thus not the only sources of variation: because we estimate subjective survival curves separately by health status for each education-marital group, there is additional variation attributable to health. Moreover, for each group, the distribution of subjective life expectancy relative to objective extends across zero, indicating that even when the

group average is positive or negative, some individuals perceive themselves as living longer or shorter lives than predicted by life tables.

Figure F.4: Life expectancy gap by education, marital status
(subjective life expectancy - objective life expectancy)



Source: Health and Retirement Study, authors' calculations

Notes: figure shows the difference between subjective and objective life expectancy by education-marital group. Life expectancies are calculated using the annual survival probabilities from Figure F.3 and the health status shares implied by MEPS health transitions. The dashed black line indicates the average life expectancy gap. Group-specific weights used to compute this average are 17.4, 14.5, 24.3, and 43.7 for non-college single; college single; non-college married; and college married men, respectively. The vertical grey line marks age 62—the first age of Social Security eligibility.

F.3 Family structure

Family structure in the model varies by education. For men born between 1931 and 1935, the share married or cohabiting is roughly stable over the life cycle. Accordingly, we model marriage as a fixed state determined at age 25, based on initial conditions drawn from the data. The married share is set to approximately 56 percent for non-college-educated individuals and about 74 percent for college-educated men.

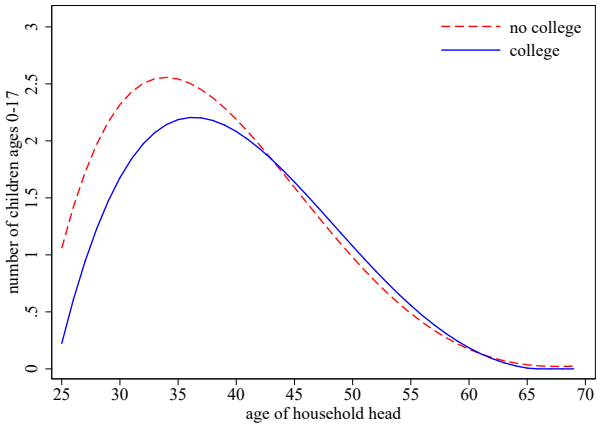
Figure F.5a displays the number of children in the household over the life cycle by education of the household head. This measure includes all children age 17 and under living in the household at any point in time. The number of children follows a hump-shaped profile, peaking around age 35 (slightly earlier for non-college-educated heads) and declining to zero by age 65. Non-college-educated men reach a peak of about 2.5 children, while college-educated men peak slightly above 2 children.

Married couples are eligible for Social Security spousal benefits, which depend on both the worker's age and the spouse's age. Consequently, the spousal age gap is an important determinant. Figure F.5b shows the distribution of this gap for married couples in the cohort. The average gap is approximately 4 years.

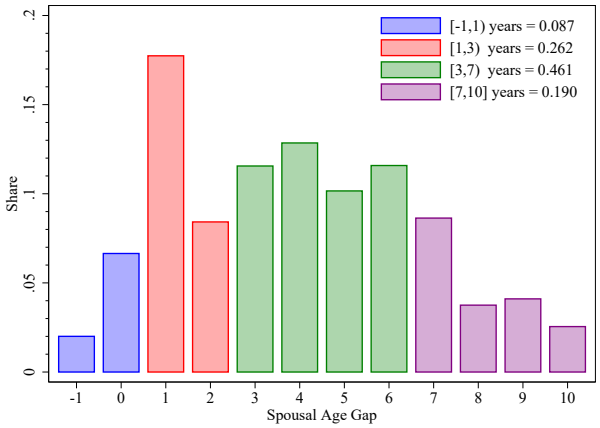
To focus on the central part of the distribution, we restrict attention to couples with an age gap between –1 (wife one year older) and +10 years, excluding the bottom 5 percent and top 10 percent. Even within this range, considerable variation remains. For computational tractability, we partition the distribution into four groups and assign each group the median gap within it. This approach captures both the prevalence of small gaps and the presence of larger ones. Estimated shares are as follows: 8.7 percent of couples have gaps of –1 to 0 years (median = 0), 26.2 percent have gaps of 1–2 years (median = 1), 46.1 percent have gaps of 3–6 years (median = 4), and 19 percent have gaps of 7–10 years (median = 8).

Figure F.5: Family statistics

(a) Number of children



(b) Spousal age gap



Source: Panel Study of Income Dynamics, authors' calculations

Notes: Panel (a) shows the number of children ages 0–17 in the household by education and age of the household head. Panel (b) displays the distribution of spousal age gaps. The color shading indicates the four groups used in the quantitative model.

F.4 Spousal income

Spousal income enters the budget constraint of married workers. Accurately estimating this income is important both for capturing total household resources and for predicting whether a spouse claims Social Security benefits on her own record or the household head's record. We estimate spousal income in two steps. First, we regress spousal labor force participation (or having positive labor income) on the state variables and income of the household head: a fourth. Second, we estimate a regression of spousal labor income—conditional on participation—on the same set of controls.

$$\mathbb{1}[y^s_{it} > 0] = X'_{it}\beta + \varepsilon_{it}$$
$$y^s_{it} = X'_{it}\beta + \varepsilon_{it}$$

where X'_{it} is a vector of controls, including a fourth-order polynomial in the age of the house-

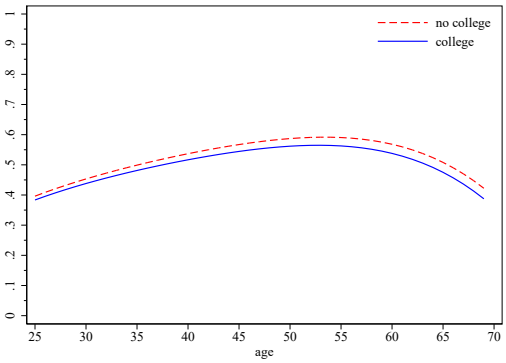
hold head; an indicator for whether the head attended college interacted with age; the head's labor income (in thousands) interacted with the fourth-order age polynomial; and the spousal age gap interacted with age.

$$\mathbb{1}[\hat{y}^s_{it} > 0] = X'_{it}\hat{\beta}$$
$$\hat{y}^s_{it} = X'_{it}\hat{\beta}$$

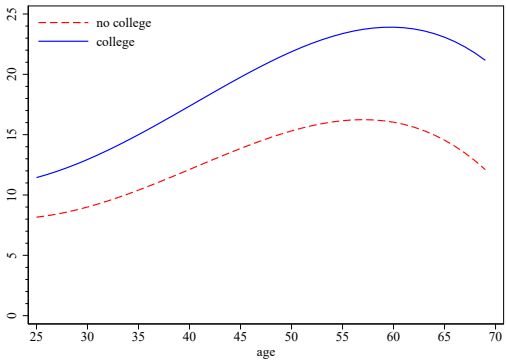
These regressions are estimated using the household earnings sample from the HRS linked to the SSA Summary Earnings file. The sample includes 75,545 observations for 1,270 households with male household heads born between 1926 and 1940. We use the estimated coefficients to impute spousal labor force participation (setting income to zero if not participating) and spousal income in the model. The estimated coefficients are reported in Table F.4. Spousal labor force participation and earnings profiles by education are shown in Figure F.6.

Figure F.6: Spousal labor force participation, earnings by age and education

(a) Labor force participation



(b) Labor earnings



Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.
Notes: average profiles for labor force participation and spousal labor income by age and education based upon the regression coefficients shown in Table F.4 using the average labor income (based on whether the spouse is working) of heads of household. Averages are computed by weighting age, education, and spousal age gap specific profiles by the distribution of age gaps: 0.087 have gaps (-1,0] years, 0.262 have gaps (0,2] years, 0.461 have gaps (2,6] years, 0.19 have gaps (6,10] years.

F.5 Wages

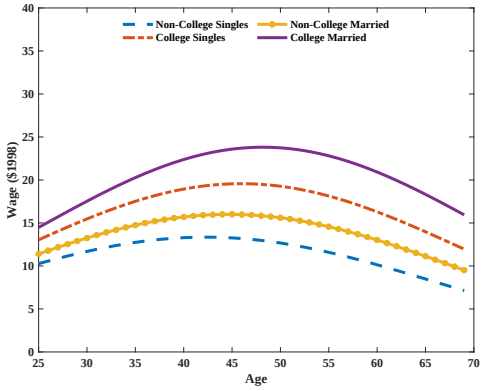
Data on wages is drawn from the PSID. Hourly wages are computed as annual earnings divided by annual hours worked. To address outliers and minimum wage distortions, we exclude observations with hourly wages below \$3.50 or above \$75 (in 1998 dollars), which trims approximately the bottom 1% and top 2% of the wage distribution. To ensure sufficient life-cycle coverage, the sample includes all workers aged 20 to 74 born between 1901 and 2000. Because wages are only observed for individuals who participate in the labor market, we estimate hourly wage profiles using a Heckman selection model.

In the first step, a probit regression estimates the probability of labor market participation and generates the inverse Mills ratio. This ratio is then included as a regressor in the second-stage wage regression. Labor market participation is modeled as a function of a quadratic in age, interacted with a college education indicator, as well as dummies for being married, having children, and being in good health—all interacted with age. We also control for cohort fixed effects. Table F.5 confirms the significance of the inverse Mills ratio, indicating that selection bias is present in the simple OLS estimates of wages.

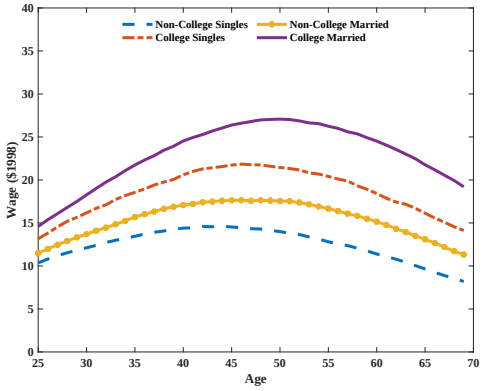
In the second stage, we estimate wage profiles by regressing the natural logarithm of hourly wages on a second-order polynomial in age, its interaction with a college education indicator, and an interaction between age and a marital status indicator. The regression also includes cohort fixed effects. Table F.5 reports the estimated coefficients, and panel (a) of Figure F.7 displays the resulting wage profiles by age, education, and marital status.

Figure F.7: Wage profiles by education and marital status

(a) Unadjusted for unemployment effects



(b) Adjusted for unemployment effects



These estimated wage profiles exhibit several interesting patterns. First, wages generally increase over the life cycle, with the rate of growth slowing at older ages. Consistent with (Rutledge et al., 2017), we find that these profiles remain fairly flat as workers approach retirement rather than declining during these years. Second, in addition to a clear skill premium associated with higher education levels, our results (aligned with Borella et al. (2023)) also reveal a marriage premium: married men earn higher wages than their non-married counterparts across both education groups.

These estimated profiles are then used to compute predicted wages for individuals in the sample. The residuals—defined as the difference between observed and predicted hourly wages—are used to estimate the stochastic component of the wage process. Specifically, the parameters of the AR(1) income process are estimated using Method of Simulated Moments, with the identity matrix as the weighting matrix. This procedure yields an estimate of $\rho = 0.968$ for persistence and $\sigma^2 = 0.022$ for the variance of the wage innovation.

Unemployment shock and wage scarring

Two parameters govern employment shocks in the model: the probability of receiving a shock, λ, and the wage penalty associated with it, ξ.

We set λ equal to the annual involuntary separation rate. Since our model explicitly captures wage scarring from unemployment, we focus solely on separations due to involuntary job loss, excluding voluntary exits and job-to-job transitions.

The separation rate is computed in two steps. First, using CPS data from 1995–2005 for men aged 25–69, we estimate the monthly separation rate among employed workers and annualize it, yielding a total separation rate of 19.9%. This rate ranges from 17.5% to 24.5% across education-marital groups and from 16.4% to 28.0% over the life cycle. Second, we adjust this figure to exclude voluntary separations. Based on Simmons (2023), who use SIPP data to classify separations, only 28.6% of separations are involuntary.[50] Applying this share yields an annual involuntary separation rate of 5.6%.

To construct the wage penalty from an unemployment shock, $\xi(\lambda)$, we follow the literature estimating the effect of job displacement on re-employment earnings. Specifically, we estimate the following regression:

$$y_{it} = x_{it}'\beta + \sum_{k\geq -2}^{20} \delta^k D_{it}^k + \alpha_i + \gamma_t + \varepsilon_{it} \tag{20}$$

Here, y_{it} is the log wage of individual i at time t; x_{it} includes controls such as education and a quadratic in experience; D_{it}^k is a set of dummies indicating the k-th year relative to a displacement event; and α_i and γ_t are individual and time fixed effects.

