Accounting for Interacting Explanations of Annuity Demand*

Derek Tharp†

Abstract

Lifecycle theory predicts that retirees should insure longevity risk with annuities, yet private ownership remains uncommon. Explanations include bequest motives, actuarially unfair pricing, pre-existing Social Security and pension income, and medical spending risk. Because the mechanisms interact, adding them one at a time credits their joint effects to whichever enters last. This paper solves a lifecycle model under all 512 combinations of nine mechanisms and uses exact Shapley values to average each contribution over every entry order. Purchase propensity falls from 100.0% in an equal-resource benchmark, with every mechanism off and existing retirement income held as liquid wealth of equal value, to 37.7% in the full model. Entered last, medical expenditures with health-dependent mortality raise purchase propensity from 29.0% to 37.7%. Averaged over every entry order, the same pair lowers it by 8.4 percentage points, so a single sequence can reverse a mechanism's apparent sign. Pricing frictions account for 28.1 points of the decline, pre-existing annuitization for 15.1, and the medical pair for 8.4. Bequest motives contribute 2.7 points and rank fifth at baseline. Bequest motives never exceed pricing frictions in any reported recomputation that treats pricing frictions as one mechanism, although other rankings change with preferences and institutions. The full model's 37.7% purchase propensity remains far above the 3.85% life-annuity ownership rate in the Health and Retirement Study sample.

JEL Codes: D14, D15, D91, G22, H55, J14

Keywords: annuity puzzle, lifecycle model, Shapley decomposition, stochastic mortality, pre-existing annuitization

*A replication package is available to the editors and referees on request and will be deposited in a public repository at acceptance.

†Department of Accounting & Finance, School of Business, University of Southern Maine, 85 Bedford Street, Portland, ME 04101, USA. Corresponding author: derek.tharp@maine.edu; telephone: +1 207 780 4837.

1 Introduction

Yaari (1965) showed that a household with no bequest motive and access to fairly priced life annuities should use them to insure its consumption against the risk of a long life. Davidoff et al. (2005) found that demand for at least partial annuitization survives many departures from that complete-market benchmark. Reported annuity ownership nonetheless remains uncommon. In the model-eligible Health and Retirement Study (HRS) sample used here, 3.85% of observations report an annuity payment that they state will continue for life, and 6.12% report any annuity income. Both figures describe payments from annuities that people already hold, and the survey does not record when or where those annuities were bought. This paper instead models whether a retiree makes a one-time purchase of a single-premium immediate annuity (SPIA).

The literature offers many mechanisms: pre-existing public and employer retirement income, bequests, medical expenditures, health-dependent mortality, survival beliefs, product pricing, nominal erosion, consumption profiles, and health-dependent preferences (Lockwood, 2012; Reichling and Smetters, 2015; Mitchell et al., 1999; O'Dea and Sturrock, 2023; Aguiar and Hurst, 2013; Finkelstein et al., 2013; Hershfield et al., 2026). Conventional within-model exercises add the mechanisms one at a time, which credits each interaction to the mechanism that enters later. Across all possible entry orders here, the contribution of medical expenditures and health-dependent mortality to the decline in purchases ranges from -8.8 percentage points at the 10th percentile to 33.8 points at the 90th. Negative values mean the pair raises purchases, and entered last it contributes -8.8 points. Averaged over every order, the pair lowers purchases by 8.4 points. A single sequence can therefore reverse the apparent sign of a mechanism.

This dependence arises because the mechanisms change one another's margins. A minimum premium matters only where a household would otherwise buy an annuity smaller than the minimum. Medical risk changes the value of liquid wealth and the cost of crossing that threshold. Pre-existing annuity income changes lifetime resources and the marginal value of an additional guaranteed dollar. A mechanism added at one point in a sequence, or removed from the full model, is measured against a single set of other mechanisms. Neither approach therefore yields an allocation that is independent of order.

This paper treats the nine mechanisms as players in a cooperative game and solves the model for all $2^9 = 512$ coalitions of players. The Shapley value averages each player's contribution over every order in which the players could enter. It is the only allocation that exhausts the total change (efficiency), treats identical players identically (symmetry), assigns zero to a player that never matters (the null-player property), and adds up across games

(additivity) (Shapley, 1953). This application brings the Shapley decompositions used in economics (Shorrocks, 2013; Sastre and Trannoy, 2002) to full solutions of a lifecycle model.

When pre-existing Social Security and pension income is switched off, the model converts each respondent's assigned schedules into liquid wealth equal to their actuarial present value at the respondent's observed age. This comparison therefore changes the form in which retirement resources are held while preserving their actuarial value. The Social Security cut studied later reduces benefits without that compensation, so it changes the amount of resources and answers a separate question.

The quantity allocated is the modeled decline from an equal-resource benchmark, in which every mechanism is off but the public consumption floor remains, to the full model. Because every coalition is solved, the contributions are exact and add up to the full decline. Averaging over entry orders settles the ordering problem, but the answer still depends on the outcome measured and on how the model is built and calibrated. The allocation also does not measure causal effects of the underlying primitives or decompose the gap between the full model and observed HRS ownership.

In the focal calibration used throughout, purchase propensity falls from 100.0% in that benchmark to 37.7% in the full model. Within that decline, pricing frictions are the largest demand-suppressing contributor, followed by pre-existing annuitization and by medical expenditures with health-dependent mortality. The three contribute 28.1, 15.1, and 8.4 percentage points and together account for 83% of the modeled decline. When the outcome is instead the mean purchased share of liquid wealth, pricing frictions and pre-existing annuitization remain first and second, and bequests rise from 2.7 to 7.9 points but remain fifth. Recomputing the full decomposition under other preferences, public floors, mortality assumptions, fixed costs, or player definitions changes both the order and the membership of the leading three. Lower risk aversion and stronger survival pessimism, for example, bring pessimism into the leading three. Calibrated bequests contribute no more than pricing frictions in every reported recomputation that keeps the three pricing features together as one player. The two tie only in the above-$350k wealth-band game, whose endpoint propensities are equal.

In a partial-equilibrium policy application, a permanent 22% Social Security reduction that households anticipate from the purchase age onward—the cut projected once the retirement trust fund runs out in 2032 if no changes are made (Board of Trustees, 2026)—raises modeled purchase propensity by 5.5 percentage points. The calculation holds the employer-pension-income proxy, annuity prices, claiming behavior, and equilibrium conditions fixed and does not forecast legislation or transition dynamics.

The nearest predecessors combine several of these explanations in one model. Pashchenko

(2013) studies seven impediments together and finds pre-existing annuitization, housing, a minimum-purchase requirement, and bequests quantitatively important in that calibration. Peijnenburg et al. (2016) study background risk and incomplete markets, and Lockwood (2012) emphasizes the interaction between loads and bequest motives. This paper places nine mechanisms in one common game and separates each mechanism's order-averaged contribution from the effect of removing it from the full model. It also finds a positive bequest–pricing interaction that is larger for the mean purchased share than for purchase participation.

2 Model

2.1 Environment

Time is discrete. An individual enters at age 65 (period $t = 1$) and may survive to a maximum age of 110 (period $T = 46$). At the purchase age, before that period's consumption choice, the individual makes a single irreversible decision: what fraction $\alpha \in [0, 1]$ of current liquid wealth W_0 to convert into a single premium immediate annuity (SPIA). Before any fixed purchase cost, the remaining wealth is $(1 - \alpha)W_0$. That cost is deducted separately in the annuitization problem below. The purchase-period payment is available for consumption immediately under the model's annual timing convention, and consumption is chosen again in every later surviving period. When the model is taken to the HRS sample of single adults aged 65–69, each respondent's one-shot decision is evaluated at the observed age using an age-appropriate SPIA price. Appendix B reports the forced-age and grid diagnostics.

The state at period t has four variables: liquid wealth $W_t \geq 0$, annuity income A (fixed after the evaluated one-time purchase), health status $H_t \in \{1, 2, 3\}$ (Good, Fair, Poor), and age t. Pre-existing annuitized income is $Y_t^{\text{pre}} = \text{SS}_t + P_t$, where Social Security is real and the employer-pension-income proxy is modeled as a nominal life-contingent flow whose real value erodes at rate π_P. The pension-proxy rate is fixed independently of the private-SPIA inflation player and equals its focal rate, $\pi_P = 2\%$. Employer-pension income reported in the HRS for the prior calendar year supplies the proxy (Section 3). That measure does not identify plan type or life contingency, so treating the proxy as life-contingent is a maintained modeling assumption.

2.2 Preferences

The individual maximizes expected discounted lifetime utility. Flow utility over consumption takes the constant relative risk aversion (CRRA) form:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \quad \gamma \neq 1. \tag{1}$$

The coefficient of relative risk aversion γ governs both risk aversion and the elasticity of intertemporal substitution, and the individual discounts future utility at rate β per period.

Bequest utility follows the De Nardi–French–Jones specification, which nests bequests as a luxury good:

$$v(b) = \theta \frac{(b+\kappa)^{1-\gamma}}{1-\gamma}, \tag{2}$$

where b is bequeathable wealth (liquid wealth at death; annuity income ceases), $\theta \geq 0$ controls bequest intensity, and $\kappa > 0$ governs how strongly bequests behave as a luxury good (De Nardi, 2004). Because $v'(b) = \theta(b+\kappa)^{-\gamma}$, marginal bequest utility varies little when b is small relative to κ; as resources rise relative to that scale, the optimal bequest share rises. De Nardi et al. (2010) estimate a luxury-good bequest specification jointly with medical-expense risk using the AHEAD sample; Lockwood (2012) subsequently imports a DFJ calibration into an annuity model. Neither source estimates the transformed values of θ and κ used here.

The parameters θ and κ must be translated together. Translating them to the focal curvature and 2014 dollars preserves the bequest threshold and marginal propensity to bequeath in DFJ's deterministic terminal problem. Appendix C derives the mapping. Pairing the transformed intensity with a zero shifter would make marginal bequest utility diverge near zero wealth.

2.3 Health dynamics and mortality

Health follows a first-order Markov process with three states: Good ($H = 1$), Fair ($H = 2$), and Poor ($H = 3$). Let $a_t = 64 + t$ denote chronological age in model period t. Transition probabilities are age-dependent and calibrated to HRS panel data following De Nardi et al. (2010):

$$Q_t(h, h') \equiv \Pr(H_{t+1} = h' \mid H_t = h, a_t). \tag{3}$$

Mortality depends on health through a proportional hazard specification. Baseline survival probabilities $s_{\text{base}}(t)$ come from a sex-blended SSA 2003 period life table matched to the sample's sex composition (Section 3; Social Security Administration, 2007). Focal survival

is

$$s(t, H) = s_{\text{base}}(t)^{\zeta_t \mu(H,t)}, \tag{4}$$

where $\mu(H, t)$ is a raw relative-hazard multiplier, normalized to one for Fair health, and ζ_t is a common age-specific anchoring factor. At nonterminal ages with $0 < s_{\text{base}}(t) < 1$, it is chosen so that $\sum_H \omega_H(t) s_{\text{base}}(t)^{\zeta_t \mu(H,t)} = s_{\text{base}}(t)$, with survivor health shares $\omega_H(t)$ evolved under the transition matrices and normalized objective survivals. Terminal survival is zero in every state, so no anchor is needed there. The anchored multiplier $\zeta_t \mu(H, t)$ therefore scales the baseline hazard, which preserves hazard ratios across health states while reproducing the population life table at each nonterminal age. Appendix A.3 gives the construction.

Raw multipliers are unadjusted ratios of estimated annual mortality hazards within HRS age bands. The focal schedule linearly interpolates $[0.49, 1.00, 3.29]$ at the 65–74 midpoint, $[0.60, 1.00, 2.77]$ at the 75–84 midpoint, and $[0.74, 1.00, 1.82]$ at the 85+ midpoint, with endpoint values outside that range. Section 5 reports alternatives that hold the multipliers constant across ages.

Health-dependent mortality paired with medical costs is the mechanism emphasized by Reichling and Smetters (2015). A negative health shock simultaneously increases mortality ($\mu(\text{Poor}, t) > 1$ reduces expected remaining annuity payments) and, through the medical expenditure process described below, increases the need for liquid wealth. The annuity becomes "valuation-risky," losing value in precisely the states where liquidity is most valuable. In this calibration, hazard ratios across health states are widest at the ages when the sample's purchase decisions are evaluated.

2.4 Subjective survival beliefs

O'Dea and Sturrock (2023) document substantial survival pessimism through respondents' 50s, 60s, and early 70s, followed by optimism at older ages. Pessimism reduces the perceived value of annuities by overweighting the probability of dying before recovering the premium.

The model allows the agent's subjective survival probabilities to differ from the objective rates used by the insurer to price annuities:

$$s^{\text{subj}}(t, H) = \psi \cdot s(t, H), \quad \psi \in (0, 1]. \tag{5}$$

The parameter ψ scales each one-year survival probability downward. Annuity pricing uses the health-independent population life table, whereas the agent's Bellman equation uses the health-conditioned, pessimism-scaled s^{subj}. The focal $\psi = 0.981$ is a constant annual scale chosen to represent the broad near-term pessimism that O'Dea and Sturrock (2023)

document. Their study does not report this value. Compounded from age 65, the constant annual scale lowers survival to ages 75, 85, and 95 from the population table's 80.4%, 45.5%, and 8.6% to 66.4%, 31.0%, and 4.9%. Belief errors vary with age (O'Dea and Sturrock, 2023; Heimer et al., 2019), so a constant wedge is a stylized specification. The reported range $\psi \in [0.96, 1]$ includes objective beliefs, and Appendix A.5 recomputes both the level and attribution at the strong-pessimism endpoint. At $\psi = 1$ the player is inactive.

2.5 Medical expenditures

Gross medical needs depend on age and health. Log medical needs follow:

$$\ln m_t = \mu_m(t, H_t) + \sigma_m(H_t) \cdot \varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \tag{6}$$

where μ_m is age- and health-dependent and σ_m is health-dependent. Residual innovations are iid conditional on health. Medical need is drawn before the model's generic public transfer covers any consumption-floor shortfall, and targeting combined out-of-pocket-plus-Medicaid spending (Section 3; Jones et al., 2018) is consistent with that timing.

A stylized public transfer provides a consumption floor $\underline{c}$. If total resources (wealth plus income minus medical costs) fall below $\underline{c}$, the transfer covers the shortfall. The individual consumes $\underline{c}$ and enters the next period with zero liquid wealth. Lockwood (2012) sets the floor amount to the \$6,180 annual SSI payment to a single elderly person in 2000. Restating that amount in the model's 2014-dollar unit with annual-average CPI-U (172.200 in 2000 and 236.736 in 2014; U.S. Bureau of Labor Statistics, 2017) gives \$8,496. The floor is a reduced-form stand-in for SSI and Medicaid, whose eligibility rules, enrollment, and actual benefit receipt the model does not represent.

Medical-expenditure expectations use state-split Gauss–Legendre quadrature over probability space. The split points follow the income, public-floor, and resource-accounting kinks, and the saturated upper tail is evaluated exactly. Focal integration uses order 32 on each subinterval. Appendix E reports the direct H0 order-48 check and the fully refined H2 order-48-to-64 last-step check.

2.6 Annuity pricing

The annuity is nominal: the insurer prices the stream of nominal payments using the nominal discount rate $r_{\text{nom}} = (1+r)(1+\pi) - 1$, where r is the real interest rate and π is the inflation rate. For a purchase at age a, let $\tau_a = a - 65 + 1$ be the model period and let $s_{\text{base}}(t)$ be objective survival from period t to $t+1$. Under the model's beginning-of-period payment

convention, the actuarially fair nominal payout per premium dollar is

$$q_a^N = \left[\sum_{k=0}^{T-\tau_a} \frac{\prod_{\ell=0}^{k-1} s_{\text{base}}(\tau_a + \ell)}{(1 + r_{\text{nom}})^k} \right]^{-1}, \tag{7}$$

where the empty product at $k = 0$ equals one. The exact Fisher relation is the maintained discounting convention for the nominal payment stream. Consequently, the initial payout is higher than under real discounting, but its real value declines after purchase:

$$A^{\text{real}}_{\tau_a+k} = \frac{A^{\text{real}}_{\tau_a}}{(1+\pi)^k}, \qquad k = 0, \ldots, T - \tau_a, \tag{8}$$

where π is the annual inflation rate.

When $\pi = 0$ (the Yaari benchmark and intermediate decomposition steps where inflation is inactive), $r_{\text{nom}} = r$ and the pricing reduces to the standard real-annuity formula.