The estimated coefficients δ^k are plotted in Figure F.8. Consistent with previous work, we find that displaced workers earn roughly 14–20 percent less than non-displaced workers, with these losses persisting for 5 to 10 years post-displacement. Based on these estimates, we set the wage penalty from an unemployment shock at 14 percent, i.e., $\xi(\lambda) = 0.86$. Additional analysis by education level reveals similar displacement effects across educational groups.

F.6 Pensions

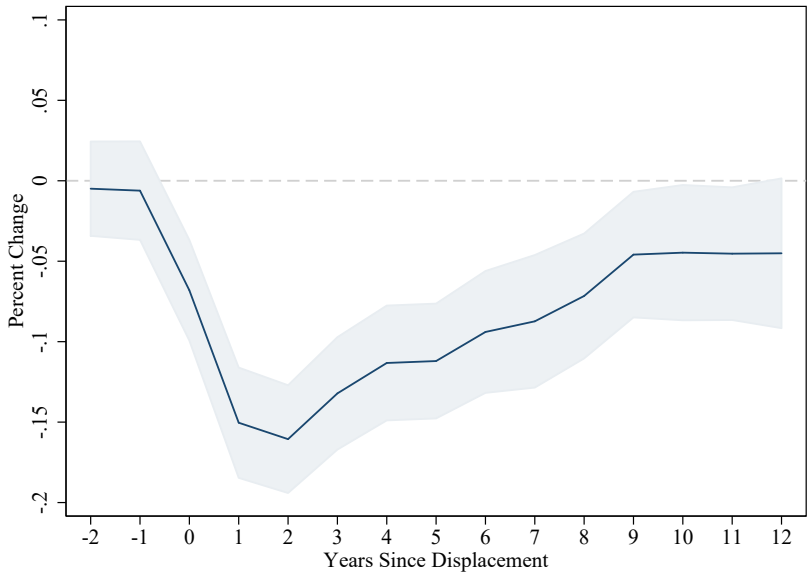
Following French (2005), we model private pensions as a regressive function of individuals' Primary Insurance Amount, pia_t, specified as:

$$pb(pia_t) = \gamma_0(\gamma_1 + \gamma_2 pia_t + \gamma_3 \max\{0, pia_t - 5000\})$$

The parameters γ_1, γ_2, and γ_3 are taken from Gustman and Steinmeier (1999), consistent with both our model and French (2005). The scale parameter γ_0 is calibrated separately for each education and marital group to match the average pension benefits received by these groups at older ages in the HRS data. Specifically, we set γ_0 to 0 for non-college singles, 0.20 for college singles, 0.18 for non-college married, and 0.25 for college married individuals. A scale of zero implies that non-college single individuals receive no pension benefits in the model, reflecting both their low pension coverage and low benefit amounts conditional on coverage. Age-specific pension accrual rates are taken from Gustman et al. (1998). The adjustment ϵ_t^p accounts for the difference in accrual patterns between pension and Social Security wealth at various ages, which tends to be negative at

[50] Simmons (2023) uses the Survey of Income and Program Participation to classify separations as involuntary (due to lay-off, discharge/firing, employer bankruptcy or sale, or temporary employment ending) or voluntary (due to employer-to-employer transitions, dissatisfaction with the job, personal circumstances, or other reasons).

Figure F.8: Wage scarring due to unemployment



Source: Panel Study of Income Dynamics, authors' calculations

Notes: figure is the graphical presentation of the estimated β coefficients from Equation 20. Solid line shows the impact of an unemployment spell on hourly wages (relative to the wages of an employed individual) for 2 years prior to and 12 year after such a spell. 90 percent confidence intervals shown.

younger ages and positive at older ages.

F.7 Taxes

Taxes are modeled using a smooth function as in Heathcote et al. (2010). Specifically, total taxes, $T(y_t)$, paid on income y_t is given by the following equation:

$$T\left(y_t\right) = y_t - \lambda_t y_t^{1-\xi_t}.$$

where λ_t determines the average level of taxes in the economy in year t and ξ_t captures the level of progressivity in year t. This function sets a relationship between pre-tax and after-tax income that we exploit in our estimation.

$$\underbrace{\overbrace{y_t}^{\text{pre-tax income}} \quad -T\left(y_t\right)}_{\text{after-tax income}} = \lambda_t y_t^{1-\xi_t}.$$

$$\ln\left(y_t - T\left(y_t\right)\right) = \ln\left(\lambda_t\right) + \left(1-\xi_t\right)\ln\left(y_t\right)$$

To estimate λ_t and ξ_t we regress the logarithm of after-tax household income on a constant and the logarithm of pre-tax household income.

PSID includes information on taxes paid up until 1991 and covers tax years up through 1990. Total taxable income of the family is measured as the sum of labor and Social Security income of the household head and the spouse and other family members (if present). Federal tax liability is constructed based upon the taxable income of the family as well as exemptions and the tax table used. We use this information to estimate parameters of the tax function which vary based upon tax year, marital status (single or married), and, if married, number of children (none, 1 child, 2 children, more than two children). Our sample consists of 40,202 taxpayers born between 1921 and 1950 who are between the ages of 20 and 75 who also had non-missing tax information. We drop households whose taxable income is below \$500 (1998 dollars) or overs \$300,000 as well as those who report paying taxes that would imply an average tax rate above the maximum marginal tax rate for that year.[51] In order to maximize the sample available to estimate tax year specific functions which incorporate changes in taxes over time, we follow the procedure in Borella et al. (2023) and estimate these time specific parameters in rolling 5-year bins.

We map these tax year and family structure specific parameters into the inputs for our structural model in two steps:

First, because we follow a single cohort in our analysis, time and age are equivalent in our setup. Therefore, we compute age- and family-structure-specific parameters based on the time-specific parameters. Specifically, parameters for tax year 1967 are used for ages 25 to 32. For ages 33 to 58, parameters are computed as the average of the parameters for the tax years during which our cohort would have been at each respective age.[52] Parameters for tax year 1990 are used for all ages 59 and above.

Second, we map these age- and family-structure-specific parameters into age, education and marital status specific parameters using the observed variation in the number of children over the life cycle. Single individuals (regardless of education level) are assigned the age-specific parameters for singles. Married individuals are assigned the profiles corresponding to the average number of children that households of their education level have at each age (see Figure F.5a). This approach allows the tax parameters to vary by education, reflecting differences in the number of dependents across groups.

F.8 Subjective social security program beliefs

We model program beliefs by including uncertainty over the adjustment factor for early claiming, $\widetilde{\gamma}_1^{ss}$. Perceived Social Security benefits (given by Equation 21) are a function of the primary insurance amount for an individual of a given SS wealth, $pia\left(a_t^{ss}\right)$, the uncertain total penalty (the accumulated annual penalties), $\widetilde{\Gamma}(t^{ss})$, and the retirement earnings test, Υ.

$$\widetilde{ssb}_t = pia\left(a_t^{ss}\right) * \widetilde{\Gamma}(t^{ss}) - \Upsilon_t \tag{21}$$

To estimate the belief distribution this uncertain penalty is drawn from, we use a question

[51] These restrictions drop only 639 observations.

[52] For example, the parameters for age 40 are the average of the parameters for tax years 1971–1975, as these are the years when the 1931–1935 cohort would be age 40.

which asks respondents about the increase in benefits from delaying claims from age 63 to 64. The question asks:

Imagine an individual, Mr. Spencer Wills, who is retired from work and turning 63 today. Spencer expects his monthly Social Security retirement benefit to be about $2,000 if he claims now. How much should he expect his monthly benefit to be if he decides to wait another year and claim when he is 64?

(a) it would stay the same

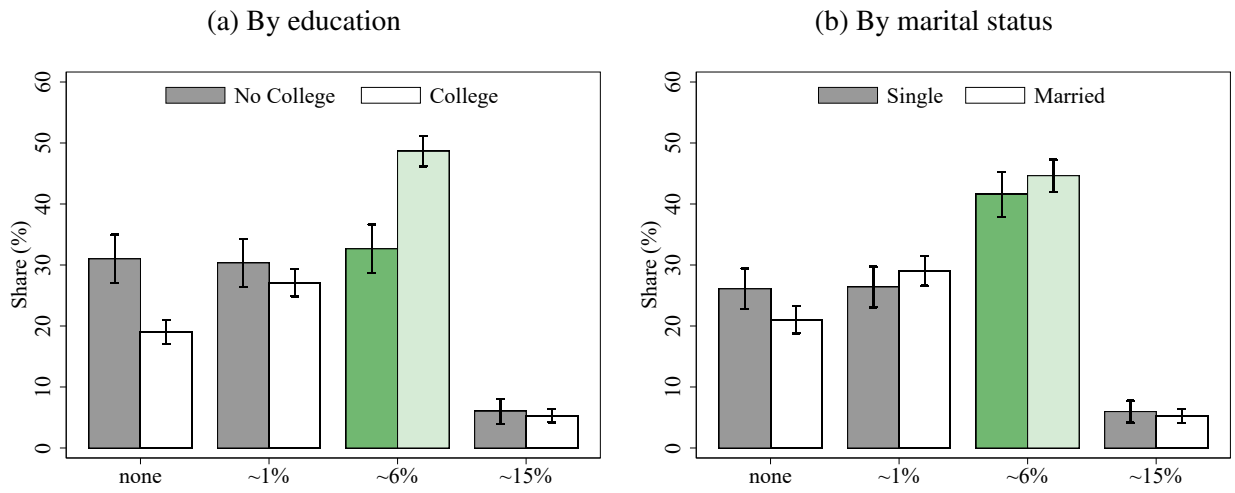
(b) approximately $2,120 (6% higher)

(c) approximately $2,020 (1% higher)

(d) approximately $2,300 (15% higher)

The answer to this question allows us to determine not only whether workers have correct information about the penalty, but also what they believe if they are misinformed. Respondents who answer (b) are classified as correctly informed, while those who answer (a), (c), or (d) are considered misinformed and assigned to the corresponding categories. Figure F.9 shows how these beliefs vary by education and marital status, and Table 4 reports the corresponding shares. Because we do not have a comparable survey question regarding the value of the delayed retirement credit, we assume that individuals are informed about this parameter.

Figure F.9: Program knowledge by education, martial status



Source: Understanding America Study (UAS), authors' calculations.

Notes: Figures show the how individuals, by education (Panel (a)) and marital status (Panel (b)), respond to a question about how much benefits increase by delaying claims from age 63 to age 64. Correct answer shown in green. 95% confidence interval shown.

We also use a question about how knowledgeable the respondent is about this penalty to determine confidence in the response. The question used reads:

Below is a list of aspects of the Social Security system. For each one do you feel you are very, somewhat, not too, or not at all knowledgeable about that aspect?

How Social Security retirement benefits change based on how old you are when you claim.

Responses to this categorical variable are converted to a numerical scale in the following way: If an agent reports the are knowledgeable we set the confidence share to 1; if somewhat knowledgeable the CS is set to 0.75; if not too knowledgeable the CS is set to 0.25; if not at all knowledgeable the CS is equal to 0. The the average of this confidence share (ACS) by education-marital status and belief type is also shown in Table 4.

Exclusion of learning

A natural thought is that individuals may learn about the rules of the Social Security program as they approach retirement. To test this hypothesis, we compute an error measure of the deviations between the perceived penalty and the true penalty. Figure F.10 shows how this measure varies over the life-cycle. If there was high amounts of learning occurring, we would expect the mean error to approach zero as retirement nears. We observe that this measure is roughly flat over the life-cycle; we do not observe clear signs of learning.

We acknowledge that this result may conflate age and cohort effects, which we are unable to disentangle given the limited panel dimension of our data. Thus, we cannot fully rule out the possibility that some learning occurs as individuals approach retirement. However, imposing a learning process in the model would require estimating an extremely noisy and uninformative signal, which would have minimal impact on beliefs over time and would not improve the model's fit to the data.