The annuity is priced with a money's worth ratio (MWR), defined as the expected present value of payouts per dollar of premium. At age a, the loaded payout rate is MWR $\times\, q_a^N$. The real-product formula replaces r_{nom} with r and has no subsequent erosion. Mitchell et al. (1999) estimated population-mortality MWRs of roughly 0.80–0.85 for the US individual market. Annuitant-mortality values were higher because they condition on selection. Wettstein et al. (2021) report an average near 0.80 and a 2001–2019 range extending to roughly 0.87. The focal MWR of 0.87 is therefore a relatively favorable price at the upper end of the reported population-mortality range. Section 5 reports alternative schedules.

A fixed cost of \$2,349 is incurred upon any purchase, and annuitization requires a minimum premium of \$10,000. The cost restates the illustrative \$2,000 value in Lockwood (2012) to 2014 dollars and has no direct transaction-cost estimate behind it. Complete games at lower fixed costs and a game that separates all three pricing components show how much the headline depends on that choice. Section 5 also reports the full minimum-premium sweep.[1]

2.7 Budget constraint

Within each period, the timing is:

1. The individual enters with liquid wealth W_t, private-annuity income A_t, pre-existing annuitized income Y_t^{pre}, and health H_t.

2. Medical expenditure m_t is realized.

[1]The implementation uses beginning-of-period annual timing, a discrete-time convention under which the first annual payment arrives at purchase. The commercial immediate annuities studied by Mitchell et al. (1999) begin payments within the next month, quarter, or year after purchase.

3. The individual chooses consumption c_t.

4. Remaining wealth earns the risk-free return r.

5. The individual survives to $t+1$ with probability $s(t, H_t)$.

6. Conditional on survival, health transitions to H_{t+1}.

Let private resources after the medical draw be

$$X_t(m_t) = \max\{0, W_t + A_t + Y_t^{\text{pre}} - m_t\},$$

and let the public transfer be $G_t(m_t) = \max\{0, \underline{c} - X_t(m_t)\}$. For $t < T$, the budget constraint is

$$W_{t+1} = (1+r)\Big[X_t(m_t) + G_t(m_t) - c_t\Big], \tag{9}$$

subject to $\underline{c} \le c_t \le X_t(m_t) + G_t(m_t)$. When the floor binds, $X_t + G_t = \underline{c}$, so the household consumes the floor and carries no liquid wealth forward. Medical needs larger than private resources are absorbed outside the household budget. The model draws private resources down to zero, which is consistent with the combined out-of-pocket-plus-Medicaid calibration target, and does not track household medical payments or public benefits separately. In the terminal period, the household divides $X_T + G_T$ between consumption and the terminal bequest, which earns no further return.

2.8 Consumption and health utility weights

The model allows flow utility to vary with age and health. Age enters through

$$w(t) = (1 - \delta_c)^{t-1}. \tag{10}$$

Consumption receives a weight that falls by a fraction δ_c each year, a reduced-form way to let spending needs decline with age. The calibration sets δ_c to zero because the model's consumption path without a private annuity already declines faster than an external spending benchmark (Section 3), so this player has no effect in the focal game. Appendix A.6 finds the same zero under every alternative fitting window and seed. The player stays in the game because declining spending needs are a candidate explanation for low annuity demand, and a player with no effect leaves every other contribution unchanged. It becomes active in the joint stress test, which draws δ_c between 0 and 1.5 percent a year (Appendix H).

Health enters through

$$u(c, H) = \varphi(H)\frac{c^{1-\gamma}}{1-\gamma}. \tag{11}$$

The focal vector $[1.00, 0.92, 0.82]$ is a calibration choice that maps the chronic-disease estimates of Finkelstein et al. (2013) onto the model's self-reported health states. Appendix A.7 reports alternative mappings and explains the zero-intercept normalization.

With the focal $\gamma > 1$, the CRRA level in equation (11) is negative. A weight $\varphi(H) < 1$ therefore makes the raw level less negative while reducing marginal utility to $\varphi(H)c^{-\gamma}$. These weights are meant to capture that lower marginal utility. The less negative level in worse health is a by-product of scaling a negative number and says nothing about whether worse health raises welfare.

Public-care aversion changes effective consumption only when the public floor binds in Poor health:

$$c_{\text{eff}}(c, H, m; \mathcal{M}) = \begin{cases} \chi_{\text{LTC}} c, & \mathcal{M} = 1,\ H = \text{Poor},\ \text{and } G_t(m) > 0, \\ c, & \text{otherwise.} \end{cases} \tag{12}$$

Here $\mathcal{M}$ indicates that the medical-risk process is active, so floor top-ups with medical risk switched off do not trigger this player. A binding floor in Poor health with medical risk active is used as a proxy for reliance on publicly financed care. The model has no Medicaid state and no explicit long-term-care need, use, or choice. This player is off at $\chi_{\text{LTC}} = 1$ and on at $\chi_{\text{LTC}} = 0.49$. Applying the parameter to consumption rather than the negative CRRA utility level preserves its consumption-equivalent interpretation and the intended welfare ordering.

2.9 Bellman equation

For $t < T$, the value function satisfies:

$$\begin{aligned} V(W, A, H, t) = \mathbb{E}_m \Big[\max_c \Big\{ & w(t) \cdot u(c_{\text{eff}}(c, H, m), H) \\ & + \beta\, s^{\text{subj}}(t, H)\, \mathbb{E}_{H'} \Big[V(W', A', H', t+1) \mid H \Big] \\ & + \beta\, (1 - s^{\text{subj}}(t, H))\, v(W') \Big\} \Big], \end{aligned} \tag{13}$$

The household observes m before choosing c, so maximization is conditional on the medical draw. Conditional on survival, health transitions inside the continuation expectation. Objective survival prices the contract and governs physical simulation, while s^{subj} enters household valuation. For $t < T$, $W' = (1 + r)[X_t(m) + G_t(m) - c]$ and $A' = A$, because the state stores a fixed nominal private-annuity payment whose real value is $A_t = A/(1 + \pi)^{t-1}$. The modeled pension-income proxy erodes at its separate fixed rate π_P. At T, the household

divides $X_T + G_T$ between consumption and a terminal bequest. The terminal bequest utility retains the factor β, but the remaining balance earns no post-decision return. Appendix K states the terminal condition exactly.

For an individual evaluated in period τ at the observed age, let ρ_τ be the age-specific first-year real payout per premium dollar under the active pricing assumptions. Because the private-annuity state stores the constant nominal payment in base-age units, the annuitization decision is

$$\alpha_\tau^* = \arg\max_{\alpha \in \mathcal{A}(W_0)} V\Big((1-\alpha)W_0 - F\mathbf{1}\{\alpha > 0\},\ \alpha W_0(1+\pi)^{\tau-1}\rho_\tau,\ H_0,\ \tau\Big), \tag{14}$$

where F is the fixed purchase cost. The nominal gross-up makes the purchase-year real payment equal to $\alpha W_0 \rho_\tau$ and then lets it erode at rate π. Let $P_{\min}$ denote the minimum premium, set to \$10,000. The feasible set is

$$\mathcal{A}(W_0) = \{0\} \cup \{\alpha \in (0,1] : \alpha W_0 \geq P_{\min},\ (1-\alpha)W_0 - F \geq 0\}.$$

The numerical search evaluates all feasible nodes of the reported 101-point annuitization grid and adds the exact minimum-premium and maximum-feasible-share endpoints whenever they are feasible. It then selects the global best candidate, including zero.

2.10 Solution method

Backward induction solves the model over the state space (W, A, H, t). The wealth grid uses 98 points: a power-spaced core of 82 points through the \$3 million knot, followed by cosine-spaced tail nodes through the \$13 million support bound. Annuity income uses a 39-point grid with the same core–tail construction (31-point cubic-spaced core). Before any population statistic is computed, a citable run certifies every initial state, equal-PV off state, and feasible post-purchase state in every evaluated coalition. The same backward pass then certifies every wealth-grid interpolation stencil reachable from those states, including all health transitions and every medical quadrature node; any material need for flat extrapolation fails immediately. At each certified grid point, a global search chooses consumption given the piecewise-linear continuation value. The search compares the boundaries and interior candidates of every segment defined by the continuation-wealth knots and the source-income kink, and derivative bisection locates each segment's interior candidate when one exists. Appendix B reports refinement, the reachable-domain certificate, and an H2-versus-H3 parameter-stress audit.

The annuitization fraction α is optimized over a 101-point grid at each respondent's

observed purchase age. Grid convergence is checked at order 32 per state-split segment under the focal wealth-band income assignment. Refinement near the fixed-cost extensive margin can be non-monotone, so Appendix Table B.1 reports both purchase propensity and channel contributions.

2.11 Cooperative game and exact attribution

Let $\mathcal{N} = \{1, \ldots, J\}$ denote the $J = 9$ focal players in Table 1. For coalition $S \subseteq \mathcal{N}$, the model is solved with exactly those players in their on states. Its value is the purchase propensity in the unweighted empirical distribution of model-eligible person-wave states:

$$v(S) = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}\{\alpha_i^*(S) > 0\}. \tag{15}$$

The same coalition solutions define the mean-purchased-share value

$$v_\alpha(S) = \frac{1}{N} \sum_{i=1}^{N} \alpha_i^*(S), \tag{16}$$

the population mean purchased share of liquid wealth, including zeros. The Shapley construction below applies separately to v and v_α. Participation is the focal statistic because fixed costs and the minimum premium create a distinct extensive margin. Table 2 reports both outcomes.

Define the demand-suppression game $g(S) = v(\varnothing) - v(S)$, so $g(\varnothing) = 0$ and positive values denote reductions in purchase propensity relative to the equal-resource benchmark. Following Shapley (1953), the Shapley contribution of player j to this game is

$$\begin{aligned} \phi_j &= \sum_{S \subseteq \mathcal{N} \setminus \{j\}} \frac{|S|!\,(J - |S| - 1)!}{J!} \left[g(S \cup \{j\}) - g(S) \right] \\ &= \sum_{S \subseteq \mathcal{N} \setminus \{j\}} \frac{|S|!\,(J - |S| - 1)!}{J!} \left[v(S) - v(S \cup \{j\}) \right]. \end{aligned} \tag{17}$$

Positive ϕ_j means that the player lowers purchase propensity on average over entry orders, and a negative value means that it supports demand on average. Efficiency gives $\sum_j \phi_j = v(\varnothing) - v(\mathcal{N})$. Exact enumeration supplies all 512 coalition values, so equation (17) contains no permutation sampling error. The values depend on how the players are defined and grouped, on the focal parameters, and on which outcome, equation (15) or (16), is decomposed. A grand-coalition removal, $v(\mathcal{N} \setminus \{j\}) - v(\mathcal{N})$, is a different comparative static because it evaluates only one predecessor set.

The pre-existing-annuitization off state preserves retirement resources. Let $b(i)$ be respondent i's wealth band based on original wealth, let a_i be the observed purchase age, and let $\tau_i = a_i - 65 + 1$ be its model period. Define $q^R_{a_i}$ and $q^N_{a_i}$ as the actuarially fair, objective-survival payout rates at that age for a real flow and for a nominal flow eroding at π_P. At that age, the pension-income proxy in current real dollars is $P_{b,a_i} = P_b(1+\pi_P)^{-(a_i-65)}$, with $\pi_P = \pi$ in the focal calibration. The commuted endowment is

$$E_i = \frac{\mathrm{SS}_{b(i)}}{q^R_{a_i}} + \frac{P_{b(i),a_i}}{q^N_{a_i}}. \tag{18}$$

When the player is off, the model sets both flows to zero and evaluates the respondent at liquid wealth $W_i + E_i$. When it is on, the endowment is removed and the real Social Security flow plus modeled nominal, life-contingent pension-income proxy are restored. The same objective mortality, payment timing, and real-versus-nominal conventions price both sides. Pricing the commutation fairly leaves the amount of retirement resources unchanged and keeps the market's pricing wedge inside the pricing-frictions player. At market rates, the commutation would also change the size of the off-state endowment.

3 Calibration and Evaluation Population

Table 1 lists the players with their on and off states, and Appendix Table A.1 collects the remaining preference, demographic, health, medical, and market parameters and labels the source of each.

Preferences. The focal risk aversion is set to $\gamma = 2.5$, above the $\gamma = 2$ used by Lockwood (2012). Section 5 reports a 16-value sensitivity grid from $\gamma = 1.5$ to 5.0, denser around the focal value.

Age-varying consumption needs. The age-needs parameter δ_c is calibrated against external spending evidence, subject to $\delta_c \geq 0$. Aguiar and Hurst (2013) report an approximately 0.30-log-point decline in household nondurable expenditure excluding housing over the latter half of their observed life cycle, roughly ages 45–75. Spread over 30 years, that decline averages about 1% a year, a figure computed here from their total. Their statistic is a cohort- and family-adjusted age profile of log household expenditure, while the model's is the endpoint log change in survivor-conditioned pooled mean consumption over ages 65–90, so the comparison gives only a sense of scale. With every other focal mechanism active and $\delta_c = 0$, consumption without a private annuity already falls by 1.95% a year, faster than

Table 1: Players in the Focal Cooperative Game

Player	Off state	On state
Pre-existing annuitization	Retirement flows commuted to equal-PV liquid wealth	Real Social Security and nominal pension-income proxy
DFJ bequests	$\theta = 0$	$\theta = 158.68$, $\kappa =$ \$396,496
Medical-expenditure process and health–mortality correlation	Zero medical expenditure; population survival	Calibrated health-dependent medical expenditures and normalized health hazards
Survival pessimism	$\psi = 1$	$\psi = 0.981$
Age-varying needs	$\delta_c = 0$	Calibrated δ_c, zero at the focal calibration
State-dependent utility	$\varphi(H) = 1$	$[1.00, 0.92, 0.82]$
Pricing frictions	MWR $= 1$, $F = 0$, no minimum	MWR $= 0.87$, $F =$ \$2,349, minimum \$10,000
Nominal-payment erosion	Real private SPIA	Nominal private SPIA at $\pi = 2\%$
Public-care aversion	$\chi_{\text{LTC}} = 1$	$\chi_{\text{LTC}} = 0.49$

The empirical population and public consumption floor are fixed in every coalition. Turning the medical player off sets medical expenditure to zero, so the player activates the entire calibrated expenditure process—its mean, dispersion, and health dependence—together with health-dependent mortality. Appendix D.2 reports an alternative partition. The pension-income proxy uses employer-pension income reported in the HRS. Nominal life contingency is a maintained model assumption because the variable does not identify plan type or payment duration.

that benchmark. The calibration therefore sets δ_c at the zero boundary, so the age-needs player is null in the focal game.

State-dependent utility. The health-utility weights $\varphi = [1.00, 0.92, 0.82]$ translate the midpoint of the health-dependent-utility evidence into the model's three states. Good health receives unit weight, and Fair and Poor health reduce marginal utility proportionally. Because multiplicatively weighted CRRA utility is cardinal, the implementation also fixes all state-specific additive intercepts at zero. Such choice-independent intercepts would not change policies in this model, but they would change raw value levels and any CEV formula that failed to remove them. The reported CEVs are therefore conditional on the explicit zero-intercept convention. Appendix A.7 reports stronger $[1.00, 0.90, 0.75]$ and moderate $[1.00, 0.95, 0.85]$ mappings.

Public-care aversion (χ_{LTC}). Both the value $\chi_{\text{LTC}} = 0.49$ and the proxy state are calibration choices motivated qualitatively by Ameriks et al. (2020). The player's net effect on demand is ambiguous, because both liquid wealth and guaranteed income can help a household avoid relying on the modeled floor. Appendix A.8 states the proxy's program and care-state limitations.

Bequests. De Nardi et al. (2010) estimate $\gamma = 3.84$, $\beta = 0.97$, a bequest intensity of 2,360, and a shifter of \$273,000 in an AHEAD model whose monetary quantities are expressed in 1998 dollars. They also emphasize that the bequest parameters are imprecisely estimated. The numerical intensity is not directly portable to another curvature. In the unweighted deterministic terminal problem, the relevant allocation wedge is $q \equiv (\beta\theta)^{-1/\gamma}$. The implementation holds q at its DFJ value, setting $\theta(\gamma, \beta) = (0.97 \times 2360)^{\gamma/3.84}/\beta$, and restates the shifter from 1998 to 2014 dollars with annual-average CPI-U (U.S. Bureau of Labor Statistics, 2017). These choices give $\theta = 158.68$ and $\kappa = \$396{,}496$. Holding q fixed preserves the source terminal bequest threshold $q\kappa$ and marginal propensity $1/(1+q)$ in that stripped-down allocation problem. This normalization of units and curvature is a calibration choice, and neither DFJ nor Lockwood (2012) estimates the resulting values. State-dependent utility and the full stochastic lifecycle model need not preserve those summaries. Appendix C gives the derivation, and Appendix C.1 reports the portability audit.