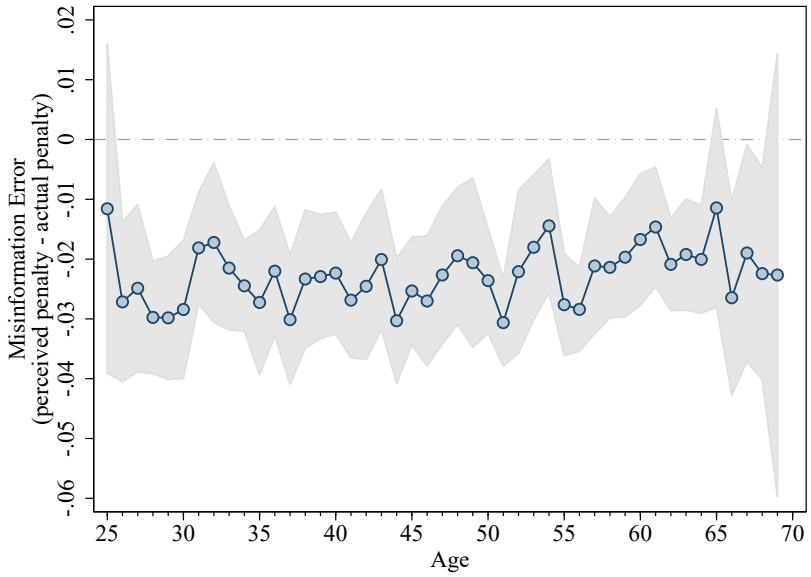
Results

The results of our estimation procedure described in Section 4.2 give us uncertainty weights that individuals, who vary by education, marital status, and belief type, use to weigh the potential annual penalties. Using these, weights, we can can compute the uncertainty-weighted penalty defined by:

$$E\left[\widetilde{\gamma}_1^{ss}|k\right] = \varepsilon_{k,k}^{ss}\widetilde{\gamma}_1^{k} + \sum_{j\neq k}\varepsilon_{k,j}^{ss}\widetilde{\gamma}_1^{j}, \text{ where } \sum_{j}\varepsilon_{k,j}^{ss} = 1\ \forall k \tag{22}$$

Deviations of this uncertainty weighted penalty from the true statutory penalty aggregated to education groups, marital statuses, and the average are included in Figure F.11. We observe that on average, this deviation is negative (the uncertainty weighted penalty is larger than the true penalty. This means that the beliefs push the cost of early claiming above the true level and incentivizes earlier claiming. Additionally, while both education levels similarly overestimate this penalty (meaning the uncertainty incentivizes earlier claiming), the pattern deviates by marital status. While married individuals strongly overestimate the penalty, singles strongly underestimate the penalty. This indicates that subjective program knowledge is pushing claims earlier among the married and later among the singles.

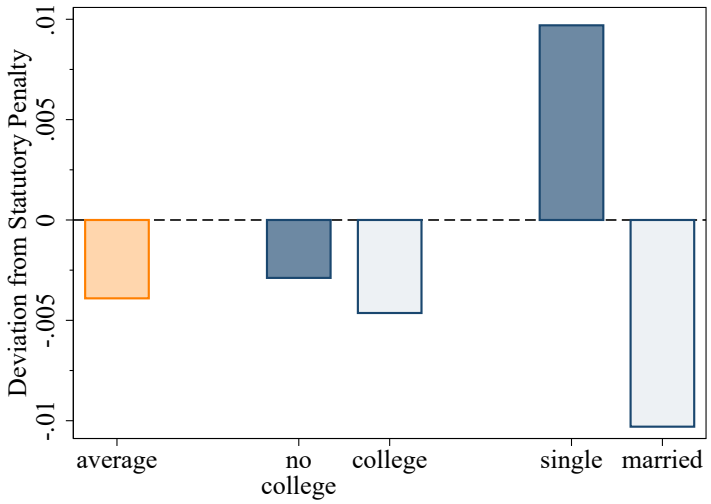
Figure F.10: Average penalty error by age



Source: Understanding America Study (UAS), authors' calculations.

Notes: Figure does the average deviation, by age, between the perceived increase in benefits for delaying claims from age 63 to age 64 and the true increase for delaying claims this single year. A negative value indicates that respondents believe the penalty for claiming one year earlier is smaller than in the policy rules; a positive value indicates that respondents believe this penalty is larger than the policy prescribes. 95% confidence interval shown.

Figure F.11: Deviations from Statutory Penalty



Notes: figure shows the deviation between the average perceived penalty and the statutory penalty of 6.67% for overall, by education group, and by marital status

G Second-step estimation

We estimate the model by matching moments that vary by age, education, and marital status. We set the population shares for these groups based on the shares observed at age 25 in the PSID. Table G.1 shows the share of the population within each fixed type.

G.1 Target profiles

The target life-cycle profiles of participation, hours, and assets are estimated from the Panel Study of Income Dynamics (PSID) covering the years 1968–2016 (annually from 1968–1995 and biennially from 1996–2016), using the method described in French (2005).

In each year t, PSID respondents report their labor market variables for year $t-1$. Participation is measured as a binary variable equal to 1 if the respondent worked at least 300 hours (25 hours per month) in the previous year. Hours are measured as a continuous, self-reported variable capturing total annual hours worked. When constructing hours profiles, we restrict the sample to respondents reporting more than 1,000 hours worked annually.

We use wealth data from the 1984, 1989, 1994, 1999, 2001, 2003, 2005, and 2007 PSID wealth surveys, as well as the 2010–2016 PSID family files. Household assets are measured as the net value of farm and business equity, vehicles, checking and savings accounts, real estate, stocks, mutual funds, IRAs, Keoghs, and investment trusts, after subtracting debts. We drop observations with household wealth above $1 million (in 1998 dollars).[53]

To measure mean life cycle profiles, we estimate the following fixed effect regression:

$$Y_{it} = f_i + \sum_{c\in\{0,1\}} \sum_{m\in\{0,1\}} \left(\sum_{k=1}^{J} \beta_c^m \mathbb{I}\left[age_{it} = k\right] \times \mathbb{I}\left[married_i = m\right] \times \mathbb{I}\left[college_i = c\right] \right) + \sum_{f=1}^{F} \beta_f \mathbb{I}\left[family\ \ size_{it} = f\right] + \beta_u u_t + \varepsilon_{it}$$

where Y_{it} represents an observation of the outcome variable (labor force participation, hours worked, or assets) for individual i in time period t; f_i is an individual fixed effect;

$$\mathbb{I}\big[\text{age} = k\big] \times \mathbb{I}\big[\text{married}_i = m\big] \times \mathbb{I}\big[\text{college}_i = c\big]$$

are education- and marital-status-specific age-group dummies[54];

$$\mathbb{I}\big[\text{family size}_{it} = f\big]$$

are family size dummies; and u_t is the unemployment rate in period t. The estimated parameters

[53] This restriction removes only 1,302 observations and leaves a sample of 93,232 observations. We also experimented with dropping the 1,723 observations with negative wealth, since the model does not allow borrowing. The estimated wealth profiles are not overly sensitive to including these cases. Among households with negative wealth, average debt is slightly over $29,000, and median debt is roughly $9,000.

[54] We could have alternatively used age dummies for each year. We choose five-year age-group dummies to reduce noise in the estimated profiles.

are

$$\{\beta_{nc}^{nm}, \beta_{c}^{nm}, \beta_{nc}^{m}, \beta_{c}^{m}, \beta_{f}, \beta_{u}\}.$$

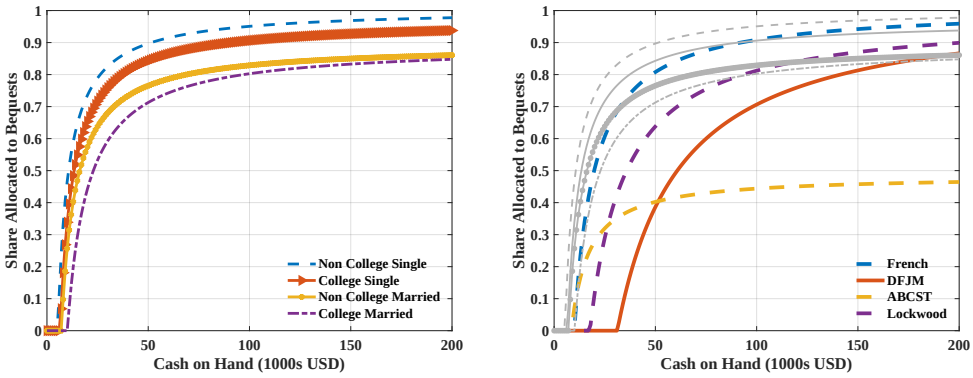
In constructing these profiles, we make three adjustments to control for year, family size, and cohort effects: (1) we set the unemployment rate to 6.5%; (2) we fix family size at 1 for single individuals and 4 for married individuals; and (3) for each education and marital status group, we use the individual-specific effect of those born in 1930 and aged 45–49.

For the wealth moments, we report median in addition to the mean, computed by estimating a quantile regression on the set of education- and marital-status-specific age dummies.

G.2 Additional estimation results and model fit

Here we compile results related to estimation and model fit that are not presented in the main text. These include: (i) Table G.2, which reports the estimated weight parameters for Social Security program uncertainty; (ii) Figures G.1 and G.2, showing the estimated bequest motives and the time cost of working, respectively; (iii) Figures G.3–G.5, illustrating the model fit for life-cycle participation, hours worked, and wealth; (iv) Figure G.6, showing the fit for cumulative benefit claiming; and (v) Figure G.7, showing the model fit for early claiming rates by age 62 labor force participation and self-reported health status.

Figure G.1: Estimated bequest motives

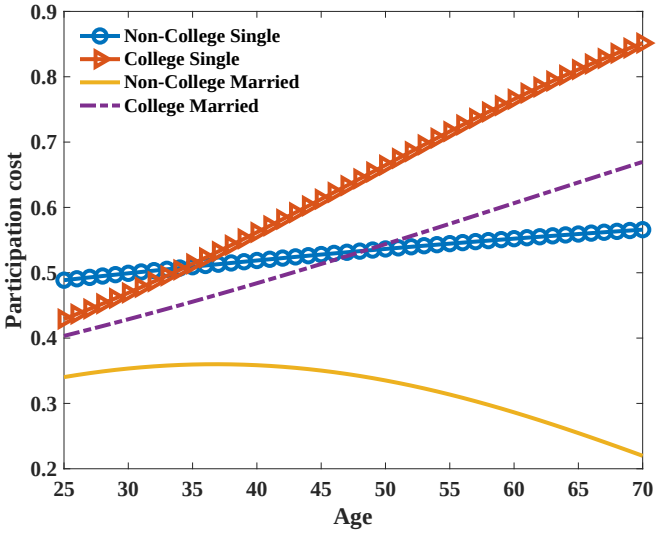


(a) (b)

Notes: Panel (a) presents expenditure share allocated to bequests for different education and marital groups facing certain death in the next period (refer to De Nardi et al. (2021) for calculation details). Panel (b) provides comparison with the literature. Grayed lines refer to groups in panel (a) and colored lines: (French) French (2005); (DFJM) De Nardi et al. (2021) for singles; (ABCST) Ameriks et al. (2020); and (Lockwood) Lockwood (2012).

With regard to characteristics that evolve over the life cycle, the model captures the correlations between health status, employment, and claiming behavior (see Figure G.7). It closely matches the early claiming rates of individuals in good health at age 62, though it slightly underpredicts the early claiming rates of those in poor health. The model also reproduces the gradient in early claiming rates between individuals who are working and those who are not working at age 62.

Figure G.2: Time cost of working



Notes: Time costs of labor market participation estimated from the structural model, expressed as a fraction of the time endowment, is reported for all college and marital groups.

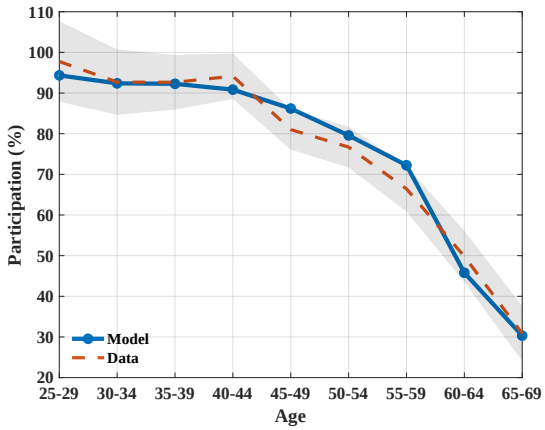
H Counterfactual experiments

In order to arrive at our no mechanisms model (Baseline A), we start from our full specification model detailed in Section 3 and switch off all potential drivers of Social Security benefit claiming. Switching off the unemployment shocks implies that unemployed individuals experience no wage-scarring effects if they choose to immediately re-enter the labor market. This results in somewhat higher life-cycle wage profiles as shown in Figure H.1.

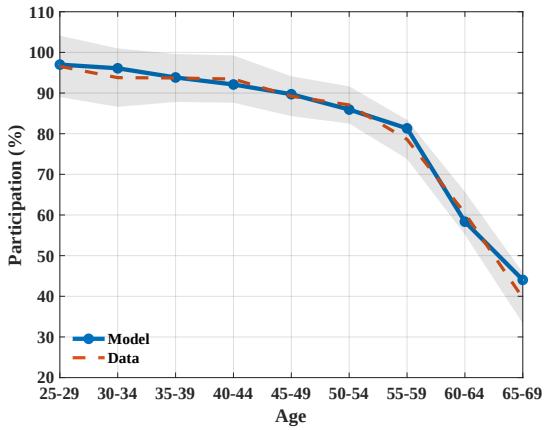
We switch off the effect of bad health on the time endowment for all individuals in the model for the health channel. Specifically, in our model solution, individuals with the bottom health state observe time endowment of those in good health (within their age-education group). As changes in life expectancy have a distinct effect on early claiming, we do not eliminate the impact of health on mortality. In other words, we still allow for the bottom health state to affect survival probabilities. We turn off bequest motives by setting the end-of-life flow utility from positive liquid assets to zero for all individuals. Note that switching off bequest motive by itself does not switch off survivors benefits for married individuals. For them, we still allow end-of-life utility from bequeathing their Social Security benefits to their spouses as survivors benefits unless marital benefits are switched off as well. For switching off subjective beliefs about early application penalty, we allow all our simulated individuals to be perfectly informed about penalty and remove any wrong priors or uncertainty around priors. For switching off subjective beliefs about survival, we give individuals their objective survival probabilities, conditional on health, instead of the ones based on their subjective evaluations. Finally, we switch off marital benefits by switching off both spousal and survivors benefits. For the latter, we do not allow any warm glow over leftover benefits after the head has deceased. Note that spouses are still allowed to claim on their own earnings record in the absence of these marital benefits.

Figure G.3: Benchmark: Participation by education and marriage
Men born in 1931-1935

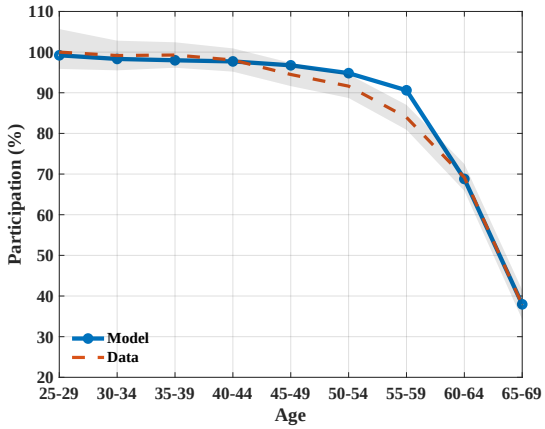
(a) Non-college, singles



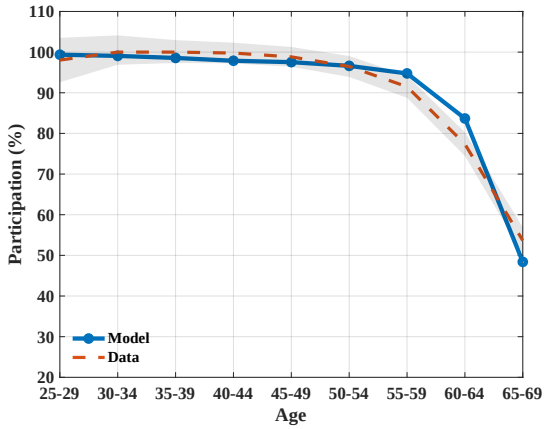
(b) College, singles



(c) Non-college, married



(d) College, married



I Robustness

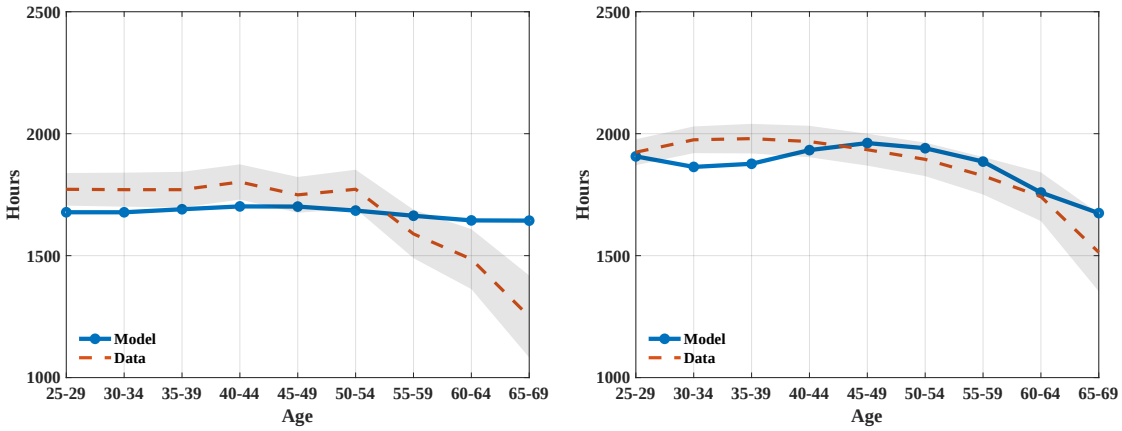
We assess the robustness of our key results along two main dimensions: (1) assumptions about preference heterogeneity, and (2) the inclusion of additional model features. First, we constrain preference heterogeneity by fixing bequest motives across groups—an important factor given their influence on claiming decisions. Second, we enrich the model by incorporating uncertain medical expenditures, which offers a complementary check on the role of precautionary savings versus altruistic motives in late-life wealth accumulation. In additional experiments, we also relax the assumption of joint benefit claiming within couples and allow for misinformation about the Social Security Earnings Test.

We focus on the robustness of two core findings: (1) the level of claiming ages—overall and by education and marital status—and (2) the size of the claiming age gaps across these groups. Table I.1 presents the results. Columns 1, 2, 3, 5, and 6 report average claiming ages for the full population, non-college group, college group, singles, and married households, respectively. Columns 4 and 7 report the education- and marital-based claiming age gaps across specifications. We find

Figure G.4: Benchmark: hours by education and marriage
Men born in 1931-1935

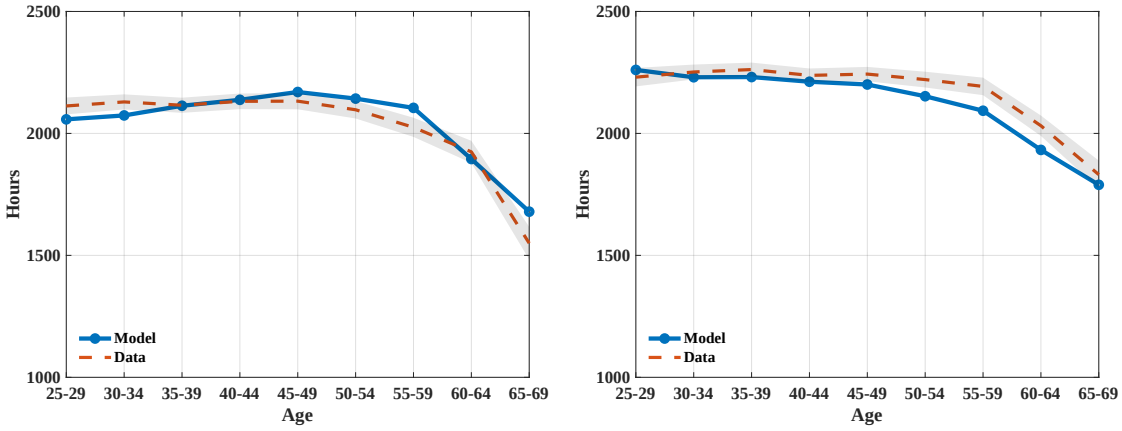
(a) Non-college, singles

(b) College, singles



(c) Non-college, married

(d) College, married



that, after full re-calibration, both the levels and the gaps in claiming ages remain remarkably stable across all alternative specifications. This suggests that our main conclusions—the overall claiming behavior and the limited heterogeneity across socioeconomic groups—are not sensitive to these modeling choices.

I.1 Fixed bequest motives

We fix bequest motives $\kappa_{beq} = 89,570, \theta_{beq} = 1.29$ for all groups and re-estimate the model to match the same wealth, labor supply and Social Security cumulative claim-age moments. The estimated second-step preference parameters for this baseline are provided in Table I.2 below:

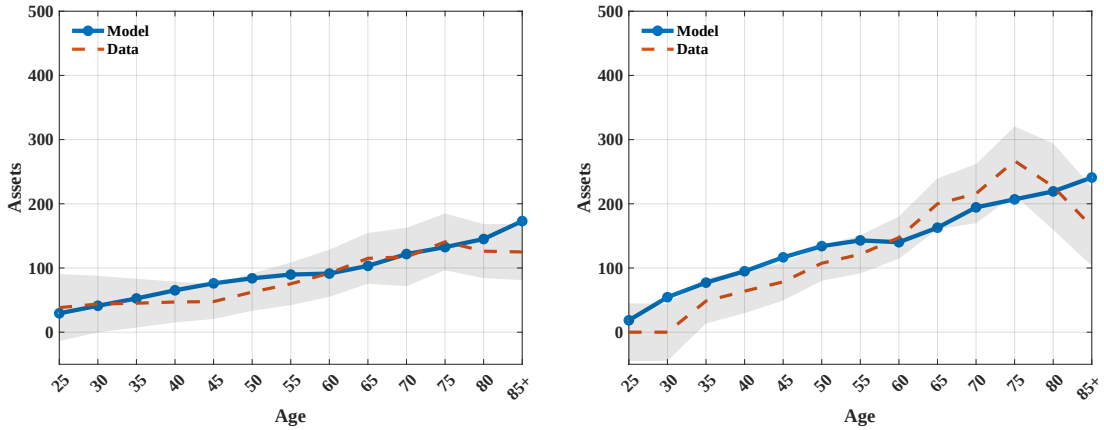
I.2 Medical expenditures

We estimate annual out-of-pocket medical expenditures from the Health and Retirement Study data using all years from 1992-2018, following the same approach as detailed in Borella et al.

Figure G.5: Benchmark: wealth by education and marriage

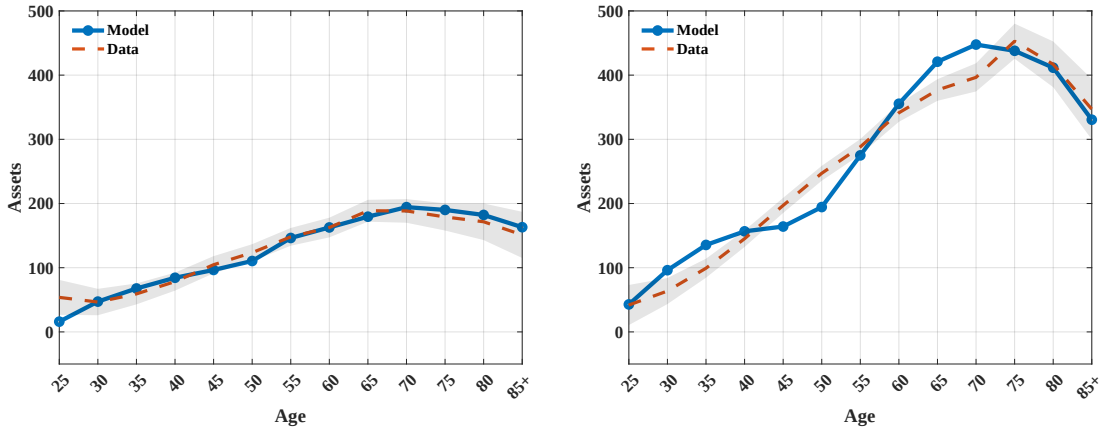
(a) Non-college, singles

(b) College, singles



(c) Non-college, married

(d) College, married



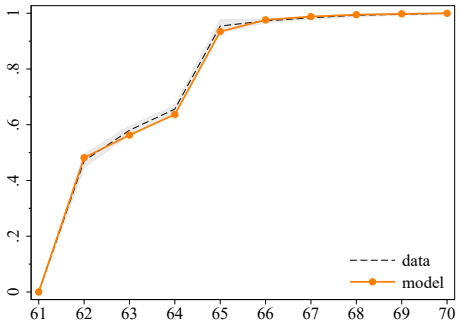
(2023). We use their approach to medical expenditures for each education and marital groups. We allow these expenditures to vary health. However, using the approach in Borella et al. (2023), we allow expenditures to vary by good and poor health only. In other words, we combine the first three groups into one. There are two key details in the estimation of these profiles from the data. First, it is important to include medical expenses in the last year of life. This comes from the HRS exit interviews. Second, expenses covered by public or private insurance are not included in the measure, because they are not directly incurred by the individual. Figure I.1 reports the estimated expenditures for male household heads, born between 1931-1935, by age and health, for different education and marital status. The estimated second-step parameters for this baseline are provided in table I.3 below:

I.3 Spousal benefit claiming

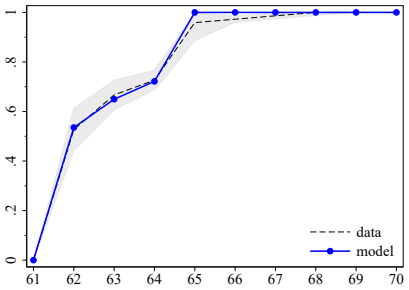
Next, we allow husbands and wives to make independent claiming decisions. For over half of the households where spouses claim auxiliary benefits instead of their own, the decision remains

Figure G.6: Benchmark: SS benefit claiming
by education and marital status

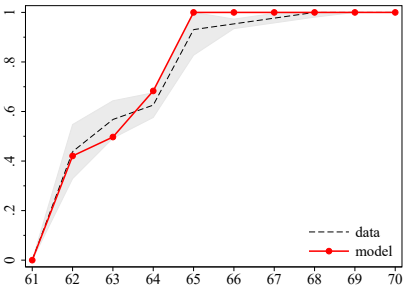
(a) Average



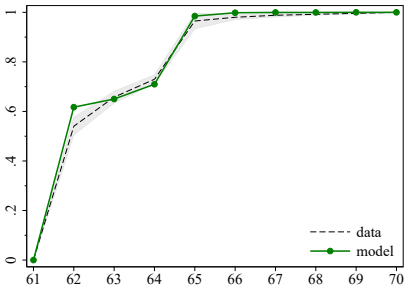
(b) Non-college, singles



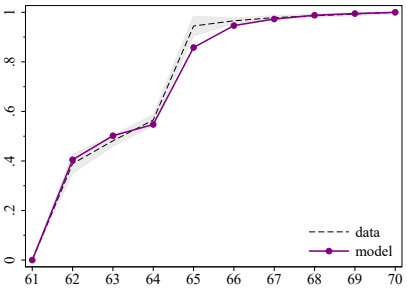
(c) College, singles



(d) Non-college, married



(e) College, married



Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.