Health and mortality. Health transition matrices are calibrated to HRS panel data following De Nardi et al. (2010). Five HRS self-reported health categories are mapped to three states: Good (Excellent/Very Good), Fair (Good), and Poor (Fair/Poor). Baseline phys-

ical survival uses a sex-blended life table: the survivorship mixture of the SSA 2003 male and female period tables at the model-eligible sample's female share (70.8%; Social Security Administration, 2007). RAND's binary sex field supplies that mixture weight (Appendix P gives the coding). The field enters only this life-table blend and is not interpreted as a richer measure of sex or gender identity. As a historical period table, the blend does not incorporate later mortality improvement. Another convention could change both demand levels and the complete Shapley ranking. A common age-specific factor rescales the health hazard multipliers so population-weighted physical survival reproduces the blended table while preserving hazard ratios across health states. Objective survival prices the annuity and governs forward simulation. The survival-pessimism player changes household beliefs in the Bellman equation and so deliberately separates the survival used for valuation from the survival used for pricing. As a robustness comparison, the prior convention in this literature uses the male SSA table from the Lockwood (2012) replication files ($q_{65} = 0.0187$, against 0.0120 for females) with unnormalized hazards. Appendix M recomputes the full nine-channel game under that convention.

Hazard multipliers are estimated from unweighted, complete-case person-wave intervals among all living HRS respondents aged 65 or older in the RAND HRS Longitudinal File 2022 (V1; waves 4–16) (Health and Retirement Study, 2025c), a broader sample than the focal eligibility subsample. Current-wave respondents must report valid age and health, the next wave's interview status identifies death or survival, and indeterminate interval outcomes are excluded (Appendix P lists the codes). Two-year death fractions map to constant annual hazards, with Fair health as the within-age-band reference. Respondents may contribute multiple intervals, and the broader sample stabilizes the age–health cells. Although the window includes the 2020–22 transition, only relative mortality across health states enters the model, because the age-specific normalization restores aggregate survival to the 2003 population table. The model uses the unrounded ratios, and the table below rounds them to two decimals:

Age band	Good	Fair	Poor
65–74	0.49	1.00	3.29
75–84	0.60	1.00	2.77
85+	0.74	1.00	1.82

Hazard ratios across health states are closer to one at older ages, and the calibration does not attribute that pattern to any single mechanism. Focally, the model linearly interpolates all three rows over age, using the first row below its midpoint and the last row above its midpoint (band midpoints 69.5, 79.5, and 90.0, the last assigned to the open 85+ band).

Future health-contingent mortality in the continuation value therefore reflects the older-age estimates. Separate checks impose the entry-age row $[0.49, 1.0, 3.29]$ at every age, use the same HRS estimates pooled over all ages 65 and older ($[0.57, 1.0, 2.7]$), apply a steeper researcher-specified mapping ($[0.45, 1.0, 3.5]$) to stress the Reichling and Smetters (2015) mechanism, and use a conservative mapping.

Medical expenditures. The medical process takes its scale from combined out-of-pocket-plus-Medicaid gross medical need in Jones et al. (2018), who report annual amounts over two-year intervals among surviving households. To map that household benchmark into the annual single-person model, the calibration averages two consecutive medical draws and conditions on objective survival through both interval endpoints. Two log-mean coefficients then match the reported means for the intervals beginning at ages 70 and 100 ($5,100 and $29,700). Holding the assumed dispersion and health shifts fixed implies 95th percentiles of $17,943 and $103,509, against source values of $13,400 and $111,200. Neither tail is a fitted moment, and Appendix A.2 reports both discrepancies. Health shifts both the mean and variance of log spending, with Poor health raising expected costs relative to Fair, and the public floor determines how much of a realized gross need is privately financed. Because the source mixes singles and couples and has a richer spending process, these matches only project the source onto the model.

Calibration vintages. HRS waves 5–9 define the 2000–2008 evaluation population, and all monetary amounts are restated to 2014 dollars. The SSA 2003 period table describes survival in that historical environment and eases comparison with Lockwood (2012). Later HRS waves enlarge the cells used to estimate relative health hazards, after which age-specific normalization restores aggregate survival to the 2003 table. Market evidence from the mid-1990s through 2019 supplies the reported MWR range. The calibration therefore pairs a historical evaluation population with more precise evidence on health-related mortality and annuity prices, so it does not reconstruct any one cohort's annuity market at the time.

Population sample. The wealth distribution and most filter variables come from the RAND HRS Longitudinal File 2022 (V1) (Health and Retirement Study, 2025c). Construction of the q286 ownership measure additionally uses the 2000, 2002, 2004, 2006, and 2008 RAND HRS Fat Files (Health and Retirement Study, 2021, 2017, 2019, 2025a,b). The sample pooled waves 5–9 (2000–2008) and so extends the 2006 cross-section in Lockwood (2012). It retained single, unpartnered respondents aged 65–69 who were partly retired, retired, disabled, or not in the labor force and who had valid self-reported health and respondent

weights. Because the model has no labor income, the screen also required prior-calendar-year earnings no greater than the constant-real 2014-dollar equivalent of Lockwood's $3,000 cutoff in the 2006 HRS. Missing partnership, labor-force status, earnings, health, weight, or baseline wealth failed the screen, and Appendix P lists the variable codes and price deflators. The resulting sample contains 3,377 person-wave observations. RAND total nonhousing wealth is the baseline asset measure. It excludes primary and secondary residences and their associated mortgage and home-loan debt while retaining IRAs, business and vehicle equity, other real estate, and financial assets net of other debt. Wealth is restated from each interview year to 2014 dollars before applying the $5,000 eligibility floor and the $30,000, $120,000, and $350,000 band boundaries. This measure remains broader than strictly liquid or immediately annuitizable wealth. The main analysis restricts attention to the 1,896 observations with baseline nonhousing net worth of at least $5,000. Both ownership-stock rates below are computed on that same wealth-restricted subsample, but neither enters calibration. The model maps each household's current state to a one-time predicted purchase decision. Focal evaluation is unweighted, and Appendix D.3 reports survey-weighted, person-balanced, alternative-wealth, and pension-recipient-mixture specifications.

Retirement income. Four fixed wealth bands map empirical states to retirement-income schedules. Focal Social Security is the survey-weighted mean of RAND retirement-benefit income among positive recipients in each band. Because the model has no claiming state, it assigns that amount to every state in the band as an entitlement-flow proxy rather than the respondent's observed current benefit. A sensitivity instead uses the all-respondent band mean, including zero current payments. For the pension-income proxy, the construction first converts each respondent's nonnegative prior-year employer-pension income, in 2014 dollars (Appendix P gives the variable code), to the model's age-65 nominal base by multiplying it by $(1 + \pi_P)^{a_i - 65}$, where a_i is interview age, and then takes the unconditional survey-weighted band mean, including zeros. The raw pooled mean at observed ages is kept only as a diagnostic. Because the pension-income question refers to the prior calendar year and exact payment dates are unavailable, the use of interview age is a timing convention. The data also do not identify the payment's age profile, plan type, or life contingency. These assigned schedules, rather than person-level flows, enter both the on state and the equal-PV commutation. Appendix D.3 reports the current-payment and pension-recipient-mixture sensitivities.

The unweighted focal statistic gives equal mass to each eligible person-wave state. Repeated observations of the same person therefore count as separate evaluation states, and the Shapley values describe the model over that set of states. Survey-weighted and earliest-wave

person-balanced evaluations in Appendix D.3 show how the full-model level changes under alternative weighting conventions.

Ownership measures. Two HRS ownership-stock measures give loose empirical context, and neither is a calibration or validation target. The preferred one is a life-contingent-payment indicator from the HRS annuity-income questions in waves 5–9, the q286 series. It identifies respondents reporting at least one annuity payment that they state will continue for life. In the pooled sample of 1,896 eligible person-waves, 73 report such a payment, or 3.85% (descriptive Wilson-type 95% interval $[2.85\%, 5.18\%]$). Because repeated observations make a person-wave binomial interval too narrow, the interval uses an effective sample size of 1068 under the conservative assumption of within-person correlation one (Appendix P). At the person level, 63 of 1,269 sample members (4.96%) report a qualifying payment in at least one eligible observation. The question identifies payout duration but not purchase source or date, so the reported stock may reflect earlier or nonretail decisions and cannot isolate a current retail SPIA purchase.

The conventional measure is the RAND any-annuity-income variable `r{w}iann`, which records positive private annuity income separately from pension income but does not establish that every payment is life-contingent. Pooled across waves 5–9, 6.12% of person-waves report positive annuity income (descriptive Wilson-type 95% interval $[4.83\%, 7.72\%]$, using the same effective-sample-size convention).

Lifecycle simulation diagnostic. Appendix Table O.1 reports the model-implied mean, median, and fraction of bequests above \$10K from 100,000 simulated trajectories. Initial wealth and health are drawn jointly from the weighted model-eligible HRS sample, all trajectories begin at age 65, and households receive their wealth-band Social Security flow and modeled nominal pension-income proxy but purchase no private annuity. Modeled pension-proxy income uses the age-65 base just described and then applies nominal erosion. The exhibit is descriptive. Estate values reported in HRS exit interviews do not share the model's asset definition or imputation conventions, and a calibration that targeted bequests would require a separately constructed exit sample and joint re-estimation of θ and κ.

4 Results

4.1 Focal attribution

Table 2 reports complete decompositions for two outcomes. Panel A allocates the 62.3 pp decline in purchase propensity from the equal-resource benchmark to the full model. Pricing frictions lead with 28.1 percentage points, followed by pre-existing annuitization (15.1), the medical-expenditure process plus health-dependent mortality (8.4), pessimism (6.3), and DFJ-style bequests (2.7). As shares of the modeled decline, pricing frictions account for 45%, the leading three for 83%, and bequests, ranked fifth, for 4%. Every focal average contribution is nonnegative, although some players have negative marginal effects in particular predecessor sets.

Panel B changes the outcome to the population mean purchased share of liquid wealth. The benchmark-to-full-model decline is 76.2 pp. Pricing frictions remain first at 29.1 points and pre-existing annuitization remains second at 20.2. Pessimism moves to third (11.9), the medical/mortality player moves to fourth (9.4), and bequests remain fifth while rising to 7.9 points. Inflation supports the purchased share on average, with a contribution of -3.1 points. Across both outcomes, pricing frictions remain first, pre-existing annuitization second, and bequests fifth, while the ordering and signs of the remaining channels change.

Figure 1 displays one sequential path through the nine players for intuition, and Table 2 supplies the order-averaged allocation. Table 2 also reports how each participation marginal varies across predecessor sets. The weighted 10th–90th percentiles and suppressing-order shares use the exact random-order probabilities in the Shapley formula, and the added-last column is the marginal in the full model. A wide range or a sign change shows how much a sequential increment depends on the mechanisms that entered before it. The replication file `order_sensitivity.csv` retains the median, full range, and exact positive, negative, and zero order probabilities for every player.

The pricing player combines three features of the modeled retail contract: a proportional money's-worth wedge, a fixed purchase cost, and a minimum premium. All three affect the extensive-margin opportunity, so the 28.1-point bundled value cannot be assigned to the proportional wedge alone. At a \$500 fixed cost, the complete game still puts pricing first at 22.4 points. Medical/mortality contributes 11.6, pre-existing annuitization 10.6, and bequests 1.8 (Appendix Table D.4).

Appendix L replaces the pricing bundle with three players. In that game, pre-existing annuitization ranks first at 16.2 points, followed by the MWR wedge at 14.9, the fixed cost at 9.7, and medical/mortality at 9.2. Minimum-premium restrictions contribute 1.3 points and bequests 2.5. The split confirms that both the evidence-based payout wedge and the

Table 2: Shapley Decomposition by Outcome Statistic

Panel A: Purchase propensity

Channel	Shapley (pp)	Share (%)	Order p10–p90 (pp)	Added last (pp)	Suppressing orders (%)
Pricing frictions	+28.08	45.1	+0.0 to +45.6	+42.4	71
Pre-existing annuitization (SS+pension-income proxy)	+15.08	24.2	+0.0 to +39.6	+24.1	68
Medical process and health-dependent mortality	+8.38	13.5	-8.8 to +33.8	-8.8	53
Pessimism	+6.32	10.1	+0.0 to +13.6	+9.7	63
Bequests	+2.66	4.3	+0.0 to +5.0	+5.0	54
Inflation	+1.31	2.1	+0.0 to +2.7	+2.5	46
Public-care aversion	+0.38	0.6	-9.3 to +5.4	+5.4	33
State utility	+0.07	0.1	-0.8 to +0.9	-0.6	17
Age needs	+0.00	0.0	+0.0 to +0.0	+0.0	0
Total modeled drop	+62.29	100.0	—	—	—

Panel B: Mean purchased share of liquid wealth

Channel	Shapley (pp)	Rank
Pricing frictions	+29.10	1
Pre-existing annuitization (SS+pension-income proxy)	+20.25	2
Pessimism	+11.92	3
Medical process and health-dependent mortality	+9.45	4
Bequests	+7.94	5
State utility	+0.45	6
Public-care aversion	+0.27	7
Age needs	+0.00	—
Inflation	-3.14	8
Total modeled drop	+76.24	—

Exact Shapley values over all $2^9 = 512$ subsets of the nine structural channels (SDU and PED held off). Positive values are demand-suppressing contributions; negative values are demand-raising contributions. Each panel orders channels by its signed Shapley value; exact null players receive no rank.
The order range is the weighted 10th–90th percentile of the channel's marginal contribution over uniformly random entry orders. 'Added last' is the grand-coalition marginal. 'Suppressing orders' is the exact probability that the marginal is positive under that same random-order measure. The machine-readable order-sensitivity file also reports the median, full range, and zero and negative probabilities.
Panel A endpoints are 100.0% in the equal-resource, floor-inclusive benchmark and 37.7% in the full nine-channel model. Panel B endpoints are 95.6% and 19.3% of liquid wealth.
Table A.5 adds the two exploratory behavioral channels (SDU and PED) to the nine-channel accounting.

assumed fixed cost matter. It also shows why component values and ranks must be reported under their own partition.

The medical player also bundles two linked primitives. Its off state sets medical expenditure to zero, so activation adds the calibrated mean, dispersion, and health dependence of spending. That expenditure process raises the option value of liquid wealth, while health-dependent mortality reduces the expected duration of annuity payments in states with high medical need. Separating expenditures from the health–mortality correlation changes the predecessor sets in which each can operate, and Appendix Table D.5 reports that alternative game.

The pre-existing-annuitization value depends on the equal-PV treatment. Deleting Social Security and pension-proxy flows would make the off state poorer and would combine a resource effect with an asset-form effect. Commutation instead gives each respondent liquid wealth of the same actuarial value, so the comparison isolates the liquidity, survival-contingency, and payment-path consequences of holding retirement resources in annuitized form. The result remains conditional on the assumption that the pension-income proxy is nominal and life contingent, which the HRS pension-income measure cannot verify.

The bequest value depends on translating the DFJ intensity and shifter together. DFJ's large shifter makes bequests a luxury good and avoids the spike in marginal utility near zero wealth that pairing the transformed intensity with a homothetic shifter would produce. Because the player's importance moves with curvature, the complete-game sensitivity varies risk aversion.

Two sequential paths in Appendix Tables A.6 and A.7 give different increments, because each path assigns interactions to the player that enters later. Equation (17) averages those increments over every entry order, and Appendix Tables F.1 and D.5 report interaction and partition diagnostics.

Across complete Shapley recomputations for $\gamma \in [2, 3]$, full-model propensity ranges from 30.3% to 38.5%. Lower curvature puts pessimism ahead of the medical/mortality player and can raise bequests above it. Changes to the public floor, mortality convention, fixed cost, or player partition also alter the allocation. Complete games on the three nested grids preserve the focal top-three ordering, and Section 5 reports each comparison.

Every player stays in the game regardless of sign, because efficiency requires it. In another calibration, a player could support demand on average and would then receive a negative demand-suppressing value. The age-needs player remains in the game with a zero focal parameter because the constrained calibration selects the boundary. Its zero value is a calibration outcome.