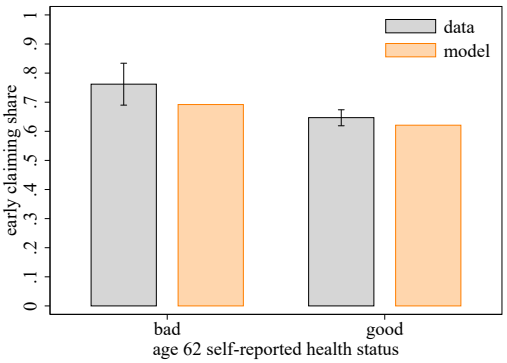
tied to the husband's claiming status, as auxiliary benefits are only available once the husband has claimed. The key change is that the model now permits spouses to delay claiming even after the husband has filed. While this extension adds significant computational burden to our already rich framework, it does not materially affect our main results. Note that even though we fully recalibrate the model, the preference parameter estimates do not change in any significant way.

I.4 Retirement earnings test misinformation

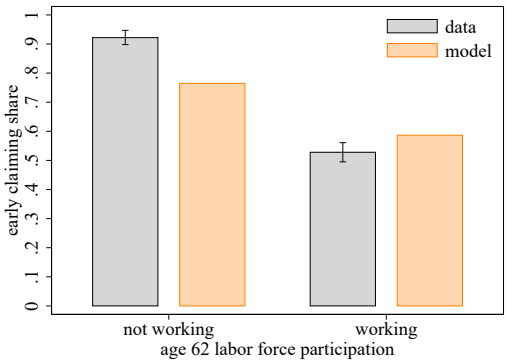
We allow individuals with fixed belief types other than the accurate penalty belief ($k = 6.7\%$) to also hold misinformation about the Earnings Test—specifically, they perceive no Earnings Test

Figure G.7: Early claiming by age 62 labor force participation, self-reported health

(a) Self-reported health status



(b) Labor force participation



Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.
Notes: grey bars the mean estimate and the 95% confidence interval of the share of men born 1926-1939 of a given labor force participation status (Panel (a)) or self-reported health (Panel (b)) at age 62 who claimed Social Security retirement benefits prior to the normal retirement age (65). Orange bars represent these shares from our life-cycle model.

when making their Social Security claiming decisions. Note that even though we fully recalibrate the model, the preference parameter estimates do not change in any significant way under this specification as well.

J Alternative Modeling of Subjective Program Beliefs

J.1 Reference Dependence

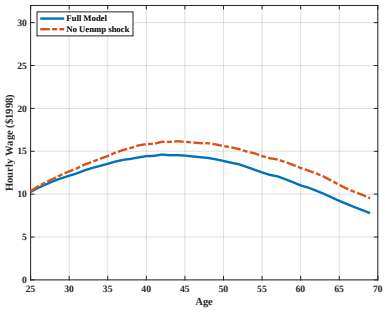
A growing empirical literature provides evidence that salient statutory ages can act as reference points for claiming and retirement decisions. In the U.S. context, Behaghel and Blau (2012) document that the spike in claiming at age 65 moved with the legislated increase in the FRA and interpret this as evidence of reference dependence with loss aversion. Brown et al. (2016) show experimentally that expected claiming ages are sensitive to the framing of benefit information. These studies support the premise that reference points may matter, but they do not by themselves deliver a unified life-cycle framework that can simultaneously match savings, labor supply, and the full claiming-age distribution that we target.

A practical challenge for a parsimonious structural implementation is that U.S. claiming behavior features *two pronounced spikes*—at the earliest eligibility age (62) and at the NRA (65)—so a formulation with a *single* reference age is generally not sufficient on its own. A single reference at 65 can generate bunching at 65, but it does not naturally generate an additional spike at 62 without further structure (e.g., multiple reference points, heterogeneity in reference ages, or other mechanisms such as liquidity constraints or misperceptions of the claiming adjustment).

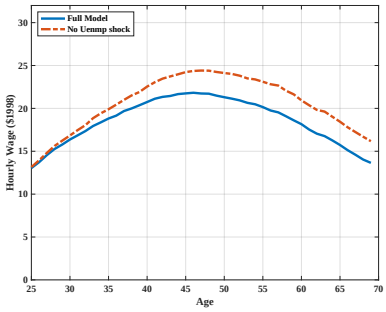
To make this point transparently within our own quantitative framework, we implement two one-parameter reference-dependent variants of the benchmark model, switching off subjective program-rule beliefs entirely (agents are fully informed about the statutory adjustment factor γ_1^{ss}) while keeping all other features of the model intact. We then re-estimate the model, replacing the

Figure H.1: Simulated wage profiles:
Full specification vs. no employment shocks

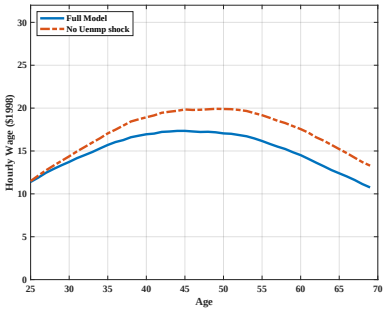
(a) Non-college, singles



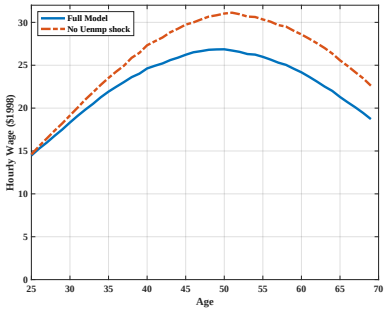
(b) College, singles



(c) Non-college, married



(d) College, married



program-beliefs channel with a single reference-dependent parameter governing the strength of reference-point utility in the claiming decision:

1. *Reference age 65 (NRA).* We use a standard "simple loss aversion" specification (losses-only around the reference outcome), following the functional form emphasized in Reck and Seibold (2023): claiming below the reference generates disutility proportional to the shortfall relative to claiming at 65.

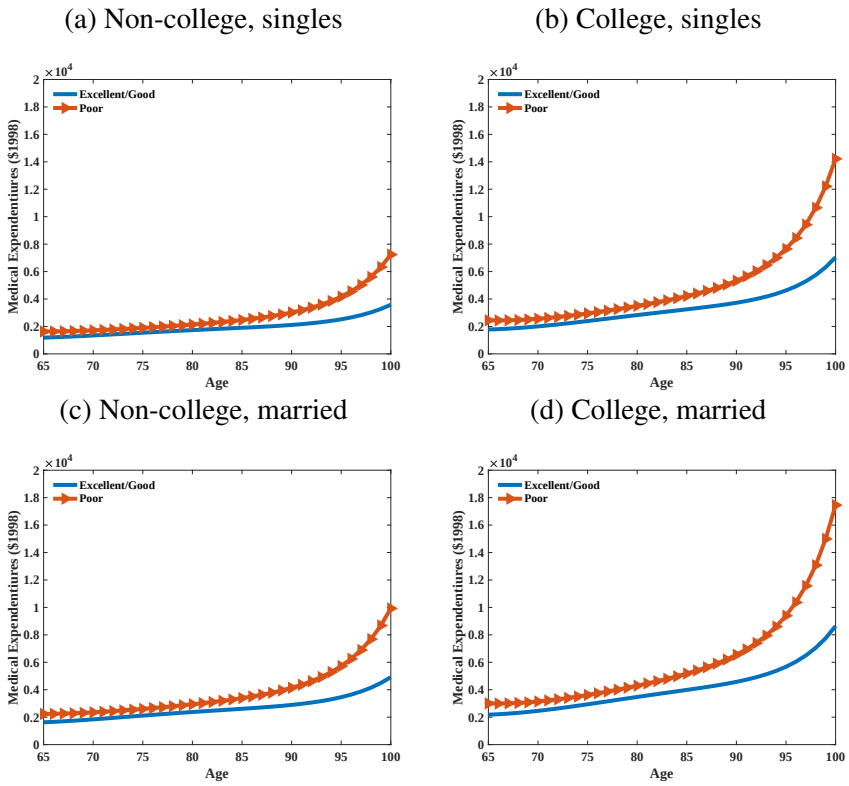
2. *Reference age 62 (EEA).* To give the reference-point channel maximal ability to generate a spike at 62, we also consider a (non-standard) "distance" specification in which any deviation from claiming at 62 is costly.[55]

Figure J.1 reports the resulting claiming fits for college and non-college singles.[56] The results are clear: with a *single* reference age, the model fits the peak at the reference age but necessarily misses the other peak. When 65 is the reference, the model over-predicts bunching at 65 and under-predicts the mass at 62; when 62 is the reference, the model matches early claiming but fails to generate the pronounced spike at 65. Thus, matching the two-peaked claiming distribution within a reference-point framework would require additional assumptions (e.g., two reference ages,

[55]This formulation is intentionally generous to the hypothesis that 62 is a focal anchor; it is closer to an inertia/adjustment-cost model than to standard loss aversion.

[56]Model fit results for married group is available upon request.

Figure I.1: Annual medical expenditures by health, education and marriage
Men born in 1931-1935



heterogeneity in reference ages, or other interacting mechanisms), at which point the approach becomes substantially less parsimonious and would require additional data to discipline the added structure. By contrast, our baseline approach is explicitly data-driven: we measure heterogeneity and uncertainty in program-rule expectations directly and show that this channel is quantitatively important for matching both the claiming distribution and the cross-group claiming-compression puzzle.

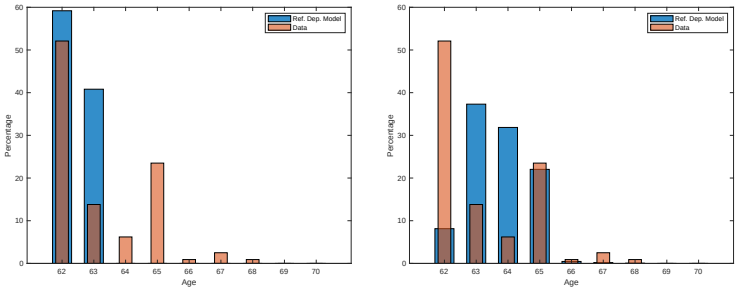
J.2 HRS Belief Distribution

In the benchmark model, we discipline the support of fixed program-belief types and their group-specific distribution using the Understanding America Study (UAS): we infer a discrete menu of perceived annual early-claiming penalties (including the true value) and estimate the corresponding shares for each education–marital group.

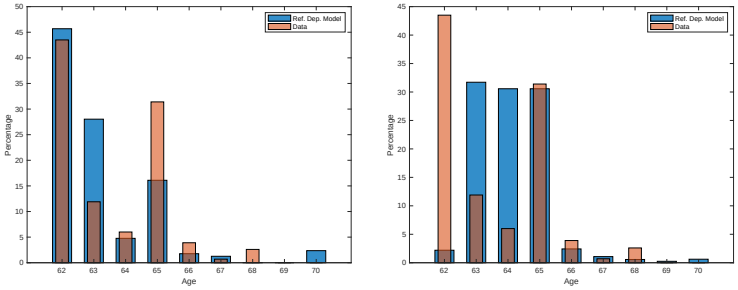
In this robustness exercise, we instead construct both objects using the Health and Retirement Study (HRS). The HRS elicits respondents' expected benefit amounts if they were to claim at multiple ages. By combining these responses across ages, we recover an implied perceived annual early-claiming penalty for each respondent (as described in Section 2). Unlike the UAS question, this implied penalty is not restricted to a small set of discrete response options. To keep the state space and computational burden comparable to the benchmark specification, we discretize the empirical HRS distribution of implied penalties into four bins and represent each bin by a single

Figure J.1: Model Fit for Singles with Reference-Dependent Claiming

(a) Ref. Age 62: Non-College (b) Ref. Age 65: Non-College



(c) Ref. Age 62: College (d) Ref. Age 65: College



Notes: Claiming-age model fit for single individuals under two re-estimated reference-dependent variants (program-misbelief channel switched off). The "Ref. Age 65" version uses a one-parameter losses-only (simple loss aversion) specification around the NRA; the "Ref. Age 62" version imposes a distance penalty from the EEA to allow 62 to be focal. With a single reference age, the model matches the peak at the reference age but misses the other peak.

support point (the group-level mean penalty within the bin). We then compute the corresponding shares within each education–marital group. The four belief types and their empirical distribution from the HRS is given in figure J.2 and Table J.1 for each education and marital group.