4.2 Bequests, pricing, and outcome choice

Lockwood (2012) shows that realistic loads can make a bequest motive much more consequential for annuity demand. The two-player projection from the empty coalition in Appendix Table F.1 hides that mechanism on the participation outcome. Bequests alone, pricing frictions alone, and the pair together all leave this model's equal-resource population at the 100% participation ceiling, so the finite difference is 0.0 points.

Averaging the bequest–pricing second difference over every coalition of the other seven players, as in the predecessor-averaged pair interaction of Appendix equation (F.1), gives a positive interaction of 3.4 percentage points for participation and 5.3 points for mean purchased share. The interaction emerges after other mechanisms move households away from the participation ceiling. Larger pairwise interactions exist elsewhere in the game, and the pre-existing-annuitization/pricing pair is the largest in absolute value under both outcomes. The bequest–pricing pair matters here because of its bearing on Lockwood (2012). Bequests also receive a larger mean-share allocation (7.9 points) than participation allocation (2.7 points), although they rank fifth under both outcomes. Pricing frictions therefore have the larger average contribution, while bequests remain economically active through purchased shares and interactions.

4.3 Purchase propensity across the wealth distribution

The model-eligible population is divided into fixed nonhousing-net-worth bands of \$5–30k, \$30–120k, \$120–350k, and above \$350k. Full-model purchase propensity is 0.0%, 4.9%, 76.0%, and 100.0% across those bands. By contrast, the q286 ownership stock is much smoother. The model evaluates a current retail purchase, while q286 can reflect an earlier purchase from an unidentified source. Appendix Table J.3 places the profiles side by side.

Aggregate modeled purchase propensity (37.7%) is numerically well above both the q286 stock (3.85%) and the any-annuity-income stock (6.12%). Because the HRS figures describe existing annuity payments, with purchase dates and sources the survey does not record, the difference between the model and either stock cannot be read as a residual. The distance is still economically important, because the nine modeled channels leave much more immediate purchase than either HRS stock records. A similarly stark discrepancy appears in the increase from 4.9% to 76.0%, a model-implied extensive-margin cliff that the observed q286 profile does not share.

The gap also marks the limits of the accounting, because the state space omits the option value of delaying purchase, couples and family risk sharing, housing illiquidity, risky assets, taxes, defaults and inertia, product complexity and trust, persistent residual medical costs,

and repeated purchase. Several could lower immediate purchase or flatten its wealth profile, but their effects on a larger Shapley allocation remain unknown until they are defined and solved as part of the game. The exploratory eleven-player behavioral extension (Appendix A.9) shows how a large purchase-event penalty can push the model to a boundary, but the size of that penalty is the researcher's choice, so the exercise does not estimate the missing gap.

The fixed cost makes this distributional profile more than a scaled version of the aggregate result. Low wealth can make a positive purchase infeasible or leave too little liquid wealth after paying the fixed cost. Middle-band households cross those feasibility and value thresholds rapidly as resources rise, while the top band is saturated. Purchase propensity stays at 37.7% whether the minimum premium is removed or raised to $25,000 (Appendix Table D.2), so the minimum premium by itself does not change the modeled level. Economic value continues to reflect bequests, existing annuity income, medical expenditures, and the payout schedule. An extensive-margin diagnostic distinguishes three reasons for a zero. In the $5–30k band, 35% have no feasible positive purchase and 65% reject every feasible purchase even at zero fixed cost. Corresponding fixed-cost-defeated shares in the $30–120k and $120–350k bands are 30% and 20%, and Appendix J.1 reports the remaining shares and definitions.

Separate by-band Shapley games show how the aggregate allocation varies across the distribution. Pricing frictions are weakly first in all four bands, with an exact tie with pre-existing annuitization in the $120–350k band. In the above-$350k band both endpoint propensities are 100%, so the contributions sum to zero and offset one another, and the tied top rank there is uninformative.[2] Cells with little purchase propensity remain sensitive to the fixed-cost threshold, so Appendix Tables J.1 and J.2 report feasibility, heterogeneous-cost, and grand-coalition-removal diagnostics.

The by-band values do not add up to the aggregate table without weighting. Each band has its own empty- and full-coalition endpoints, and nonlinear thresholds can make a player important in a band even when few observations purchase there. The aggregate Shapley value instead averages coalition purchase indicators over the full evaluation sample before allocating the endpoint change. Reporting both objects separates where a mechanism operates from how much it contributes to the population change.

[2]The replication package lists every band's values in `tables/csv/shapley_by_wealth.csv`.

4.4 Within-model access value

Welfare is measured by the subjective ex ante consumption-equivalent value (CEV) of access to the modeled annuity. Under DFJ bequests, population mean CEV is 0.9%, 37.7% of evaluated model-eligible person-waves have positive CEV, and 27.5% exceed 1% of consumption. Household valuation uses subjective survival, while insurer pricing and physical simulation use objective survival. Appendix Table I.1 reports the age-65 wealth–health grid, and Appendix Table J.5 reports price, indexation, and belief sensitivities.

Purchase propensity and access value answer different questions. The purchase indicator compares the best feasible positive purchase with no purchase after applying the fixed and minimum costs. CEV asks how much scaled consumption in the fixed no-access allocation would deliver the value available with the market. A household can therefore have a small access value even when a grid or transaction-cost threshold changes its discrete purchase indicator. Conversely, a positive modeled purchase does not imply that an unmodeled reform generating that opportunity would pass a social cost-benefit test.

The calculation scales consumption felicity in the fixed no-access allocation and leaves the bequeathed-wealth stock unscaled. Appendix I derives the closed form and reports an alternative convention that holds consumption in public-floor-binding periods fixed. That convention raises CEV in the diagnostic cells with positive access value and leaves the non-purchasing cells at zero. Both conventions value access within the stated preferences and transfer environment. Neither includes the fiscal cost of the floor or equilibrium responses.

4.5 Social Security benefit reduction

The policy experiment applies a permanent 22% reduction to Social Security that households anticipate from the observed purchase age onward. Pre-retirement saving and claiming, the employer-pension-income proxy, and annuity prices stay fixed. Aggregate purchase propensity changes from 37.7% to 43.2%. Table 3 reports the focal cut and the surrounding response schedule.

Purchase propensity rises monotonically over the reported 0–50% cut range, reaching 49.0%, but complete elimination lowers it to 41.7%. Because the schedule is discrete, it does not locate the turning point between 50% and full elimination. This reversal is consistent with opposing longevity-insurance and liquid-wealth-value margins, although the exercise does not separately identify their contributions.

The calculation isolates an uncompensated resource reduction at the full model and is not the inverse of the equal-PV pre-existing-annuitization player. It holds the evaluated initial wealth state, pension-income proxy, legislation, claiming, cohort transitions, insurer

Table 3: Private Annuity Demand Response to Social Security Benefit Reductions

SS Benefit Cut	Propensity (%)	Mean α	Δ (pp)
0% (baseline)	37.7	0.193	—
10%	40.1	0.216	+2.4
15%	41.6	0.229	+3.9
22% (anticipated cut)	43.2	0.247	+5.5
30%	45.0	0.268	+7.3
40%	46.9	0.293	+9.2
50%	49.0	0.319	+11.3
100% (elimination)	41.7	0.325	+4.0
HRS q286 ownership stock	3.85		
HRS any-annuity-income proxy	6.12		

Structural model (rational, preference, and public-care aversion channels; behavioral channels off). The permanent anticipated cut changes Social Security only; the modeled nominal pension-income proxy is held fixed. Baseline SS = [\$13K, \$14K, \$15K, \$15K], pension proxy = [\$7K, \$10K, \$13K, \$13K]. The 22% permanent cut is motivated by projected financing pressure but is not a forecast of the timing, transition rule, legislation, or general-equilibrium incidence. 100% cut eliminates Social Security entirely (the pension proxy remains). The two HRS rows are distinct, nonnested stock comparators, not an interval.

pricing, and labor supply fixed, while the lifecycle consumption–saving policy is re-solved for each benefit schedule. The exercise therefore traces partial-equilibrium responses to lower benefits and does not forecast the transition after the retirement trust fund runs out (Board of Trustees, 2026). Appendix Table J.1 reports how the wealth-band response changes under alternative extensive-margin treatments.

Two forces can oppose one another in this experiment. Lower Social Security reduces pre-existing longevity insurance and can raise the marginal value of a private annuity, but it also raises the shadow value of retaining liquid wealth once the public income flow is smaller. Initial liquid wealth, and therefore mechanical minimum-premium feasibility at purchase, is unchanged. At the 22% cut, the band responses are 0.0%, 8.5 pp, 14.3 pp, and 0.0 pp from the lowest to highest wealth band. The lowest band does not respond on the focal extensive margin, and the highest is already saturated. Because the policy changes the amount of resources and does not toggle the equal-PV player, these band figures are net responses of the full model.

Appendix J reports separate price, indexation, belief, and lower-price product sensitivities. Employer-pension receipt is used there only as an illustrative mixture weight and does not measure group-annuity access or imply that employer markets would replace a public-benefit cut.

5 Robustness

Appendix Table D.1 reports the full sensitivity analysis.

5.1 Public consumption floor

The public consumption floor is fixed outside the focal game, so the attribution is conditional on that institution. On the common extended H2 grid, the focal floor yields 37.1% full-model purchase propensity. Pricing frictions rank first at 27.9 percentage points, pre-existing annuitization second at 15.9 points, and the medical-expenditure/mortality player third at 8.5 points. With a near-zero floor, propensity falls to 8.8%, while pre-existing annuitization ranks first at 49.4 points, the medical-expenditure/mortality player second at 21.6 points, and pricing third at 14.7 points. The same three channels remain the leading set in this comparison, but their order and magnitudes depend heavily on the public safety net (Appendix Table D.3).

5.2 Risk aversion

Appendix Table D.1 reports the level sensitivity to risk aversion, and Table 4 reports the joint risk-aversion–inflation cells. The note to Table 4 lists the entry-age row of the focal age-varying hazard schedule, which every cell retains. At the focal $\gamma = 2.5$, the full model predicts 37.7%. Purchase propensity is 1.6% at $\gamma = 1.5$, 30.3% at $\gamma = 2.0$, 38.0% at $\gamma = 2.6$, and 38.5% at $\gamma = 3.0$.

Table 4: Predicted Purchase Propensity (%) by Risk Aversion and Inflation Rate

	$\pi = 1\%$	$\pi = 2\%$	$\pi = 3\%$
$\gamma = 2.4$	38.6	37.0	35.3
$\gamma = 2.5$	39.1	37.7	35.9
$\gamma = 2.6$	39.5	38.0	36.6
$\gamma = 3.0$	40.0	38.5	36.7

Baseline: $\gamma = 2.5$, $\pi = 2\%$, DFJ bequests, MWR = 0.87.
Every cell retains the focal age-varying hazard schedule;
its entry-age multipliers are [0.49, 1.0, 3.29].

The complete games at $\gamma = 1.5$ and at the tested values of γ between 2 and 3 recompute the contributions as well as the endpoints. At $\gamma = 2$, pricing frictions remain first, but pessimism moves to second and pre-existing annuitization to third. Bequests contribute 6.4 percentage points and rank fourth, ahead of medical/mortality. At the lower-curvature sensitivity $\gamma = 1.5$, bequests rise to 14.2 points and rank third, behind pricing frictions and pessimism. Bequests therefore need not remain fifth even though bundled pricing frictions exceed them in every reported curvature game.

Changing γ affects several margins simultaneously: consumption smoothing, the value of longevity insurance, the cost of losing liquid wealth in adverse medical states, and the DFJ bequest intensity through the portability mapping. Through that mapping, the stripped-down terminal allocation wedge remains fixed. Because these margins need not move demand in the same direction, there is no a priori sign or monotonicity restriction on the resulting purchase propensity.

In the nominal-annuity calibration, inflation erodes the real value of later payments but raises the initial nominal payout, because the insurer prices with the nominal discount rate. At the baseline $\gamma = 2.5$, predicted purchase propensity is 39.1% at $\pi = 1\%$ and 35.9% at $\pi = 3\%$. This joint path changes both the initial payout and its subsequent erosion, so it does not identify either effect in isolation.

5.3 Health-mortality correlation

The hazard multipliers govern the strength of the Reichling–Smetters mechanism. Focally, the age-varying specification produces 37.7% and reflects the narrowing of hazard ratios across health states at older ages (the Poor-health multiplier falls from 3.29 at ages 65–74 to 1.82 at 85+). By comparison, the constant entry-age row produces 36.8%, the steeper researcher-specified mapping $[0.45, 1.0, 3.5]$ produces 36.4%, the HRS self-reported-health mapping $[0.57, 1.0, 2.7]$ produces 37.0%, and the conservative specification $[0.60, 1.0, 2.0]$ produces 37.8%. Appendix Figure O.2 reports the level comparison, and Appendix N reports a complete-game recomputation under the constant entry-age row.

Appendix M also restores the earlier male-table, unnormalized-hazard convention used in this literature. In that complete game, purchase propensity falls to 21.9%. Pricing frictions remain first at 26.5 points, while medical/mortality rises to second at 23.9 and pre-existing annuitization moves to third at 17.2. The comparison changes both the population table and hazard normalization, so it establishes sensitivity to the mortality convention without isolating a single mechanism behind the rotation.

5.4 Pricing and survival beliefs

Varying MWR over a sensitivity range that extends above the cited population-mortality values moves purchase propensity from 27.8% at 0.82 to 49.9% at 0.95, with the fixed cost and minimum premium retained at every price point. Appendix Table D.1 reports each one.

Predicted purchase propensity is 26.1% at $\psi = 0.96$, 37.7% at the focal $\psi = 0.981$, and 47.4% under objective beliefs. At the strong-pessimism endpoint, pricing frictions remain first, pessimism ranks second, and pre-existing annuitization ranks third. Bequests move to fourth and medical/mortality to fifth. Appendix A.5 reports that complete-game recomputation.

The belief experiment changes household valuation but not physical mortality or insurer pricing. It therefore opens a deliberate wedge between the survival probabilities used in the Bellman equation and those used to price and simulate the contract. The endpoint game tests whether the allocation changes when that wedge is strengthened.

5.5 Additional robustness checks

Joint parameter stress design. Across the 1,000 fixed-seed design points, median purchase propensity is 26.2%, the central 90% spans 2.7%–43.0%, and the full range is 0.3%–50.6%. Each point replaces the focal age-varying health-hazard schedule with constant

Good–Fair–Poor multipliers whose Poor-health value is drawn uniformly on $[2.0, 3.5]$, so the design does not nest the focal mortality schedule. The spread summarizes researcher-specified parameter draws and has no confidence-interval or posterior interpretation (Appendix H).

Grid convergence. Refinement near the fixed-cost extensive margin changes participation levels. Predicted purchase propensity is 37.7% at the focal piecewise 98×39 grid and 37.7% under nested tail refinement (114×47). Full nested refinement (195×77) gives 37.1%, while a conventional single-power 170×50 grid gives 37.7% (Appendix Table B.1). The complete H0, H1, and H2 games preserve the focal top-three ordering of pricing frictions, pre-existing annuitization, and medical/mortality despite movement in levels and contribution magnitudes (Appendix Table B.2).

Quadrature convergence. Appendix Table B.1 (Panel B) varies the state-split Gauss–Legendre order per probability segment over $\{8, 12, 16, 24, 32, 48\}$. At the focal grid, purchase propensity is 37.7% at order 32 and 37.7% at order 48, a difference of 0.0 pp. Panel C reports the fully refined H2 grid at orders 48 and 64, with order 64 serving as the combined grid-and-quadrature reference. Order 32 remains the focal production rule, although the extensive-margin indicator can still change when small expectation errors move a household across $\alpha^* > 0$.

Annuitization grid. Holding the wealth grid and quadrature at focal resolution and varying n_α over $\{51, 101, 201, 401\}$ leaves purchase propensity within a 0.053 pp band. Appendix Table B.1, Panel D, reports the four values before display rounding.

Outcome statistic and evaluation age. Table 2, Panel B, reports the complete mean-purchased-share vector. Pricing frictions and pre-existing annuitization remain first and second. Pessimism moves to third, medical/mortality to fourth, and inflation becomes demand-supporting. Forcing every observation's evaluation age to 65 changes purchase propensity to 33.8%, with pricing frictions still first at 29.4 points. Pre-existing annuitization ranks second, pessimism third, medical/mortality fourth, and bequests fifth at 3.5 points. Both diagnostics replace medical/mortality with pessimism in the leading three, so outcome choice and evaluation age affect the focal ranking.

Bequest specification. Under the nine-channel structural baseline, the CPI-restated DFJ luxury-good specification ($\theta = 158.68$, $\kappa = \$396{,}496$) produces 37.7% purchase propensity; removing bequests changes this to 42.7%. Keeping the source shifter at its unrestated

1998-dollar value instead produces 31.0%. That run is included only as a units sensitivity. Section 2 discusses why θ and κ must be interpreted jointly.

Numerical accuracy. Wealth–annuity and annuitization-grid refinement, the H0 order-32-to-48 and H2 order-48-to-64 quadrature checks, limiting-case tests, and the endogenous next-period-wealth boundary audit assess accuracy (Appendices B and E). A citable run fails unless the diagnostic rosters are complete, independently generated focal cells agree up to serialization precision, every selected policy transition remains inside the wealth grid up to a scale-aware numerical tolerance, and the participation and mean-share convergence budgets in Appendix B hold. The production boundary certificate evaluates 68,107,988 transitions and records zero material overshoots. No single Euler-residual scalar is reported because lower and upper consumption corners require Karush–Kuhn–Tucker (KKT) inequalities and the continuation-value approximation is kinked and piecewise linear, so a smooth finite-difference residual has no uniform interpretation across states.

Exact coalition enumeration removes sampling error from the Shapley formula, but it cannot correct error in the coalition values. The H0 order-48 and H2 order-64 quadrature audits, grid refinements, exact annuitization-boundary candidates, and fail-closed empirical-state coverage address those values directly. Residual movement near a discrete participation threshold is reported as a sensitivity.

6 Discussion

Exact Shapley accounting removes the investigator's entry order from a multi-channel life-cycle decomposition. In the focal participation game, pricing frictions, pre-existing annuitization, and medical/mortality rank first through third and together account for 83% of the modeled decline. Calibrated bequests rank fifth and contribute no more than pricing frictions in any reported recomputation that keeps the three pricing features together. Their focal contribution is small but grows under lower risk aversion, and bequests interact positively with pricing frictions.

The allocation depends on the definition of a player. Bundling the medical process with health-dependent mortality assigns their interaction to one named mechanism. Splitting them distributes that interaction between two average marginal contributions. The same issue applies to the three pricing terms and to Social Security versus the pension-income proxy. Alternative-partition games show this dependence and clarify the economic meaning of each label. In the pricing split, pre-existing annuitization leads the MWR wedge, and the assumed fixed cost remains quantitatively important.

Lowering the public consumption floor puts pre-existing annuitization first and pricing frictions third. Lower risk aversion and stronger pessimism bring survival beliefs into the leading three, while alternative mortality, fixed costs, and player partitions produce further changes. The three complete nested-grid games preserve the focal top-three order despite changes in levels and contribution magnitudes.

Because switching off pre-existing annuitization converts retirement income into liquid wealth of equal actuarial value, that player measures the effect of the form in which resources are held. The 22% policy experiment instead removes Social Security resources at the full model with prices, claiming, legislation, and equilibrium adjustment held fixed, so its result is a partial-equilibrium comparative static.

Eligible observations are single respondents aged 65–69 who meet the labor-force, earnings, health, weight, and $5,000 nonhousing-wealth screens. Each observation maps to one irreversible retail purchase at the observed age. The two HRS measures are current ownership and income stocks with no purchase date and incomplete product identification, so they provide scale and distributional context for the model prediction.

The difference between 37.7% modeled purchase propensity and the HRS stocks remains a substantive limitation even though the quantities are not matched moments. Omitted mechanisms include deferral, marriage and family insurance, housing choice, equity, taxes, insurer credit risk, defaults, product complexity, trust, and repeated purchase. Family risk sharing can substitute for annuity purchases (Kotlikoff and Spivak, 1981), housing can alter annuity demand (Pashchenko, 2013), and a risky asset can change the opportunity cost of annuitization (Inkmann et al., 2011). The current results cannot assign the level gap to those omitted channels or predict their positions in a larger game. RAND nonhousing net worth also includes assets with limited immediate liquidity. Appendix D.3 varies the asset definition and weighting in full-model level checks, without recomputing the coalition game under those alternative state constructions. The reported complete games do not jointly vary the fixed cost, money's-worth ratio, and bequest parameters, so they do not locate where the pricing-over-bequests ordering would reverse.

Health is Markov and residual medical shocks are iid conditional on health. The annual health matrices are principal-square-root transformations of biennial transitions observed among respondents alive and reporting health at both endpoints, then combined separately with current-health mortality. This reduced-form annualization is not a jointly estimated health–death process, and the game does not propagate attrition or sampling uncertainty. Persistent medical or long-term-care cost states could raise the value of liquidity and change the medical-channel allocation, but signing that change and its effect on the rankings would require a model with that additional state. The public-care player represents the stated

interaction of medical risk, Poor health, and a generic floor because Medicaid participation and long-term-care setting, insurance, and utilization are absent. Survival beliefs use a constant proportional wedge despite the documented age pattern in belief errors (O'Dea and Sturrock, 2023). Public-care aversion, state-dependent utility, and the two exploratory behavioral channels rest on assumed mappings.

Player definitions, numerical tolerances, and empirical screens were fixed before the final production run, and the full game was enumerated, so the accounting is exact given those choices. Every complete-game sensitivity exercise recomputes both endpoints. The open questions are economic: which states and mechanisms belong in the lifecycle problem, how their parameters should be identified, and what observed purchase measure could match the modeled decision in timing and product scope.

7 Conclusion

Sequential decompositions credit each interaction to whichever mechanism enters later, so their answer depends on an order the investigator chooses. Solving all 2^9 coalitions of a nine-player lifecycle game removes that choice. Entered last, medical expenditures with health-dependent mortality raise purchase propensity from 29.0% to 37.7%. Averaged over every order, the same pair lowers it by 8.4 points. At the focal calibration, purchase propensity falls from 100.0% in the equal-resource benchmark to 37.7% in the full model. Pricing frictions, pre-existing annuitization, and medical/mortality rank first through third and together account for 83% of the modeled decline. The first two keep their positions under the mean-purchased-share outcome. Calibrated bequests rank fifth under both outcomes and contribute no more than pricing frictions in any reported recomputation that keeps the three pricing features together.

Recomputing the full decomposition under alternative preferences, institutions, and player definitions changes both the ordering and the membership of the leading set, and the rankings depend on which mechanisms the model includes. The three nested-grid games preserve the focal top-three participation ordering. Because fair commutation holds the amount of retirement resources fixed, the pre-existing-annuitization player measures only the effect of their form. The separate 22% Social Security experiment instead cuts benefits without compensation and, with prices and initial wealth held fixed, raises purchase propensity by 5.5 percentage points in partial equilibrium. Purchase propensity in the full model, 37.7%, remains far above the HRS ownership stocks. In the focal calibration, then, pricing frictions and the retirement income households already hold do more than bequests to lower modeled demand for private annuities.

Data Availability

A replication package contains the Julia code, the redistributable aggregate calibration inputs, the generated results, and reconstruction instructions, with a checksum for every file. It is available to the editors and referees on request and will be deposited in a public repository at acceptance. Health transitions, population evaluation, and ownership-measure files are constructed from RAND HRS and RAND HRS Fat File inputs available to registered HRS users at https://hrsdata.isr.umich.edu. The author does not redistribute the source records or three person-level intermediate extracts through the replication package under the HRS conditions. Registered users can regenerate them from their own downloads with the documented local reconstruction pipeline, which records a checksum of every input file so users can confirm they reproduced the same extracts. All remaining parameters are tied to cited sources or labeled as calibration choices.

Declaration of competing interests

The author owns and is employed by a financial advisory firm (Conscious Capital, LLC) that is primarily compensated through asset-under-management fees, and is licensed to sell insurance and annuities. In addition, the author is employed by Income Lab and has an ownership interest in a software company that provides retirement spending projections, including Social Security and annuity analysis.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

Generative AI tools were used to assist with code development, manuscript preparation, and editing. All analysis, interpretation, and conclusions are the sole responsibility of the author, who reviewed and edited the content and takes full responsibility for the published article.

Acknowledgments

The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant numbers NIA U01AG009740 and NIA R01AG073289) and is conducted by the University of Michigan. Separate citations cover the RAND HRS Longitudinal File 2022 (V1) and the five RAND HRS Fat Files used here (Health and Retirement Study, 2025c; RAND Center for the Study of Aging, 2025; Health and Retirement Study, 2021, 2017, 2019, 2025a,b). This analysis uses Early Release data from the Health and Retirement Study, RAND HRS Longitudinal File 2022 (V1), sponsored by the National Institute on Aging (grant number NIA U01AG009740) and conducted by the University of Michigan. These data have not been cleaned and may contain errors that will be corrected in the Final Public Release version of the dataset.

References

Aguiar, M., Hurst, E., 2013. Deconstructing life cycle expenditure. Journal of Political Economy 121, 437–492. doi: 10.1086/670740.

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M.D., Tonetti, C., 2020. Long-term-care utility and late-in-life saving. Journal of Political Economy 128, 2375–2451. doi: 10.1086/706686.

Board of Trustees, 2026. The 2026 annual report of the board of trustees of the federal old-age and survivors insurance and federal disability insurance trust funds. U.S. Social Security Administration, Washington, DC. URL: https://www.ssa.gov/OACT/TR/2026/tr2026.pdf. accessed July 18, 2026.

Davidoff, T., Brown, J.R., Diamond, P.A., 2005. Annuities and individual welfare. American Economic Review 95, 1573–1590. doi: 10.1257/000282805775014281.

De Nardi, M., 2004. Wealth inequality and intergenerational links. Review of Economic Studies 71, 743–768. doi: 10.1111/j.1467-937X.2004.00302.x.

De Nardi, M., French, E., Jones, J.B., 2010. Why do the elderly save? The role of medical expenses. Journal of Political Economy 118, 39–75. doi: 10.1086/651674.

Finkelstein, A., Luttmer, E.F.P., Notowidigdo, M.J., 2013. What good is wealth without health? The effect of health on the marginal utility of consumption. Journal of the European Economic Association 11, 221–258. doi: 10.1111/j.1542-4774.2012.01101.x.

Health and Retirement Study, 2017. 2002 RAND HRS Fat File (V2C) public use dataset. Developed by the RAND Center for the Study of Aging under subcontract from the Health and Retirement Study and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740) and the Social Security Administration. Ann Arbor, MI. URL: https://hrsdata.isr.umich.edu/data-products/2002-rand-hrs-fat-file.

Health and Retirement Study, 2019. 2004 RAND HRS Fat File (V1C) public use dataset. Developed by the RAND Center for the Study of Aging under subcontract from the Health and Retirement Study and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740) and the Social Security Administration. Ann Arbor, MI. URL: https://hrsdata.isr.umich.edu/data-products/2004-rand-hrs-fat-file.

Health and Retirement Study, 2021. 2000 RAND HRS Fat File (V1D) public use dataset. Developed by the RAND Center for the Study of Aging under subcontract from the Health and Retirement Study and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740) and the Social Security Administration. Ann Arbor, MI. URL: https://hrsdata.isr.umich.edu/data-products/2000-rand-hrs-fat-file.

Health and Retirement Study, 2025a. 2006 RAND HRS Fat File (V4B) public use dataset. Developed by the RAND Center for the Study of Aging under subcontract from the Health and Retirement Study and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740) and the Social Security Administration. Ann Arbor, MI. URL: https://hrsdata.isr.umich.edu/data-products/2006-rand-hrs-fat-file.

Health and Retirement Study, 2025b. 2008 RAND HRS Fat File (V3B) public use dataset. Developed by the RAND Center for the Study of Aging under subcontract from the Health and Retirement Study and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740) and the Social Security Administration. Ann Arbor, MI. URL: https://hrsdata.isr.umich.edu/data-products/2008-rand-hrs-fat-file.

Health and Retirement Study, 2025c. RAND HRS Longitudinal File 2022 (V1) public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant numbers NIA U01AG009740

and NIA R01AG073289). Ann Arbor, MI. URL: https://hrsdata.isr.umich.edu/data-products/rand-hrs-longitudinal-file-2022-v1.

Heimer, R.Z., Myrseth, K.O.R., Schoenle, R.S., 2019. YOLO: Mortality beliefs and household finance puzzles. The Journal of Finance 74, 2957–2996. doi: 10.1111/jofi.12828.

Hershfield, H.E., Shu, S.B., Brown, J.R., Hurwitz, A., Milevsky, M.A., Mitchell, O.S., Toland, T., 2026. The Annuity Puzzle Revisited: Barriers, Behavior, and Policy Paths to Lifetime Income. Working Paper 35145. National Bureau of Economic Research. doi: 10.3386/w35145.

Inkmann, J., Lopes, P., Michaelides, A., 2011. How deep is the annuity market participation puzzle? The Review of Financial Studies 24, 279–319. doi: 10.1093/rfs/hhq080.

Jones, J.B., De Nardi, M., French, E., McGee, R., Kirschner, J., 2018. The Lifetime Medical Spending of Retirees. Working Paper 24599. National Bureau of Economic Research. doi: 10.3386/w24599.

Kotlikoff, L.J., Spivak, A., 1981. The family as an incomplete annuities market. Journal of Political Economy 89, 372–391. doi: 10.1086/260970.

Lockwood, L.M., 2012. Bequest motives and the annuity puzzle. Review of Economic Dynamics 15, 226–243. doi: 10.1016/j.red.2011.03.001.

Mitchell, O.S., Poterba, J.M., Warshawsky, M.J., Brown, J.R., 1999. New evidence on the money's worth of individual annuities. American Economic Review 89, 1299–1318. doi: 10.1257/aer.89.5.1299.

O'Dea, C., Sturrock, D., 2023. Survival pessimism and the demand for annuities. The Review of Economics and Statistics 105, 442–457. doi: 10.1162/rest_a_01048.

Pashchenko, S., 2013. Accounting for non-annuitization. Journal of Public Economics 98, 53–67. doi: 10.1016/j.jpubeco.2012.11.005.

Peijnenburg, K., Nijman, T., Werker, B.J.M., 2016. The annuity puzzle remains a puzzle. Journal of Economic Dynamics and Control 70, 18–35. doi: 10.1016/j.jedc.2016.05.023.

RAND Center for the Study of Aging, 2025. RAND HRS Longitudinal File 2022 (V1). Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa Monica, CA.

Reichling, F., Smetters, K., 2015. Optimal annuitization with stochastic mortality and correlated medical costs. American Economic Review 105, 3273–3320. doi: 10.1257/aer.20131584.

Sastre, M., Trannoy, A., 2002. Shapley inequality decomposition by factor components: Some methodological issues. Journal of Economics 77, 51–89.

Shapley, L.S., 1953. A value for n-person games, in: Kuhn, H.W., Tucker, A.W. (Eds.), Contributions to the Theory of Games, Volume II. Princeton University Press, Princeton, pp. 307–317.

Shorrocks, A.F., 2013. Decomposition procedures for distributional analysis: A unified framework based on the Shapley value. Journal of Economic Inequality 11, 99–126.

Social Security Administration, 2007. Period life table, 2003. Annual Statistical Supplement, Table 4.C6. URL: https://www.ssa.gov/policy/docs/statcomps/supplement/2007/4c.html. accessed July 18, 2026.

U.S. Bureau of Labor Statistics, 2017. Historical consumer price index for all urban consumers (CPI-U): U.S. city average, all items. CPI Detailed Report, Table 24. URL: https://www.bls.gov/cpi/tables/historical-cpi-u-201707.pdf. accessed July 18, 2026.

Wettstein, G., Munnell, A.H., Hou, W., Gok, N., 2021. The Value of Annuities. Working Paper CRR WP 2021-5. Center for Retirement Research at Boston College. URL: https://crr.bc.edu/wp-content/uploads/2021/02/wp_2021-5.pdf.

Yaari, M.E., 1965. Uncertain lifetime, life insurance, and the theory of the consumer. The Review of Economic Studies 32, 137–150. doi: 10.2307/2296058.