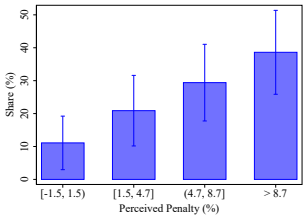
A further difference is that, while the UAS elicits confidence in the reported beliefs, the HRS does not include an analogous confidence measure. This leaves two approaches: (i) jointly estimate the diagonal "own-belief" weight (confidence) along with the off-diagonal uncertainty weights using claiming moments, or (ii) fix the diagonal weight using external information and estimate only the off-diagonal weights, as in the benchmark. We adopt the second approach to keep the number of free parameters comparable to the benchmark model: we set the diagonal confidence weight, $\varepsilon^{ss}_{k,k}$, to the corresponding group mean from the UAS and estimate only the off-diagonal weights to match the same Social Security claiming moments.

Section (I) of Table J.2 compares this HRS-based belief specification to the benchmark (UAS-based) specification using three summary statistics. First, we report the education and marital claiming-age gaps implied by the *full model*. Second, we report the corresponding gaps in *Baseline B + program beliefs* under each belief specification, which isolates the role of program beliefs holding all other mechanisms fixed. Third, we compute the implied contribution of program beliefs to gap compression, defined as the reduction in each gap relative to Baseline B as a share of the total gap reduction between Baseline B and the full model.

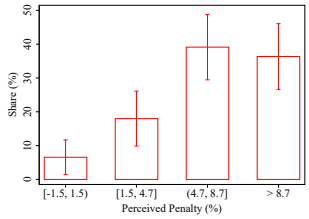
Across all three statistics, the HRS-based results are very close to the benchmark. In particular, the claiming-age gaps implied by the full model, the gaps implied by Baseline B with program beliefs, and the fraction of gap compression attributable to program beliefs are all remarkably similar under the two belief distributions. This robustness exercise therefore indicates that our main conclusions do not depend on whether the empirical distribution of perceived penalties is taken from the UAS or inferred from the HRS.

Figure J.2: HRS Distribution over Belief Types

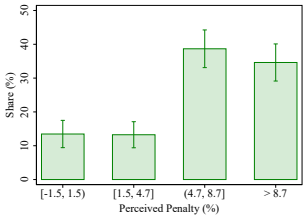
(a) Non-college, singles



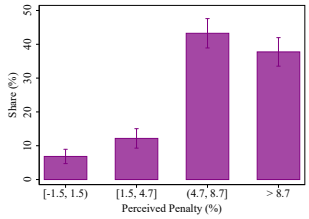
(b) College, singles



(c) Non-college, married



(d) College, married



J.3 Distance-Based Uncertainty Weighting

In the benchmark model, conditional on the survey-measured program belief type and reported confidence in stated belief, we estimate a set of off-diagonal probabilities governing how the residual uncertainty mass is distributed across the remaining penalty options. In this robustness, we replace these free off-diagonal weights with a *parametric distance kernel* that is governed by a single scalar for each stated belief type.

Concretely, for each education–marital group $g \equiv (e, m)$ and belief type k, we continue to pin down the on-diagonal weight using the UAS confidence/knowledge measure:

$$\epsilon^{ss}_{k,k}(g) \equiv ACS_{gk}.$$

We then restrict the off-diagonal weights to follow a symmetric distance-based logit kernel over the fixed survey menu $\mathcal{G}$:

$$\epsilon^{ss}_{k,j}(g) = (1 - \epsilon^{ss}_{k,k}(g)) \cdot \frac{\exp\{-\tau_{gk} \, |\gamma^{ss}_{1,j} - \gamma^{ss}_{1,k}|\}}{\sum_{h \neq k} \exp\{-\tau_{gk} \, |\gamma^{ss}_{1,h} - \gamma^{ss}_{1,k}|\}}, \qquad j \neq k, \tag{23}$$

where $\gamma^{ss}_{1,j} \in \mathcal{G}$ denotes the discrete annual penalty values in the survey and τ_{gk} is a type-specific concentration parameter. This formulation preserves the probability structure by construction

($\epsilon^{ss}_{k,j} \geq 0$ and $\sum_{j \in \mathcal{G}} \epsilon^{ss}_{k,j} = 1$) and does not permit beliefs outside the survey support. Importantly, it collapses the beliefs flexibility from eight free off-diagonal weights per education and marital group (two per belief type, with the third implied by adding-up) *to four* parameters per group (one τ_{gk} per type), which are then *overidentified* by our eight claiming moments (cumulative claiming at ages 62–69).

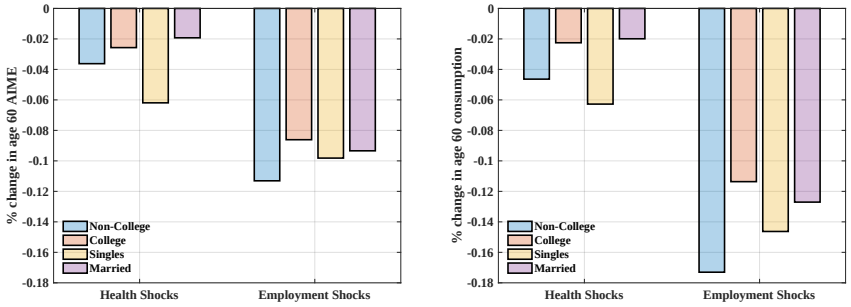
Despite the substantial restriction on program belief flexibility, the model continues to match the claiming profile closely and the decomposition conclusions are unchanged. Section (II) of Table J.2 compares this restricted distance-based uncertainty weighting belief specification to the benchmark specification using same three summary statistics as above. Across all three statistics, the distance-based results are very close to the benchmark. In particular, the claiming-age gaps implied by the full model, the gaps implied by Baseline B with program beliefs, and the fraction of gap compression attributable to program beliefs are all remarkably similar under the two belief structures. This robustness exercise therefore indicates that our main conclusions do not depend on how we model the uncertainty in program beliefs.

We further experimented with even tighter restrictions (e.g., a single τ_g common across types, or pooling across types), but these restrictions materially worsen claiming fit in some groups. This pattern is informative: it indicates that the data require some heterogeneity across stated belief types in how residual uncertainty is allocated, and the 4-τ kernel is a disciplined and parsimonious way to accommodate this without reverting to free-form off-diagonal weights.

K Additional figures and tables

Figure K.1: Effect of health and employment shocks

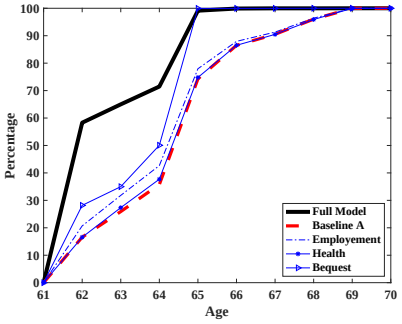
(a) AIME (b) Consumption



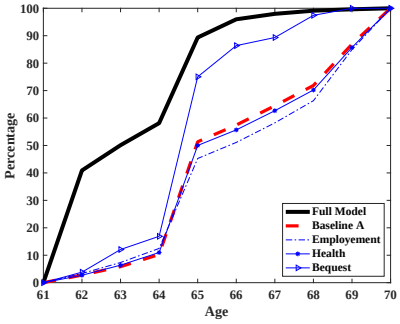
Notes: For each experiment involving the addition of health or employment shocks over the no mechanisms baseline, the percentage change in Age 60 AIME and consumption is reported relative to the no-mechanisms baseline.

Figure K.2: Effect of standard mechanisms: SS benefit claiming by education and marital status

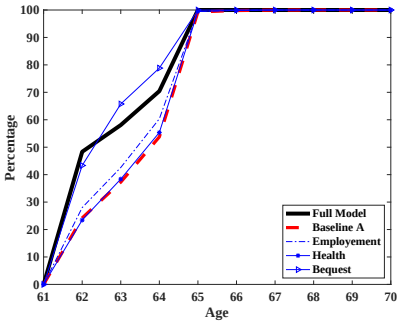
(a) Non-college



(b) College



(c) Singles



(d) Married

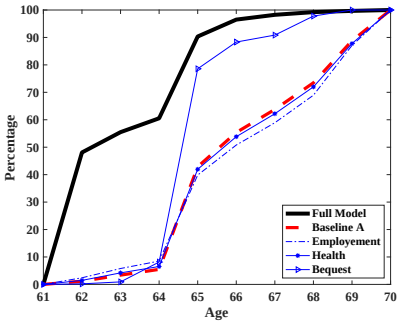


Table B.3: Regression Results

	Cohort Controls	+ Demographic Controls	+ Additional Controls
	Dependent variable: SS Claiming Age		
Survival Beliefs			
very pessimistic	-0.220*	-0.195*	-0.183
	(0.119)	(0.118)	(0.163)
pessimistic	-0.00743	0.0642	0.216
	(0.112)	(0.111)	(0.153)
optimistic	0.148	0.181*	0.309**
	(0.105)	(0.104)	(0.145)
very optimistic	0.130	0.322**	0.450**
	(0.122)	(0.128)	(0.175)
Penalty Beliefs			
much larger penalty	0.378***	0.363***	0.258*
	(0.103)	(0.102)	(0.144)
larger penalty	0.276***	0.277***	0.232**
	(0.0881)	(0.0874)	(0.118)
smaller penalty	-0.0817	-0.0672	-0.0910
	(0.107)	(0.107)	(0.147)
much smaller penalty	-0.186	-0.143	-0.0217
	(0.114)	(0.113)	(0.172)
Control Variables			
female	-0.144**	-0.134**	0.231**
	(0.0681)	(0.0679)	(0.108)
normal retirement age	0.636**	0.546*	0.936**
	(0.311)	(0.308)	(0.412)
born 1941-1950	-0.465**	-0.447**	-0.489**
	(0.184)	(0.182)	(0.241)
born 1951-1960	-1.746***	-1.737***	-1.649***
	(0.218)	(0.216)	(0.290)
college, single		0.339**	0.362**
		(0.152)	(0.173)
non-college, married		-0.215	0.00267
		(0.133)	(0.205)
college, married		0.284**	0.322
		(0.136)	(0.210)
share spousal age gap 1-2 years			0.224
			(0.168)
share spousal age gap 3-6 years			0.104
			(0.170)
share spousal age gap 7-10 years			0.0670
			(0.215)
working at age 62			1.466***
			(0.105)
bad health at age 62			-0.230
			(0.179)
log of real wealth at age 62			0.0710**
			(0.0297)
holds a private pension			0.123
			(0.0930)
constant	22.27	28.01	0.780
	(20.33)	(20.17)	(27.02)
N	2517	2517	1142
R^2	0.078	0.096	0.250

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Public-Use Health and Retirement Study (2016, v1), authors' calculations.

Notes: Columns show the regression coefficients from regression in Equation 1. Columns represent regression specifications which differ in control: (1) controls for gender, the normal retirement age, and cohort, (2) additionally controls for education and marital status, (3) additionally controls for labor force participation, self-reported health, and the logarithm of real wealth prior to claiming; the age gap between spouses (assumed to be 0 if single); and whether the individual holds a private pension.

Table B.4: Regression Results

	Cohort Controls	+ Demographic Controls	+ Additional Controls
	Dependent variable: SS Claiming Age		
Beliefs			
penalty belief	4.119***	3.807***	2.875***
	(0.705)	(0.700)	(0.990)
survival belief	0.00502***	0.00613***	0.00678***
	(0.00139)	(0.00142)	(0.00203)
Control Variables			
female	-0.144**	-0.133**	0.223**
	(0.0677)	(0.0675)	(0.108)
normal retirement age	0.649**	0.560*	0.921**
	(0.310)	(0.308)	(0.412)
born 1941-1950	-0.458**	-0.445**	-0.482**
	(0.183)	(0.182)	(0.241)
born 1951-1960	-1.742***	-1.731***	-1.635***
	(0.217)	(0.216)	(0.290)
college, single		0.346**	0.345**
		(0.151)	(0.171)
non-college, married		-0.200	0.00337
		(0.133)	(0.204)
college, married		0.273**	0.288
		(0.133)	(0.207)
share spousal age gap 1-2 years			0.242
			(0.167)
share spousal age gap 3-6 years			0.0921
			(0.169)
share spousal age gap 7-10 years			0.0745
			(0.214)
working at age 62			1.463***
			(0.105)
bad health at age 62			-0.202
			(0.179)
log of real wealth at age 62			0.0675**
			(0.0296)
holds a private pension			0.120
			(0.0928)
constant	21.19	26.98	1.843
	(20.30)	(20.14)	(26.99)
N	2517	2517	1142
R^2	0.078	0.095	0.246

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Public-Use Health and Retirement Study (2016, v1), authors' calculations.
Notes: Columns show the regression coefficients from regression in Equation 1. Columns represent regression specifications which differ in control: (1) controls for gender, the normal retirement age, and cohort, (2) additionally controls for education and marital status, (3) additionally controls for labor force participation, self-reported health, and the logarithm of real wealth prior to claiming; the age gap between spouses (assumed to be 0 if single); and whether the individual holds a private pension.