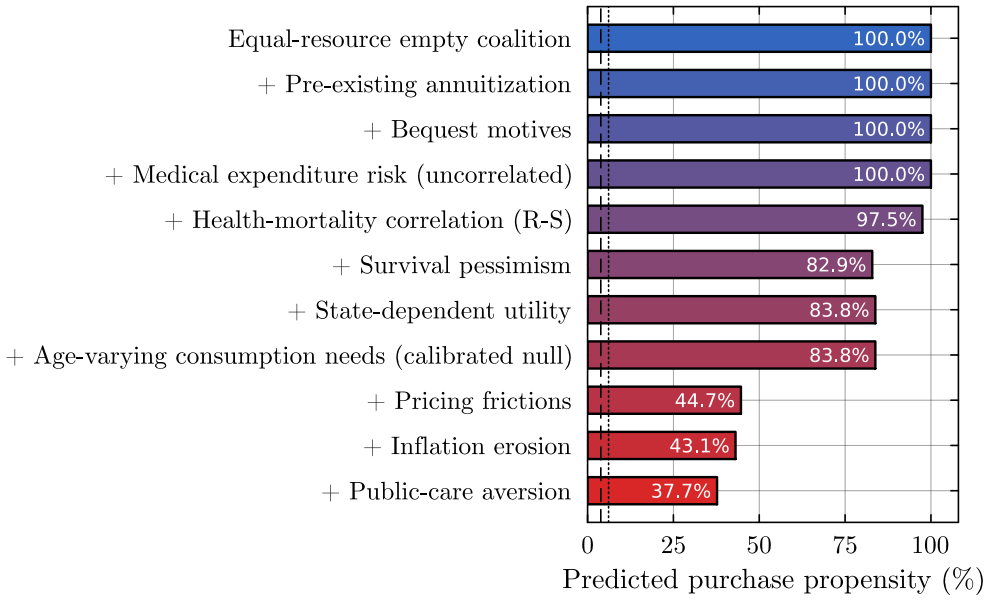


Figure 1: Sequential decomposition of predicted purchase propensity (nine-channel structural model). Each bar shows the prediction after adding the labeled channel to all previous channels; for visual transparency, the bundled medical/mortality player appears as successive medical-risk and health–mortality-correlation steps. The final bar is the nine-channel structural baseline. This single ordering is shown for intuition only—the order-independent attribution is the Shapley decomposition (Section 4). The dashed and dotted lines mark the untargeted HRS q286 life-contingent-payment stock (3.85%) and any-annuity-income proxy (6.12%).

Online Appendix for "Accounting for Interacting Explanations of Annuity Demand"

Derek Tharp

A Model Details

Table A.1 collects the focal parameters, their sources, and the calibration choices. Quantities labeled as mappings or choices were not estimated in this paper or in the cited studies.

A.1 Health transition matrices

Health transitions are calibrated to the RAND HRS Longitudinal File 2022 (V1; consecutive-wave pairs from waves 5–12, 2000–2014, with two-year rates annualized via matrix square root) (Health and Retirement Study, 2025), with the five self-reported health categories mapped to three states: Good (Excellent/Very Good), Fair (Good), and Poor (Fair/Poor). The historical window is closer to the 2000–08 evaluation population than later HRS waves are, and it excludes pandemic-period transitions. Estimates pool unweighted complete-case person-wave transitions among all HRS respondents aged 65 or older who respond and report valid health at both endpoints, without the focal marital-status, partnership, labor-force, earnings, or wealth screens. Respondents may contribute more than one interval. Transition probabilities are age-band dependent: distinct 3×3 transition matrices are estimated for six age bands (65–69, 70–74, 75–79, 80–84, 85–89, 90+) and applied piecewise within each band. In the estimated matrices, entry into Poor health accelerates around ages 75–85. The six-band specification keeps that pattern, which a two-anchor linear interpolation would smooth over. Each annual matrix must be stochastic and must square back to its two-year matrix within the builder's numerical tolerance. The three-state structure follows De Nardi et al. (2010), and the transition matrices are re-estimated here.

Write $a_t = 64 + t$ and $Q_t(h, h') = \Pr(H_{t+1} = h' \mid H_t = h, a_t)$ for the row-h, column-h' element of the one-year transition matrix used in model period t.

The two-year health matrices are survivor- and response-conditional. Each is annualized with its principal matrix square root. The builder accepts that root only if it is real,

Table A.1: Model Parameters

Parameter	Symbol	Value	Source or status
Preferences			
Risk aversion	γ	2.5	Calibration choice; Lockwood (2012) uses 2
Discount factor	β	0.97	Standard calibration
Bequest intensity	θ	158.68	Portability mapping; derivation in Section C.1
Bequest shifter	κ	\$396,496	CPI-restated from De Nardi et al. (2010)
Demographics and return			
Entry age		65	Retirement evaluation age
Maximum age		110	Lockwood (2012)
Risk-free rate	r	0.02	Real
Health and mortality			
Health states		3	Good, Fair, Poor
Hazard multipliers	$\mu(H,t)$	See text	Estimated from Health and Retirement Study (2025)
Medical-integration order		32 per probability segment	State-split probability-domain Gauss–Legendre
Survival-belief scale	ψ	0.981; range $[0.96, 1]$	Mapping motivated by O'Dea and Sturrock (2023)
Medical expenditures and retirement income			
Log mean at 65 (Fair)	μ_m	6.979338	Solved to match the projected age-70/100 two-year mean targets
Annual log-mean growth		0.0596923	Solved to match the projected age-70/100 two-year mean targets
Log standard deviation	σ_m	1.4	Calibration choice; upper-tail fit reported as an implied diagnostic
Health log-mean / log-SD shifts	$\delta_\mu(H)$	$[-0.5, 0, 0.7]$	
	$\delta_\sigma(H)$	$[-0.2, 0, 0.2]$	Assumed mapping; not estimated
Consumption floor	$\underline{c}$	\$8,496	Lockwood (2012); year-2000 SSI, CPI-restated
Pre-existing income by band	Y^{pre}	\$20,089; \$23,826; \$27,685; \$27,804	Assigned band means from Health and Retirement Study (2025), waves 5–9; conventions in main text
Annuity market			
Money's worth ratio	MWR	0.87	Mitchell et al. (1999); Wettstein et al. (2021)
Fixed purchase cost	F	\$2,349	Calibration: Lockwood (2012)'s illustrative \$2,000 experiment, treated as 2006 dollars and CPI-restated to 2014 (U.S. Bureau of Labor Statistics, 2017)
Minimum premium		\$10,000	Calibration choice; sweep reported
Inflation rate	π	0.02	Nominal-payment calibration
Preference and structural channels			
Age-needs decline	δ_c	0.0000	Constrained residual boundary check against an external expenditure-profile scale from Aguiar and Hurst (2013)
Health-utility weights	$\varphi(H)$	$[1.00, 0.92, 0.82]$	Assumed mapping; Section A.7
Public-care aversion	χ_{LTC}	0.49	Reduced-form proxy motivated by Ameriks et al. (2020); no observed care-program state
Exploratory behavioral channels			
Portfolio-consumption weight	λ_W	0.625	Mapping informed by Blanchett and Finke (2024, 2025)
Purchase-event penalty	ψ_{purchase}	0.05	Reduced form motivated by framing evidence

Notes: Dollar quantities are in 2014 dollars. The final two parameters enter only the exploratory eleven-player extension. Employer-pension income reported in the HRS supplies the pension-income proxy, and the model assumes that it is nominal and life contingent.

nonnegative, row stochastic, and recomposes the unrounded two-year matrix within a 10^{-9} maximum-cell tolerance. It neither clips entries nor projects an inadmissible root onto the stochastic simplex. Machine-readable probabilities retain full floating-point precision, while displayed values are rounded. The health root therefore recomposes its two-year input, although the principal square root is only a computational annualization and does not identify a latent annual process. Annualization precedes the interleaving of those one-year health transitions with separately estimated current-health mortality. This reduced-form separability only approximates a jointly estimated health-and-death process, because state-dependent selection into survival means that the modeled health distribution among two-year survivors need not reproduce every cell of the original survivor-conditional matrix. That discrepancy is not separately calibrated or used as identifying variation. Attrition and sampling uncertainty in the transition estimates are not propagated through the demand game.

The transition matrices exhibit two standard features. First, health deterioration is more likely than improvement: transition rates from Good to Fair exceed those from Fair to Good in every age band from 70–74 onward (in the 65–69 band the two rates are nearly equal, 0.144 versus 0.147). Second, deterioration accelerates with age: the annual probability of remaining in Good health falls from 0.84 at ages 65–69 to 0.67 at ages 90+.

A.2 Medical expenditure calibration

Gross medical needs, scaled to combined out-of-pocket-plus-Medicaid benchmarks and realized before the model's public-floor transfer, follow:

$$\ln m_t = \mu_m^{\text{base}} + g \cdot (a_t - 65) + \delta_\mu(H_t) + \left[\sigma_m^{\text{base}} + \delta_\sigma(H_t)\right] \cdot \varepsilon_t, \quad \varepsilon_t \sim N(0,1),$$

where a_t is age in period t, $\mu_m^{\text{base}} = 6.979338$ is the mean of log need at age 65 for Fair health, $g = 0.0596923$ is its annual growth, $\delta_\mu(H)$ shifts that mean by health state $[-0.5, 0, 0.7]$ for Good, Fair, and Poor respectively, $\sigma_m^{\text{base}} = 1.4$ is the base log standard deviation, and $\delta_\sigma(H)$ shifts the log standard deviation by health $[-0.2, 0, 0.2]$. Both shift vectors and the base dispersion are assumed values that Jones et al. (2018) do not report.

The source's Figure 1a reports annualized expenditures over two-year intervals for surviving households, combining initially single households and couples. A couple that loses one spouse remains a surviving household. By contrast, the model has one person and iid annual residual shocks conditional on Markov health. The model's projection is therefore the distribution of $(m_a + m_{a+1})/2$ conditional on the model individual surviving from age a through age $a + 2$, using objective health-conditioned survival and the intervening health transition. Two log-mean coefficients solve exactly for the source means at intervals beginning at ages

70 and 100 under that operator. This mapping respects the source's measurement interval and survivor conditioning, but it neither reproduces the source household unit nor imports its persistent medical-spending process.

Table A.2 reports the resulting projected moments against the source benchmarks.

Table A.2: Medical Expenditure Moments: Model vs. Target

Moment	Model	Source benchmark (Jones et al., 2018)
Projected two-year mean, interval beginning age 70	\$5,100	\$5,100
Projected two-year 95th pctile, interval beginning age 70	\$17,943	\$13,400
Projected two-year mean, interval beginning age 100	\$29,700	\$29,700
Projected two-year 95th pctile, interval beginning age 100	\$103,509	\$111,200

The log-mean coefficients match the two projected means exactly. With dispersion and health shifts held at their assumed values, the two upper-tail rows are implied diagnostics, and the table reports both misses.

A.3 Survival calibration

Baseline survival probabilities $s_{\text{base}}(t)$ use a sex-blended SSA 2003 period life table: the survivorship mixture of the male and female SSA 2003 tables at the model-eligible sample's female share (70.8%), $S_{\text{base}}(a) = w_f S_f(a) + (1 - w_f) S_m(a)$, with one-year conditional survivals derived from the mixed survivorship (Social Security Administration, 2007). The prior convention—the male table from the Lockwood (2012) replication code (BAP_wtp.m)—is retained as the mortality-convention robustness of Appendix M.

Health-conditional survival is anchored to this aggregate table. The raw proportional-hazard form is $s_{\text{base}}(t)^{\mu(H,t)}$, but the multipliers are relative hazards rather than table-anchored levels, so their population-weighted average departs from $s_{\text{base}}(t)$. A common age-specific factor ζ_t rescales every multiplier so the health mixture reproduces the table exactly:

$$s(t, H) = s_{\text{base}}(t)^{\zeta_t \, \mu(H,t)}, \qquad \zeta_t \text{ solves } \sum_H \omega_H(t) \, s_{\text{base}}(t)^{\zeta_t \, \mu(H,t)} = s_{\text{base}}(t),$$

where $\omega_H(t)$ is the health distribution among survivors in period t (age $64 + t$), initialized from the model-eligible sample's pooled ages 65–69 counts (737 Good, 649 Fair, and 510

Poor of 1,896) and evolved forward under the health transition matrices weighted by the normalized survivals. For $0 < s_{\text{base}}(t) < 1$ and strictly positive multipliers, the factor ζ_t is unique because the left side is strictly decreasing in ζ_t; the code solves it by bisection. At the terminal age, $s_{\text{base}}(T) = 0$ in every state and the factor is immaterial. Hazard ratios across health states are preserved at nonterminal ages; only their common level is anchored. The multipliers $[0.49, 1.0, 3.29]$ for Good, Fair, and Poor health are informed by unweighted complete-case person-wave intervals from waves 4–16 of the RAND HRS Longitudinal File 2022 (V1) (Health and Retirement Study, 2025). At the interval start, the respondent must be alive and interviewed (RAND `RwIWSTAT`=1) with valid age and health. Next-wave code 5 (died since the preceding interview) is a death and codes 1 or 4 are survival; codes 0, 6, 7, 9, and missing statuses are excluded. Code 6 denotes a death that occurred before a previous wave. For each age–health cell, the two-year mortality fraction m maps to the constant annual hazard $-\log(1-m)/2$, with Fair health as the within-band reference. Respondents may contribute multiple intervals. Unrounded ratios enter the model, and the table rounds them to two decimals. The estimation window includes 2020–22, but the anchor above keeps aggregate mortality fixed at the 2003 table:

Age band	μ(Good)	μ(Fair)	μ(Poor)
65–74 (midpoint 69.5)	0.49	1.00	3.29
75–84 (midpoint 79.5)	0.60	1.00	2.77
85+ (midpoint 90.0)	0.74	1.00	1.82

The focal schedule linearly interpolates the three rows at their reported midpoints, with the endpoint rows applied below 69.5 and above 90.0. A complete-game sensitivity holds the 65–74 row fixed at every age, and level checks use researcher-specified constant mappings from $[0.45, 1.0, 3.5]$ to $[0.60, 1.0, 2.0]$. The focal convention varies with age because annuity value depends on health-contingent mortality over the whole continuation horizon, including ages well after purchase.

A.4 Subjective survival beliefs

The agent uses subjective survival probabilities in the Bellman equation:

$$s^{\text{subj}}(t, H) = \psi \cdot s(t, H), \quad \psi \in (0, 1].$$

Annuity pricing uses objective survival from the population life table, with no health rating. A household with $\psi < 1$ overweights the probability of dying before collecting and so values

annuities less than under objective beliefs. Because the same ψ multiplies every health state's one-year survival, it does not change conditional survivor health shares. Cumulative subjective survival over k years is exactly ψ^k times cumulative objective survival. From age 65, the focal mapping changes survival to ages 75, 85, and 95 from 80.4%, 45.5%, and 8.6% to 66.4%, 31.0%, and 4.9%, respectively.

O'Dea and Sturrock (2023) document substantial survival pessimism at younger-old ages and optimism at sufficiently old ages. That study does not report a value for this scaling, so the focal ψ is a constant annual scale chosen to represent the near-term pessimism they document. Heimer et al. (2019) likewise document systematic age variation in mortality-belief errors, but the model does not estimate their age profile or impose a source-identified crossover age. Reported values span $\psi \in [0.96, 1.0]$, including objective beliefs, and the endpoint Shapley game shows how both participation and attribution respond under this stylized constant-wedge family.

A.5 Channel ranking at the strong-pessimism endpoint

Because both participation and attribution can depend on ψ, Table A.3 recomputes the full nine-player game at the strong-pessimism endpoint $\psi = 0.96$. The table reports every value and sign, ordered by signed value. Holding the focal $98 \times 39 \times 101$ grid fixed isolates the survival-belief endpoint from grid resolution.

Table A.3: Nine-Channel Shapley Decomposition at the Strong-Pessimism Endpoint ($\psi = 0.96$)

Channel	Shapley (pp)
Loads	+29.3
Pessimism	+17.7
Pre-existing annuitization (SS+pension-income proxy)	+16.8
Bequests	+5.2
Medical+R-S	+4.6
Inflation	+0.8
State utility	+0.3
Age needs	+0.0
Public-care aversion	-0.7

Exact Shapley values over all $2^9 = 512$ subsets of the nine-channel game at the strong-pessimism endpoint $\psi = 0.96$ of the reported range, computed on the focal production grid. Rows are ordered by signed Shapley value. Positive values are demand-suppressing and negative values are demand-raising; the table reports the resulting signs and ordering without imposing either ex ante.

A.6 Age-varying consumption needs

The age-varying needs weight is $w(t) = (1 - \delta_c)^{t-1}$. Here δ_c is a nonnegative reduced-form residual, calibrated against external spending evidence. Aguiar and Hurst (2013) report an approximately 0.30-log-point decline in household nondurable expenditure excluding housing over the latter half of their observed life cycle, roughly ages 45–75. Spread over 30 years, that decline averages about 1% a year, a figure computed here from their total. Their repeated-cross-section statistic is an age coefficient for log household expenditure after cohort and family controls. The model instead uses the endpoint log change in survivor-conditioned pooled mean consumption from ages 65 to 90, so the comparison gives only a sense of scale. With every other focal mechanism active and $\delta_c = 0$, that model path already declines by 1.95% a year. The constrained calibration therefore sets δ_c to zero, and age-varying needs is a null player in the focal game.