Table B.5: Average Belief by Quintile

Quintile	Subjective Survival Deviation	Perceived Penalty
Q1	-39.184	-0.013
Q2	-13.715	-0.053
Q3	0.084	-0.074
Q4	14.663	-0.094
Q5	36.814	-0.144

Notes: Shows the percentage point deviation between respondents' subjective probability of survival from age 60 to age 75 and the objective probability for their education and marital status group and the average perceived annual penalty for early Social Security claims within each quintile.

Figure K.3: Effect of subjective beliefs: SS benefit claiming by education and marital status

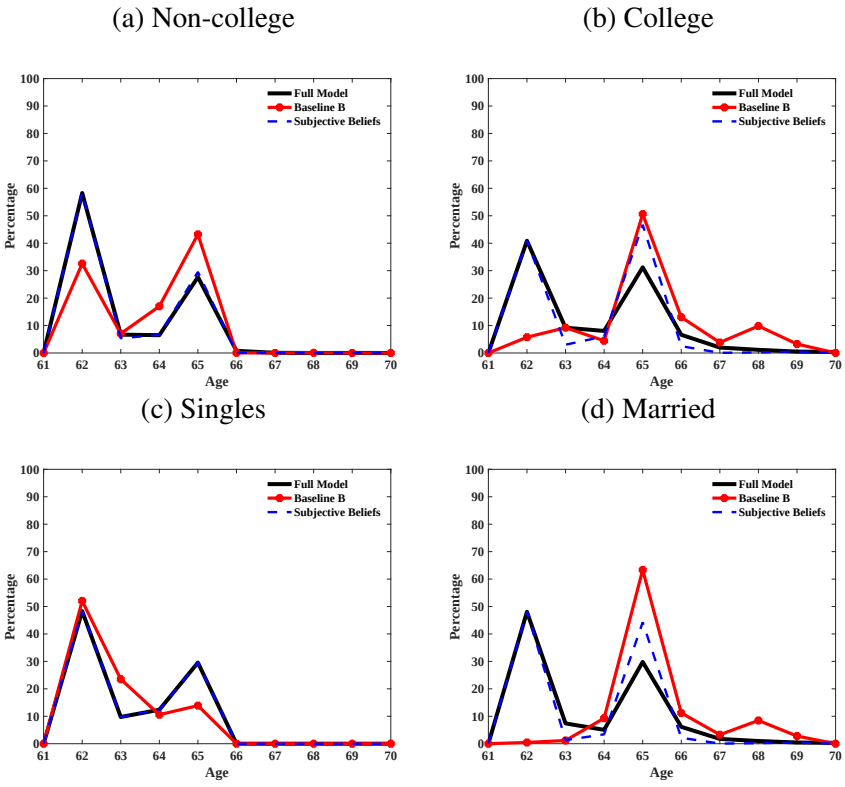


Table B.6: Regression Results

	Cohort Controls	+ Demographic Controls	+ Additional Controls
	Dependent variable: SS Claiming Age		
Survival Beliefs			
Q1	-0.281***	-0.273***	-0.278*
	(0.106)	(0.106)	(0.146)
Q2	-0.107	-0.0670	0.0713
	(0.102)	(0.102)	(0.141)
Q4	0.0201	0.0150	0.129
	(0.101)	(0.0999)	(0.140)
Q5	0.0145	0.188	0.267*
	(0.110)	(0.115)	(0.156)
Penalty Beliefs			
Q1	-0.110	-0.0978	-0.111
	(0.110)	(0.110)	(0.156)
Q2	0.0877	0.0705	0.0407
	(0.111)	(0.110)	(0.154)
Q4	0.200*	0.180*	0.0487
	(0.106)	(0.106)	(0.147)
Q5	0.474***	0.452***	0.322**
	(0.110)	(0.109)	(0.154)
Control Variables			
female	-0.144**	-0.135**	0.240**
	(0.0680)	(0.0678)	(0.108)
normal retirement age	0.657**	0.576*	0.939**
	(0.311)	(0.308)	(0.413)
born 1941-1950	-0.463**	-0.451**	-0.489**
	(0.184)	(0.182)	(0.242)
born 1951-1960	-1.751***	-1.749***	-1.649***
	(0.218)	(0.216)	(0.291)
college, single		0.328**	0.322*
		(0.152)	(0.173)
non-college, married		-0.226*	-0.0289
		(0.133)	(0.205)
college, married		0.263*	0.273
		(0.136)	(0.210)
share spousal age gap 1-2 years			0.242
			(0.168)
share spousal age gap 3-6 years			0.0960
			(0.170)
share spousal age gap 7-10 years			0.0714
			(0.216)
working at age 62			1.461***
			(0.105)
bad health at age 62			-0.225
			(0.179)
log of real wealth at age 62			0.0678**
			(0.0297)
holds a private pension			0.127
			(0.0932)
constant	20.91	26.12	0.772
	(20.35)	(20.19)	(27.07)
N	2517	2517	1142
R^2	0.078	0.095	0.248

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Public-Use Health and Retirement Study (2016, v1), authors' calculations.

Notes: Columns show the regression coefficients from regression in Equation 1. Columns represent regression specifications which differ in control: (1) controls for gender, the normal retirement age, and cohort, (2) additionally controls for education and marital status, (3) additionally controls for labor force participation, self-reported health, and the logarithm of real wealth prior to claiming; the age gap between spouses (assumed to be 0 if single); and whether the individual holds a private pension.

Table E.1: Panel Study of Income Dynamics sample selection

	Observations	Individuals
Initial sample	3,385,493	82,573
SRC sample	1,429,096	34,856
1968-2017	1,359,384	34,856
Interviewed twice	1,010,175	33,725
Male household head (and wife, if applicable)	122,317	9,683
Ages 25-99	115,057	9,208
Born 1902-1965	95,534	5,593

Table E.2: Health and Retirement Study sample selection: male sample
(Respondent Cross-Year Benefits Files)

	Observations
Initial sample	43,559
Male household head	14,394
Received SS retirement benefits between age 62-70	5,097
Born 1926-1939	2,296

Table E.3: Health and Retirement Study sample selection: spouse sample
(Respondent Cross-Year Benefits Files)

	Observations
Initial sample	10,384
Male household head	5,280
Head claims SS retirement, dual, or auxiliary benefits	3764
Born 1926-1940	1014
Head claims SS retirement benefits on his own record between ages 62-70	539
Spouse claims SS retirement, dual, or auxiliary benefits	254

Table E.4: Health and Retirement study sample selection
(Summary Earnings File)

	Observations	Individuals
Initial sample	889,627	43,558
Male household head	103,040	1,457
Born 1921-1939	89,810	1,270
Spouse earnings ≥ 0	72,545	1,270

Table E.5: Understanding America Study sample selection

	Observations	Individuals
Initial sample	17,520	10,727
Nationally representative sample	12,899	7,975
Ages 25-69	10,593	6,651
Male	4,305	2,702
Not currently receiving SS retirement benefits	3,663	2,347
Answered misinformation questions	3,648	2,339

Table F.1: Ordered probit estimates of health transitions

	(1)
Health Lag=1	1.454***
Age	0.0213
$\text{Age}^2/10^2$	0.0186
$\text{Age}^3/10^3$	-0.00337*
College=0 $\times$ married=1	-0.149***
College=1 $\times$ married=0	-0.369***
College=1 $\times$ married=1	-0.623***
cut1	2.232***
Observations	85,853
Pseudo R^2	0.22

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

Table F.2: Life expectancy comparisons: HRS vs. SSA Life Tables

Life expectancy at	HRS objective	Life Tables
age 50	28.0	27.1 - 28.1
age 60	20.3	20.0 - 21.0
age 70	13.2	13.1 - 14.1

Notes: life expectancy at ages 50, 60, and 70 from both the HRS and life tables is reported for men born between 1931 and 1935. The group-specific weights used to compute average life expectancy from the HRS are 17.4, 14.5, 24.3, and 43.7 for non-college single, college single, non-college married, and college married men, respectively. Because life expectancy from life tables can be calculated for each single birth cohort, we report a range spanning those born in 1931 through 1935.

Table F.3: Life expectancy comparisons: Objective and Subjective

Life expectancy at	All		Bad Health		Good health	
	Objective	Subjective	Objective	Subjective	Objective	Subjective
age 50						
Non-college singles	21.9	27.3	15.1	26.4	31.1	29.0
College singles	27.3	29.7	17.5	26.8	33.1	30.9
Non-college married	26.9	27.7	19.0	26.2	35.2	29.3
College married	31.7	29.9	19.4	25.9	36.0	31.2
age 60						
Non-college singles	15.4	17.5	11.3	16.6	23.2	19.3
College singles	19.1	19.8	12.0	17.2	24.3	20.9
Non-college married	19.2	17.9	14.2	16.4	26.4	19.4
College married	23.3	20.0	14.6	16.2	27.2	21.2
age 70						
Non-college singles	9.9	9.4	7.8	8.6	15.5	11.2
College singles	12.3	11.0	8.2	9.6	16.3	11.8
Non-college married	12.2	9.7	9.4	8.6	17.7	10.9
College married	15.3	11.1	9.9	8.4	18.4	12.1

Notes: objective and subjective life expectancy at ages 50, 60, and 70 from the HRS are reported for men born between 1931 and 1935, both for the overall population in this birth cohort and conditional on health status at the given ages. "Good health" refers to the top three categories of self-reported health, while "bad health" corresponds to the bottom two categories.

Table F.4: Regression results: spousal income estimation

	Spousal LFP		Spousal Labor Earnings (thousands of 1998 dollars)	
age	0.0351**	(0.0136)	-2.081**	(0.918)
$age^2/10^2$	-0.0183	(0.0500)	8.241**	(3.316)
$age^3/10^4$	-0.0692	(0.0738)	-11.32**	(4.930)
$age^4/10^6$	0.0492	(0.0376)	4.619*	(2.595)
college*age	-0.000506	(0.000553)	0.131***	(0.0274)
labor income (thousands of 1998 $)*$age$	0.00130***	(0.000228)	0.0362***	(0.00738)
labor income *$age^2/10^2$	-0.00966***	(0.00139)	-0.226***	(0.0469)
labor income *$age^3/10^4$	0.0212***	(0.00279)	0.446***	(0.0974)
labor income *$age^4/10^6$	-0.0142***	(0.00184)	-0.278***	(0.0661)
spousal age gap (0,2] years*age	0.00272***	(0.00105)	0.0733	(0.0471)
spousal age gap (2,6] years*age	0.00554***	(0.000999)	0.123***	(0.0451)
spousal age gap (6,10) years*age	0.00967***	(0.00126)	0.304***	(0.0733)
constant	-0.398***	(0.119)	19.37**	(8.817)
N	61158		25848	
R^2	0.110		0.173	

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Source: Health and Retirement Study (2016, v1) linked to SSA Administrative Records, authors' calculations.

Notes: regression coefficients from regressing spousal labor force participation (left columns) and real spousal labor income—conditional on positive labor income—in thousands of 1998 dollars (right columns) on a set of controls. All controls variable values are determined characteristics the head of household.

Table F.5: Estimated coefficients from logarithm of hourly wages

	Wage Equation	
age	0.0743***	(0.00249)
$age^2/10^2$	-0.0879***	(0.00375)
college*age	0.0108***	(0.000358)
college*$age^2/10^2$	-0.00475***	(0.000755)
married*age	0.00417***	(0.000201)
born 1946-2000	-0.0665***	(0.0102)
born 1941-1945	-0.00117	(0.0113)
born 1936-1940	-0.0229*	(0.0124)
born 1926-1930	0.0326**	(0.0129)
born 1921-1925	0.0801***	(0.0144)
born 1916-1920	0.0228	(0.0175)
born 1900-1915	-0.0257	(0.0183)
constant	1.020***	(0.0501)
	Selection Equation	
age	0.0939***	(0.00292)
$age^2/10^2$	-0.156***	(0.00327)
college*age	0.0106***	(0.000988)
college*$age^2/10^2$	-0.00780***	(0.00176)
married*age	0.00895***	(0.000314)
kids*age	-0.000401	(0.000293)
born 1946-2000	-0.132***	(0.0256)
born 1941-1945	-0.00424	(0.0296)
born 1936-1940	-0.0191	(0.0319)
born 1926-1930	0.0765**	(0.0312)
born 1921-1925	0.0733**	(0.0327)
born 1916-1920	-0.0143	(0.0361)
born 1900-1915	0.185***	(0.0355)
constant	-0.115*	(0.0670)
Inverse Mills Ratio	0.202***	(0.0464)
N	96768	

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Source: Panel Study of Income Dynamics, authors' calculation

Notes: dependent variable of the wage equation: natural logarithm of the real hourly wage (1998 dollars); dependent variable of the selection equation: labor force participation. Standard errors in parentheses.

*p< 0.1, **p< 0.05, ***p< 0.01

Table G.1: Shares by fixed-type

Non-college, single	College, single	Non-college, married	College, married
0.18	0.15	0.23	0.43

Notes:The share within each fixed type is calculated from the PSID based on education and marital status for individuals between ages 25 and 30 in the 1931–1935 birth cohort. The numbers may not add up to one due to rounding.

Table G.2: Estimation results: uncertainty over SS penalty beliefs $\left(\epsilon_{ss}^{k,j}\right)$

Type (e, m, k)	**SS Penalty Outcome** $(\widetilde{\gamma}_{\mathbf{1}}^{\mathbf{j}})$ **No penalty**	1% **penalty**	6% **penalty**	15% **penalty**
no college, single				
no penalty	0.4640	0.1096	0.0479	0.3785
1% penalty	0.0264	0.4950	0.0021	0.4765
accurate beliefs (∼6% penalty)	0.0097	0.0308	0.5560	0.4035
15% penalty	0.0599	0.0786	0.3545	0.5070
college, single				
no penalty	0.6510	0.0005	0.0002	0.3483
1% penalty	0.0618	0.5350	0.0433	0.3599
accurate beliefs (∼6% penalty)	0.0320	0.0310	0.7100	0.2270
15% penalty	0.0017	0.0006	0.1507	0.8470
no college, married				
no penalty	0.6040	0.0993	0.1166	0.1800
1% penalty	0.1310	0.5920	0.1294	0.1475
accurate beliefs (∼6% penalty)	0.0346	0.0891	0.6580	0.2183
15% penalty	0.1517	0.1503	0.1520	0.5460
college, married				
no penalty	0.6580	0.1591	0.1091	0.0738
1% penalty	0.0194	0.6710	0.0622	0.2473
accurate beliefs (∼6% penalty)	0.0316	0.0719	0.7560	0.1404
15% penalty	0.0550	0.1067	0.1193	0.7190

Table G.3: Preference parameters

Parameter	Non college			College		
	Full Model	Restricted Beliefs Model		Full Model	Restricted Beliefs Model	
		Survival	Program		Survival	Program
Singles						
β	0.963	0.963	0.959	0.975	0.966	0.965
ρ	3.711	3.706	3.639	2.459	2.461	2.413
ν	0.578	0.576	0.567	0.579	0.565	0.558
θ_{beq}	0.782	0.781	0.797	0.988	1.068	1.054
κ_{beq} (in 000s)	217.7	214.2	215.5	106.7	110.0	101.1
$\bar{l}$	5053	5043	5084	4504	4593	4652
Married						
β	1.030	1.028	1.026	1.071	1.067	1.070
ρ	1.837	1.743	1.871	2.519	2.620	2.627
ν	0.528	0.526	0.494	0.662	0.673	0.660
θ_{beq}	1.245	1.175	0.920	1.659	1.579	1.626
κ_{beq} (in 000s)	47.6	40.3	45.0	67.4	76.1	69.3
$\bar{l}$	4574	4541	4647	4354	4410	4349

Restricted survival beliefs model refer to the full model but with objective survival probabilities. Restricted program beliefs model refer to the full model but with full information related to SS early claiming penalty and no uncertainty about those beliefs.

Table H.1: Social Security early claiming share: role of different channels

Experiment	All	Education status		Marital status	
		Non-College	College	Singles	Married
Data	65.4	72.8	57.3	67.2	62.2
Baseline A (No-Mechanisms)	21.3	35.4	11.2	53.4	6.2
Health Risk	23.2	37.6	12.8	55.9	7.8
Employment Shocks	25.2	42.3	12.9	60.0	8.8
Bequest Motive	32.5	51.5	18.9	78.8	10.7
Baseline B (Health+Employment+Bequest)	36.6	58.2	21.1	86.3	13.3
(1) Baseline B + Marital Benefits					
Spousal	58.6	69.0	51.2	86.3	45.6
Survivors	29.6	42.0	20.8	86.3	3.0
Spousal & Survivor	54.0	59.2	50.2	86.3	38.8
(2) Baseline B + Program Misbeliefs					
Fixed Beliefs	66.3	80.5	56.1	87.2	56.4
Beliefs & Uncertainty	58.2	71.1	48.9	76.2	49.7
(3) Baseline B + Mortality Misbeliefs	47.4	79.2	24.6	84.9	29.8
(4) Baseline B + Subjective Beliefs (2) + (3)	58.6	71.1	49.7	70.3	53.1
Full Model: Baseline B + (1) + (2) + (3)	64.1	71.7	58.6	70.3	61.1

Notes: *Fraction claiming Social Security benefits early (before age 65) across groups is reported for each experiment. Each experiment listed under (1)-(3) entails including given channel to the standard mechanisms of Baseline B to understand its impact on claiming. No mechanisms (Baseline A) case refers to a scenario in which all mechanisms are omitted.

Table H.2: Social Security claiming behavior: role of different channels

Experiment	Singles		Married	
	Non-College	College	Non-College	College
Data	63.1	63.5	63.1	63.6
Baseline A (No-Mechanisms)	63.5	64.3	65.6	67.2
Health Risk	63.5	64.3	65.6	67.3
Employment Shocks	63.3	64.2	65.4	67.5
Bequest Motive	62.6	63.7	64.7	65.7
Baseline B (Health+Employment+Bequest)	62.4	63.4	64.6	65.8
(1) Baseline B + Marital Benefits				
Spousal	62.4	63.4	64.3	65.1
Survivors	62.4	63.4	64.9	66.6
Spousal & Survivor	62.4	63.4	64.6	65.4
(2) Baseline B + Program Misbeliefs				
Fixed Beliefs	62.4	62.8	63.0	64.1
Beliefs & Uncertainty	62.8	63.3	63.1	64.3
(3) Baseline B + Mortality Misbeliefs	62.5	63.5	64.1	65.0
(4) Baseline B + Subjective Beliefs (2) + (3)	63.1	63.4	63.1	63.8
Full Model: Baseline B + (1) + (2) + (3)	63.1	63.4	63.0	63.8

Notes: *Average SS benefit claiming age across groups is reported for each experiment. Each experiment listed under (1)-(3) entails including given channel to the standard mechanisms of Baseline B to understand its impact on claiming. No mechanisms (Baseline A) case refers to a scenario in which all mechanisms are omitted.

Table I.1: Robustness tests
Changes in SS claiming behavior

	Average	Education			Marriage		
		Non-college	College	Gap	Singles	Married	Gap
Benchmark	63.4	63.1	63.7	+0.6	63.2	63.5	+0.3
(1) Fixed Bequest	63.5	63.0	63.8	+0.8	63.2	63.6	+0.4
(2) Medical Expenditures	63.5	63.1	63.8	+0.7	63.2	63.6	+0.4
(3) Sousal Claiming	63.6	63.1	63.9	+0.8	63.2	63.8	+0.5
(4) Earnings Test	63.4	63.0	63.7	+0.7	63.2	63.6	+0.4

Notes: *average claiming ages and claiming age gap between education and marital groups is reported for each scenario. (1) $\kappa_{beq} = 89,570, \theta_{beq} = 1.29$ for all groups (2) medical expenditures (3) separate claiming decisions for both husband and wife (4) subjective program beliefs also apply to earnings test.

Table I.2: Fixed bequest motives: preference parameters

Parameter	Description	Singles		Married	
		Non-college	College	Non-college	College
β	discount factor	0.951	0.969	1.031	1.074
ρ	relative risk aversion	3.102	2.357	1.868	2.643
ν	consumption weight	0.591	0.582	0.524	0.660
$\bar{l}$	time endowment	4939	4581	4563	4463

Table I.3: Medical expenditures: preference parameters

Parameter	Description	singles		married	
		Non-college	college	Non-college	College
β	discount factor	0.959	0.965	1.029	1.070
ρ	relative risk aversion	3.639	2.413	1.831	2.603
ν	consumption weight	0.567	0.558	0.532	0.658
θ_{beq}	bequest intensity	0.797	1.054	1.215	1.645
κ_{beq} (in 000s)	bequest curvature	215.5	101.1	42.8	69.3
$\bar{l}$	time endowment	5084	4652	4566	4499

Table J.1: HRS belief robustness: Program misinformation shares and confidence

		Belief Type, k			
		$\tilde{\gamma}_1^k = 0.0$	$\tilde{\gamma}_1^k = 0.034$	$\tilde{\gamma}_1^k = 0.067$	$\tilde{\gamma}_1^k = 0.123$
Non-college, single	Share	11.1	20.9	29.4	38.6
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.43	0.43	0.43	0.43
College, single	Share	6.6	18.0	39.1	36.3
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.60	0.60	0.60	0.60
Non-college, married	Share	13.5	13.2	38.7	38.7
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.59	0.59	0.59	0.59
College, married	Share	6.8	12.2	43.2	37.8
	ACS $\left(\varepsilon_{k,k}^{ss}\right)$	0.67	0.67	0.67	0.67

Source: The Health and Retirement Study, authors' calculations.

Notes: For each education-marital status group, the columns report stated annual early claiming penalty,γ_1^k, for belief type k. The top line (Share) reports the distribution of responses across the four discretized belief types from the survey. The second line (ACS) shows the average confidence scale for the group from UAS.

Table J.2: Alternate program belief modeling: Social Security claiming age gaps

	Claiming age-gap	
	Education	Marriage
Benchmark beliefs		
Baseline B + Program beliefs	1.1	0.9
Gap reduction explained (%)	48.6	75.1
Full Model	0.6	0.3
(I) *HRS belief distribution*		
Baseline B + Program beliefs	1.1	0.7
Gap reduction explained (%)	46.4	79.6
Full Model	0.6	0.3
(II) *Restricted belief uncertainty*		
Baseline B + Program beliefs	1.1	0.9
Gap reduction explained (%)	45.5	75.4
Full Model	0.7	0.3

Notes: claiming age gap between education and marital groups is reported for each scenario. Gap reduction explained is the difference between the claiming age gap under Baseline B and after introducing a given channel. Total gap is the difference between Baseline B and the full model (education: $0.9 = 1.5 - 0.6$; marital: $2.3 = 2.6 - 0.3$). Percent reduction explained = gap reduction / total gap (e.g., for program beliefs and education: $(1.5 - 1.1)/0.9 \approx 49\%$).

Table K.1: Share in Bad Health at Age 60

	bad health share
Non-college singles	36.0
College singles	16.8
Non-college married	29.3
College married	9.7
By Education:	
Non-college	27.2
College	16.7
By Marital Status:	
Single	32.1
Married	11.4
Overall:	20.0

Supplementary References

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., and Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, 128(6):2375–2451.

Bissonnette, L., Hurd, M. D., and Michaud, P.-C. (2017). Individual survival curves comparing subjective and observed mortality risks. *Health Economics*, 26:e285–e303.

Borella, M., De Nardi, M., and Yang, F. (2023). Are marriage-related taxes and social security benefits holding back female labour supply? *The Review of Economic Studies*, 90(1):102–131.

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2021). Why do couples and singles save during retirement? Technical report, National Bureau of Economic Research.

French, E. (2005). The effects of health, wealth, and wages on labor supply and retirement behavior. *The Review of Economic Studies*, 72(2):395–427.

Gustman, A. L., Mitchell, O. S., Samwick, A. A., and Steinmeier, T. L. (1998). Evaluating pension entitlements.

Gustman, A. L. and Steinmeier, T. L. (1999). Effects of pensions on savings: Analysis with data from the health and retirement study. In *Carnegie-Rochester Conference Series on Public Policy*, volume 50, pages 271–324. Elsevier.

Heathcote, J., Storesletten, K., and Violante, G. L. (2010). The macroeconomic implications of rising wage inequality in the united states. *Journal of political economy*, 118(4):681–722.

Lockwood, L. M. (2012). Bequest motives and the annuity puzzle. *Review of economic dynamics*, 15(2):226–243.

O'Dea, C. and Sturrock, D. (2023). Survival pessimism and the demand for annuities. *The Review of Economics and Statistics*, 105(2):442–457.

Rutledge, M. S., Sass, S. A., and Ramos-Mercado, J. D. (2017). How does occupational access for older workers differ by education? *Journal of Labor Research*, 38(3):283–305.

Simmons, M. (2023). Job-to-job transitions, job finding and the ins of unemployment. *Labour Economics*, 80:102304.