The calibration uses repeated simulated trajectories within each wealth band and deterministic band-specific shock streams. Table A.4 reports alternative seeds and fitting windows, including an endpoint at age 75 because the source data do not extend through the model's full retirement horizon. Every alternative also calibrates δ_c to zero. The focal Shapley table retains the player for comparability and reports its zero contribution.

A.7 State-dependent utility calibration

Finkelstein et al. (2013) estimate a 10–25% marginal-utility decline associated with a one-standard-deviation increase in the number of chronic diseases relative to no chronic disease. The health-utility weights $\varphi = [1.00, 0.92, 0.82]$ translate that evidence into three self-reported-health states, using the multiplicative formulation examined by Reichling and Smetters (2015). Neither paper estimates this vector or a mapping from chronic conditions to self-reported-health states. The focal vector is a midpoint mapping, and stronger and moderate alternatives are reported separately.

The implementation uses the zero-intercept cardinal normalization

$$u(c, H) = \varphi(H)\frac{c^{1-\gamma}}{1-\gamma}.$$

For $\gamma > 1$, the CRRA level is negative, so $\varphi(H) < 1$ makes that level less negative even though it lowers marginal utility from $c^{-\gamma}$ to $\varphi(H)c^{-\gamma}$. These weights are meant to lower marginal utility, and the less negative level in worse health says nothing about whether worse health raises welfare. An additive health–age intercept independent of consumption would not affect choices or the annuity-access value difference because survival and health

Table A.4: Reduced-Form Age-Needs Calibration Sensitivity

Simulation seed	Endpoint age	Decline at $\delta_c = 0$
42	70	1.258%
42	75	1.438%
42	80	1.576%
42	85	1.788%
42	90	1.950%
42	95	2.240%
123	70	1.255%
123	75	1.450%
123	80	1.638%
123	85	1.853%
123	90	2.060%
123	95	2.305%
777	70	1.274%
777	75	1.374%
777	80	1.628%
777	85	1.923%
777	90	2.146%
777	95	2.350%

Each cell reports the decline in survivor-conditioned pooled mean model consumption from age 65 to the listed endpoint, with the age-needs parameter δ_c set to zero. The external reference is 0.995% a year, derived from a 0.30-log-point decline over ages 45–75, and the source and model statistics are not unit-identical. Seed 42 and endpoint 90 form the headline cell. As a reduced-form residual, δ_c is not identified as a structural taste parameter. Because the constrained calibration allows only $\delta_c \geq 0$ and the decline at zero already exceeds the reference in every cell, the calibrated value is zero throughout. All cells use deterministic band-specific shock streams.

transitions do not depend on the household's portfolio decision. Such an intercept would still enter cardinal value levels and would have to be removed before using the homogeneous CEV formula. The code sets those intercepts to zero, and the reported CEVs are conditional on that convention.

At the focal midpoint $\varphi = [1.00, 0.92, 0.82]$, purchase propensity is 43.1% in the eight-channel specification that omits public-care aversion. A stronger $[1.00, 0.90, 0.75]$ mapping yields 43.1%, while the moderate $[1.00, 0.95, 0.85]$ mapping yields 43.0%, a moderate-minus-stronger change of -0.1 pp. State-dependent utility's exact Shapley value is 0.1 pp.

A.8 Public-care-aversion proxy

Strategic-survey evidence in Ameriks et al. (2020) motivates the public-care-aversion player, whose coefficient and care-state definition are calibration choices. Operationally, $\chi_{\text{LTC}} < 1$ applies only when medical risk is active, the generic public consumption floor binds, and model health is Poor. The floor amount is anchored to SSI, and the medical-need target includes out-of-pocket and Medicaid-paid spending, but neither identifies a respondent's Medicaid eligibility or receipt. No long-term-care need, insurance, setting, or utilization state appears in the model. Accordingly, the trigger is a narrow, assumed proxy for reliance on publicly financed care. The "LTC" subscript on the parameter matches the replication code and does not denote a modeled long-term-care state. Applying χ_{LTC} inside CRRA as $u(\chi_{\text{LTC}}c)$ ensures that a lower value reduces consumption-equivalent welfare for any $\gamma > 0$. Multiplying the negative CRRA level by $\chi_{\text{LTC}} < 1$ would instead reverse that ordering when $\gamma > 1$.

A.9 Behavioral parameter sensitivity

The behavioral parameters λ_W and ψ_{purchase} are exploratory values, and the model neither identifies them nor matches them to moments. Source-dependent utility discounts consumption financed by liquid-wealth drawdown:

$$c_{\text{eff},t} = c_{\text{income},t} + \lambda_W c_{\text{portfolio},t}, \qquad \lambda_W \in (0, 1]. \tag{A.1}$$

Income-financed consumption includes private-annuity and pre-existing annuitized income, and portfolio-financed consumption is the remainder, financed from liquid wealth. The value $\lambda_W = 0.625$ is set to the ratio of the roughly 50% utilization of non-lifetime-income sources to the roughly 80% utilization of lifetime income reported descriptively by Blanchett and Finke (2025). Related estimates in Blanchett and Finke (2024) imply that investment assets

generate about half as much additional spending as wealth held in guaranteed income. Neither study estimates a source-dependent utility coefficient, and their observational spending relationships do not identify this model parameter.

Medical spending is assigned to current income first and then to liquid wealth when equation (A.1) is evaluated, and any public-floor top-up counts as income-financed consumption. This accounting order is a modeling convention that can affect attribution within the exploratory extension but not the focal nine-player game, where $\lambda_W = 1$.

The purchase-event extension applies a one-time narrow-framing penalty equal to the discounted, objective-survival-weighted value of the annuity's running underwater balance. Let $t = 1$ denote the purchase period within this mental account, let τ denote the corresponding model period, and define

$$S_\tau(1) = 1, \qquad S_\tau(t) = \prod_{j=0}^{t-2} s_{\text{base}}(\tau + j) \quad (t \geq 2).$$

With premium $P = \alpha W_0$, fixed scheduled nominal payment $A = P\rho$ measured in purchase-period dollars, and first-year payout rate ρ, the last potentially positive underwater term within the finite model horizon is

$$t_{\max} = \min\left\{ \left\lceil \frac{1}{\rho} \right\rceil,\ T - \tau + 1 \right\}.$$

The penalty is

$$\Delta V_{\text{purchase}} = -\psi_{\text{purchase}} u'(c_{\text{ref}}) \sum_{t=1}^{t_{\max}} \beta^{t-1} S_\tau(t) \max\{0, P - A(t-1)\} \mathbf{1}\{\alpha > 0\}. \tag{A.2}$$

Here $c_{\text{ref}} = \$18{,}000$ and $u'(c_{\text{ref}}) = c_{\text{ref}}^{-\gamma}$. The balance is measured immediately before each beginning-of-period payment, so the purchase-period term has no prior payout. That balance cumulates scheduled nominal dollars and does not restate later payments for inflation. Population survival prices the contract and weights this penalty, while health-conditioned subjective survival continues to govern household continuation utility. By construction, $\psi_{\text{purchase}} = 0.05$, the reference-consumption level, and equation (A.2) are this paper's reduced form. Annuity-as-gamble framing is motivated by Hu and Scott (2007) and Brown et al. (2008), while Barberis and Huang (2009) provide a general framed-preference formulation. Chalmers and Reuter (2012) document weak demand responses to economically meaningful but less salient annuity-price variation. None of these papers estimates this equation, its reference-consumption normalization, or ψ_{purchase}.

The reported extension evaluates a single illustrative pair, $\lambda_W = 0.625$ and $\psi_{\text{purchase}} =$

0.05. Neither parameter enters the focal nine-player result of 37.7%.

For q286, the life-contingent-payment ownership stock is 3.85%, with a 95% interval of $[2.85\%, 5.18\%]$. By contrast, the conventional any-annuity-income proxy is 6.12%, with a 95% interval of $[4.83\%, 7.72\%]$. Both intervals use the conservative effective-sample-size convention described in the main text. Neither measure enters the demand calibration.

Eleven-channel extended Shapley. Table A.5 reports the exact Shapley decomposition over all $2^{11} = 2{,}048$ subsets of the game that adds the two exploratory behavioral parameters. Along the sequential extension path, source-dependent utility raises purchase propensity from 37.7% to 53.4%, after which the purchase-event penalty moves the eleven-channel specification to the 0.0% participation boundary. The exact Shapley values are -9.6 pp for source-dependent utility and 78.8 pp for the narrow-framing penalty. Under the table's sign convention, the negative value for source-dependent utility means that it raises demand. Because neither behavioral magnitude is identified and the purchase penalty reaches a participation boundary, these values and ranks describe the chosen sensitivity specification.

Table A.5: Exact Shapley-Value Decomposition of Predicted Purchase Propensity

Channel	Shapley (pp)	Share (%)
Pre-existing annuitization (SS+pension-income proxy)	+8.34	8.3
Bequests	+0.49	0.5
Medical+R-S	+4.93	4.9
Pessimism	+3.58	3.6
Age needs	+0.00	0.0
State utility	+0.68	0.7
Pricing frictions	+12.45	12.5
Inflation	-0.27	-0.3
Public-care aversion	+0.63	0.6
Source-dependent utility (SDU)	-9.64	-9.6
Narrow-framing penalty (PED)	+78.80	78.8
Total	+100.00	100.0

Exact Shapley values computed from all 2048 channel subsets. Each value represents the weighted average marginal purchase-propensity reduction (pp) across all coalition orderings. Equal-resource benchmark: 100.0%. Full model: 0.0%.

A.10 Sequential paths

A sequential path is useful for tracing the model but does not attribute interactions uniquely. One table, Table A.6, adds the focal structural players in one order. A second, Table A.7, follows a rational-channels-first path through the preference channels and then the exploratory behavioral extension. Their increments differ because the player entering later receives the interaction. The main paper therefore uses the exact Shapley allocation for attribution.

Table A.6: Sequential Decomposition of Predicted Private-Annuity Purchase Propensity

Channel	Prop. (%)	Δ (pp)	Ret. (%)	Cum.
Equal-resource empty coalition	100.0	—	—	—
+ Pre-existing annuitization	100.0	+0.0	100.0%	1.0000
+ Bequest motives	100.0	+0.0	100.0%	1.0000
+ Medical expenditure risk (uncorrelated)	100.0	+0.0	100.0%	1.0000
+ Health-mortality correlation (R-S)	97.5	-2.5	97.5%	0.9752
+ Survival pessimism	82.9	-14.6	85.0%	0.8291
+ State-dependent utility	83.8	+0.9	101.1%	0.8381
+ Age-varying consumption needs (calibrated null)	83.8	+0.0	100.0%	0.8381
+ Realistic pricing loads	44.7	-39.1	53.3%	0.4467
+ Inflation erosion	43.1	-1.6	96.5%	0.4309
+ Public-care aversion	37.7	-5.4	87.5%	0.3771
HRS q286 stock (untargeted)	3.85%	—	—	—
HRS annuity-income proxy (untargeted)	6.12%	—	—	—

Retention rate = purchase propensity after the channel divided by purchase propensity before it. The cumulative product is a descriptive summary of this ordering, not a channel attribution.
The two HRS entries are distinct, nonnested current reported ownership and income stocks that may reflect earlier or nonretail choices. They are untargeted comparators, not bounds of an observed interval and not the same estimand as the model's current-state purchase propensity.

Table A.7: Sequential Channel Decomposition

Specification	Propensity (%)	Δ (pp)
Six rational channels	42.9	—
+ Age-varying consumption needs	42.9	+0.0
+ State-dependent utility	43.1	+0.2
+ Public-care aversion (nine-channel baseline)	37.7	-5.4
+ Source-dependent utility (exploratory)	53.4	+15.7
+ Narrow-framing penalty (exploratory)	0.0	-53.4

Sequential extension path from the six rational channels through the nine-channel structural baseline to the exploratory eleven-channel specification. The final two rows use the single disclosed pair $\lambda_W = 0.625$ and $\psi_{\text{purchase}} = 0.05$; they are a point sensitivity to researcher-chosen magnitudes, not a reported range.

A.11 Nominal annuity pricing derivation

The nominal rate is defined by the exact Fisher relation $r_{\text{nom}} = (1+r)(1+\pi) - 1$. For a purchase in model period τ, the present value of one unit of beginning-of-period nominal payment, discounted at that rate, is

$$\text{PV}^N_\tau = \sum_{k=0}^{T-\tau} \frac{\prod_{\ell=0}^{k-1} s_{\text{base}}(\tau+\ell)}{(1+r_{\text{nom}})^k},$$

where the product is one at $k = 0$. The fair payout rate is $q^N_\tau = 1/\,\text{PV}^N_\tau$, and the loaded payout rate is $\text{MWR}\, q^N_\tau$.

The consumer receives a fixed nominal payment after purchase. If its real value in period τ is A^{real}_τ, its real value k periods later is $A^{\text{real}}_\tau/(1+\pi)^k$. The initial real payout under nominal pricing exceeds the payout under real pricing (which uses discount rate r), but declines in purchasing power over time.

When $\pi = 0$, $r_{\text{nom}} = r$ and the nominal and real pricing formulas coincide. The Yaari-like benchmark (Yaari, 1965), which assumes $\pi = 0$, is therefore unaffected by the pricing convention.

B Grid Convergence

Table B.1 reports predicted purchase propensity and mean annuitization share ($\bar{\alpha}$) under the full nine-player model, computed with the focal wealth-band Social Security and employer-pension-income-proxy configuration. Panel A holds medical integration at the focal state-split Gauss–Legendre order 32 and varies the wealth–annuity grid; Panel B holds the grid at 98×39 and varies the order applied on each probability segment; Panel C reports the fully refined grid at orders 48 and 64, with order 64 serving as the combined reference; Panel D varies the annuitization grid n_α.

Because the diagnostic uses the focal wealth-band assignment, the focal cell (98×39, state-split order 32) reproduces the full-model purchase propensity of 37.7% exactly. Holding the integration rule fixed in Panel A, propensity is 37.7% at the focal piecewise H0 grid, 37.7% under H1 tail refinement (114×47), 37.1% under H2 full refinement (195×77; the step between the two nested refinements is 0.6 pp), and 37.7% on a conventional single-power 170×50 grid. Refinement can be non-monotone near the discrete extensive margin at the fixed-cost threshold. Panel B separately documents the integration-order path. Forcing every evaluation to age 65 gives 33.8% purchase propensity, with pricing frictions still first at 29.4 pp and bequests still fifth at 3.5 pp.

Table B.1: Grid, Quadrature, and Annuitization-Grid Convergence (Headline Wealth-Band Configuration)

Specification	$n_W \times n_A$	GL order	Purchase prop. (%)	$\bar{\alpha}$
Panel A: Grid convergence (state-split GL, order 32 per segment)				
H0 piecewise (production)	98×39	**32**	**37.71**	**0.1935**
H1 tail-refined	114×47	32	37.71	0.1932
H2 fully refined	195×77	32	37.13	0.1925
Conventional global power	170×50	32	37.71	0.1931
Panel B: Quadrature convergence (98×39 grid)				
	98×39	8	35.76	0.1770
	98×39	12	36.81	0.1855
	98×39	16	37.13	0.1886
	98×39	24	37.55	0.1922
	98×39	**32**	**37.71**	**0.1935**
	98×39	48	37.71	0.1939
Panel C: Fully refined-grid quadrature references				
H2 piecewise	195×77	48	37.29	0.1939
H2 piecewise (combined reference)	195×77	64	37.34	0.1943
Panel D: Annuitization grid (98×39 grid, state-split GL order 32)				
$n_\alpha = 51$	98×39	32	37.66	0.1930
$n_\alpha = 101$	98×39	**32**	**37.71**	**0.1935**
$n_\alpha = 201$	98×39	32	37.71	0.1934
$n_\alpha = 401$	98×39	32	37.71	0.1936

H0 uses a piecewise nonuniform tensor grid with power-spaced core nodes through \$3 million and cosine-spaced tail nodes through the unchanged \$13 million support bound. H1 inserts tail midpoints; H2 also inserts core midpoints, so the ladder is exactly nested. The conventional global-power row retains the former 170-by-50 topology as a separate comparator. All specifications use baseline parameters ($\gamma = 2.5$, DFJ bequests, MWR $= 0.87$, $\pi = 2\%$, $\psi = 0.981$) under the production fixed-wealth-band Social Security plus pension-income-proxy assignment. The reported Gauss–Legendre order applies separately on each probability segment; the saturated upper tail is evaluated exactly. Panel C reports the H2 high-order ladder, with order 64 serving as the combined grid-and-quadrature reference. $n_\alpha = 101$ in Panels A–C; Panel D varies it.

Panel B keeps the focal 98×39 grid fixed and evaluates per-segment orders 8, 12, 16, 24, 32, and 48. The corresponding purchase propensities are 35.8%, 36.8%, 37.1%, 37.6%, 37.7%, and 37.7%. At order 32, the focal value differs from order 48 by 0.0 pp; the spread across orders 24, 32, and 48 is 0.2 pp. Order 24 is part of this separately gated high-order band. Orders 8, 12, and 16 are coarser stress diagnostics with no convergence budget. They can still move the discrete purchase threshold ($\alpha^* > 0$).

State-split Gauss–Legendre order 32 per probability segment is the focal standard. Taking the most-refined specification (Panel C: the fully refined 195×77 grid, order 64) as reference, the combined grid-and-integration-order difference is approximately 0.4 pp. Panel D varies the age-65 annuitization grid n_α over $\{51, 101, 201, 401\}$ while holding the wealth grid and medical integration at focal resolution. Quadrature and mean-α budget values were fixed before any production output on the final inputs existed. Three grid-ownership budgets were set once from a non-citable pre-production refinement pilot on those inputs, in which the net ownership-count step fell from 0.580 percentage point under the first core refinement to 0.158 percentage point under the next. These two steps show contraction over the tested grids without identifying a geometric rate or a limit.

Non-citable local checks on the fully refined grid showed quadrature sensitivity that the four-grid pilot had not measured. Panel C reports the corresponding production-run quantities. On H2, the order-32-to-48 gaps are 0.158 percentage points for participation and 0.001319 for mean α; the order-48-to-64 gaps are 0.053 percentage points and 0.000417. The first mean-α difference in the local check exceeded the strict 0.001 last-step budget, whereas the second satisfied it. No tolerance value was created or enlarged after observing that breach. Instead, the already fixed high-band budgets of 0.50 percentage point and 0.0025 govern order-32-to-48 quadrature sensitivity conditional on H2, while the already fixed strict budgets of 0.25 percentage point and 0.001 govern the H2 order-48-to-64 last step. Comparing this H2 response with the H0 order-32-to-48 response exposes the grid-by-quadrature interaction, while the H2 step itself changes only quadrature order. At H0, the direct order-32-to-48 comparison retains the strict budgets, and the separately gated order-24/32/48 band may span at most 0.50 percentage point and 0.0025. The production-run H0-to-H2/order-64 combined-reference gaps are 0.4 percentage points and 0.000770, evaluated against the existing 1.00-point and 0.005 budgets.

The realized production comparisons satisfy each corresponding gate. Specifically, the H1-to-H2 ownership step is 0.6 percentage point, below its 0.75-point budget; the H0-to-H2/order-64 combined-reference gaps are 0.4 percentage points and 0.000770, below 1.00 and 0.005; and the H2 order-48-to-64 last step is 0.053 percentage points and 0.000417, below 0.25 and 0.001. These comparisons cover only the tested grids and orders.

For the nested grid refinements (114×47 to 195×77), the budgets are 0.75 percentage point and 0.0025. Focal-to-tail-refined, focal-to-fully-refined, focal-to-single-power-grid, and focal-to-combined-reference comparisons have budgets 1.00 percentage point and 0.005, while the refined $n_\alpha \in \{101, 201, 401\}$ band has budgets 0.25 percentage point and 0.005. Exact row rosters and independently regenerated focal-cell identities are also required. Monotonic convergence is not assumed. The budget assignment, though not the budget values, post-dated the first local check, and no tolerance was enlarged after production output on the final inputs existed. Across the H0, H1, and H2 complete games (Table B.2), the focal leading order remains pricing frictions, pre-existing annuitization, and the medical-expenditure process plus health-dependent mortality. That stability concerns grid resolution only, and the floor, fixed-cost, mortality, risk-aversion, and player-partition exercises can still change channel magnitudes or ordering.

Table B.2: Nine-Channel Shapley Attribution Across Wealth–Annuity Grids

Channel	H0 98×39 production	H1 114×47 tail-refined	H2 195×77 fully refined
Full-model propensity (%)	37.7	37.7	37.1
Pricing frictions	+28.1 [1]	+28.1 [1]	+27.9 [1]
Pre-existing annuitization	+15.1 [2]	+15.1 [2]	+15.9 [2]
Medical+R-S	+8.4 [3]	+8.4 [3]	+8.5 [3]
Pessimism	+6.3 [4]	+6.3 [4]	+6.3 [4]
Bequests	+2.7 [5]	+2.7 [5]	+2.6 [5]
Inflation	+1.3 [6]	+1.3 [6]	+1.2 [6]
Public-care aversion	+0.4 [7]	+0.4 [7]	+0.2 [7]
State utility	+0.1 [8]	+0.1 [8]	+0.1 [8]
Age needs	+0.0 [—]	+0.0 [—]	+0.0 [—]

Cells report signed Shapley contributions in percentage points, with competition rank in brackets. Exact null players display an em dash. Positive values suppress purchase propensity. Every column is a complete 512-coalition nine-channel game under the production state-split integration rule. H1 refines only the H0 cosine tails; H2 also refines both cores. The H0 production column is reused bit-for-bit from the authenticated focal-game cache.

Consumption optimizer. A global search selects consumption from the feasible interval at each medical-shock node. It compares all interval boundaries and admissible interior maxima of the objective with piecewise-linear continuation, and it allows the continuation slopes to rise between wealth intervals, so the full objective need not be concave. Section K describes the interval search and its derivative-bisection step. Globality applies to the represented objective within floating-point and root-bracket tolerances. Grid interpolation, medical integration, and the discrete purchase grid remain numerical approximations assessed by the separate refinement checks.

C Bequest Specification

The De Nardi–French–Jones (DFJ) bequest utility is $v(b) = \theta(b+\kappa)^{1-\gamma}/(1-\gamma)$. In this utility, κ controls the luxury-good nature of bequests. When κ is large relative to wealth b, marginal bequest utility $v'(b) = \theta(b+\kappa)^{-\gamma}$ is nearly constant across wealth levels, making bequests a luxury good: the bequest-to-wealth ratio rises with wealth.

The parameters originate in De Nardi et al. (2010). De Nardi, French, and Jones estimate a bequest intensity of 2,360 and a shifter of $273,000 using AHEAD. Their model's monetary quantities are in 1998 dollars, and the associated estimate of the consumption-risk-aversion coefficient is 3.84. They caution that the bequest parameters are imprecisely estimated. Lockwood (2012) uses a related DFJ calibration in an annuity model with $\sigma = 2$, but neither paper reports or re-estimates the implementation pair $\theta = 158.68$ and $\kappa = \$396{,}496$. The model obtains that pair through the portability normalization below and restates the shifter with annual-average CPI-U (163.0 in 1998 and 236.736 in 2014; U.S. Bureau of Labor Statistics, 2017). It does not treat the HRS interview year as the shifter's price base.

C.1 Bequest parameter portability

In the unweighted deterministic terminal problem, the objective is $u(c) + \beta\theta u(b+\kappa)$. The first-order condition therefore depends on β and θ only through $q \equiv (\beta\theta)^{-1/\gamma}$. Bequests begin when resources exceed $q\kappa$, and the marginal propensity to bequeath above that threshold is $1/(1+q)$. At the DFJ source estimates, these quantities are approximately $36,400 in 1998 dollars and 0.88, respectively. The focal portability rule holds q fixed when curvature or the discount factor changes, so

$$\theta(\gamma, \beta) = \frac{(0.97 \times 2360)^{\gamma/3.84}}{\beta}.$$

At $\gamma = 2.5$ and $\beta = 0.97$, this rule gives the focal intensity $\theta = 158.68$, and the shifter is restated separately to $\kappa = \$396{,}496$ in 2014 dollars. This mapping preserves the two unweighted deterministic source quantities by construction. Neither De Nardi et al. (2010) nor Lockwood (2012) reports this normalization. The normalization need not preserve bequest moments or terminal allocations in the full lifecycle model, which adds state-dependent utility, medical-expenditure risk, stochastic mortality, and survival pessimism.

Table C.1 checks the resulting mapping across curvature values, verifying the invariant q, restated threshold, and marginal propensity. The focal purchase propensities and Shapley values depend on this portability choice.

Table C.1: Curvature Portability of the DFJ Bequest Pair

	$\gamma = 2.00$	$\gamma = 2.50$	$\gamma = 3.00$	$\gamma = 3.84$
$\theta(\gamma)$	57.95	158.68	434.49	2360.00
$(\beta\theta(\gamma))^{-1/\gamma}$	0.1334	0.1334	0.1334	0.1334
Bequest threshold (2014 dollars)	\$52 884	\$52 884	\$52 884	\$52 884
Marginal propensity to bequeath	0.882	0.882	0.882	0.882

DFJ estimate $\gamma = 3.84$, $\beta = 0.97$, $\theta = 2360$, and $\kappa = \$273\,000$ in 1998 dollars. The model restates κ with annual-average CPI-U (236.736/163.0) and maps the coefficient as $\theta(\gamma, \beta) = (0.97 \times 2360)^{\gamma/3.84}/\beta$.
The mapping preserves the source deterministic bequest threshold and marginal propensity. It is a units-and-curvature normalization, not an estimate for the HRS analysis population.

D Full Sensitivity Analysis

Tables D.1 and D.2 report selected one-way purchase-propensity sensitivities. Complete-game attribution checks and numerical diagnostics appear in the surrounding tables; the machine-readable `robustness_full.csv` retains the full level-sensitivity roster.

D.1 Consumption-floor and fixed-cost attribution robustness

The public consumption floor is held fixed outside the focal cooperative game, so the nine-player values are conditional on that institution. Table D.3 recomputes the complete 512-coalition game with the focal floor and with a near-zero \$1 floor (exact zero is inadmissible under CRRA utility). Both games use the predeclared extended-support floor grid, which keeps every H2 node through the \$13 million initial-state bound and adds a continuation-only wealth tail through \$125 million. Because the unextended H0, H1, and H2 grids do not certify every evaluated high-wealth/full-coalition interpolation stencil for the near-zero-floor game, the focal-floor game uses the same extended grid, so the comparison never mixes state grids. The table reports the empty- and full-coalition endpoints and the modeled drop before the values and ranks, so it shows how both the level and the attribution depend on the safety net.

On the common extended grid, full-model purchase propensity is 37.1% with the production floor and 8.8% with the near-zero floor. Under the production floor, pricing frictions contribute 27.9 pp and rank first, pre-existing annuitization contributes 15.9 pp and ranks second, and the medical-expenditure/mortality player contributes 8.5 pp and ranks third. Near-zero-floor results put pre-existing annuitization first (49.4 pp), the medical-expenditure/mortality player second (21.6 pp), and pricing frictions third (14.7 pp). The

Table D.1: Full Sensitivity Results

Category	Specification	Purchase propensity (%)
Risk aversion (γ)		
	$\gamma = 1.50$	1.6
	$\gamma = 2.00$	30.3
	$\gamma = 2.20$	34.8
	$\gamma = 2.30$	36.0
	$\gamma = 2.40$	37.0
	$\gamma = 2.5$ (baseline)	37.7
	$\gamma = 2.60$	38.0
	$\gamma = 3.00$	38.5
	$\gamma = 3.50$	38.4
	$\gamma = 4.00$	37.3
	$\gamma = 5.00$	34.7
Money's worth ratio		
	MWR = 0.82	27.8
	MWR = 0.85	33.6
	MWR = 0.87 (baseline)	37.7
	MWR = 0.90	42.9
	MWR = 0.95	49.9
Hazard multipliers $[\mu_G, \mu_F, \mu_P]$		
	Age-varying HRS estimates (baseline)	37.7
	Researcher-specified steep stress $[0.45, 1.0, 3.5]$	36.4
	Constant entry-age row $[0.49, 1.0, 3.29]$	36.8
	HRS SRH $[0.57, 1.0, 2.7]$	37.0
	Conservative $[0.60, 1.0, 2.0]$	37.8
Inflation (π)		
	$\pi = 1\%$	39.1
	$\pi = 2\%$ (baseline)	37.7
	$\pi = 3\%$	35.9
Survival pessimism (ψ)		
	$\psi = 0.960$ (strong-pessimism endpoint)	26.1
	$\psi = 0.970$	31.8
	$\psi = 0.981$ (baseline)	37.7
	$\psi = 0.990$	42.1
	$\psi = 1.000$ (objective)	47.4

Unless otherwise noted, all specifications use the focal parameters: $\gamma = 2.5$, $\beta = 0.97$, DFJ bequests ($\theta = 158.68$, $\kappa = \$396{,}496$), survival pessimism $\psi = 0.981$, MWR = 0.87, $F = \$2{,}349$, $\pi = 2\%$, age-varying health-hazard ratios, state-split probability-domain Gauss–Legendre medical integration at order 32 per probability segment, and the focal grid ($98 \times 39 \times 101$). In the γ sweep, the linked DFJ intensity $\theta(\gamma, \beta)$ is recomputed by the portability rule in Section C.1; it is not held at its focal numeric value. All sweeps begin from the nine-player specification with purchase propensity 37.7%.

Table D.2: Full Sensitivity Results (continued)

Category	Specification	Purchase propensity
Bequest specification		
	No bequest ($\theta = 0$)	42.7%
	DFJ luxury (baseline)	37.7%
Minimum purchase		
	\$0	37.7%
	\$1,000	37.7%
	\$5,000	37.7%
	\$10,000 (baseline)	37.7%
	\$25,000	37.7%

Notes as in Table D.1. The table reports the minimum-purchase sweep directly; grid and quadrature convergence are reported separately in Table B.1.

same small leading set therefore remains prominent, but both the modeled level and the ordering of the leading set depend on the safety net.

The pricing-frictions player bundles the proportional money's-worth wedge, the minimum premium, and a fixed purchase cost. Table D.4 recomputes the complete game at \$500, \$1,000, \$2,000, and the CPI-restated focal value of \$2,349 while holding the other two pricing components fixed. The table reports the resulting endpoints, contributions, and ranks. A separate loads-split game in Appendix L changes the player partition.

D.2 Robustness to channel partition

The nine-player game combines the calibrated medical-expenditure process with the Reichling–Smetters correlation and combines Social Security with the employer-pension-income proxy. Turning the expenditure component off sets medical costs to zero, so the toggle removes the mean as well as the risk. Within the first bundle, mortality correlation operates through health states that also shift spending. In the second, the retirement-income flows share the maintained life-contingency assumption but differ in payment path (real SS, nominal pension proxy). Table D.5 recomputes two ten-player games, first separating the medical-expenditure process from the correlation and then separating SS from the pension proxy. The HRS pension-income measure does not identify DB plan type or life contingency, so the second partition separates the model's two income components without verifying any pension contract.

Table D.3: Nine-Channel Attribution under Alternative Consumption Floors

	Production floor		Near-zero floor	
Outcome or channel	Value	Rank	Value	Rank
Equal-resource benchmark (%)	100.0	—	100.0	—
Full-model propensity (%)	37.1	—	8.8	—
Total modeled drop (pp)	62.9	—	91.2	—
Pricing frictions	+27.9	1	+14.7	3
Pre-existing annuitization	+15.9	2	+49.4	1
Medical+R-S	+8.5	3	+21.6	2
Pessimism	+6.3	4	+3.2	4
Bequests	+2.6	5	+1.1	5
Inflation	+1.2	6	+0.7	6
Public-care aversion	+0.2	7	+0.3	7
State utility	+0.1	8	+0.3	8
Age needs	+0.0	—	+0.0	—

Exact Shapley values over all $2^9 = 512$ subsets with the public consumption floor fixed outside the game, at its production value ($8496 in 2014 dollars; CPI-U-restated from Lockwood's $6,180 year-2000 amount) and at $1. Both games use the predeclared extended-support floor grid: every H2 node through the audited $13 million initial-state bound, plus a 128-interval continuation-only cosine tail through $125 million (323 wealth by 77 annuity-income nodes overall). H0, H1, and the unextended H2 grid do not certify every evaluated high-wealth/full-coalition interpolation stencil under the near-zero floor. The endpoint rows show that the modeled level and attribution are conditional on the safety net. Positive values suppress demand. Competition ranks are shown; numerical nulls and endpoints receive an em dash. Values sum to each reported empty-to-full drop.

Table D.4: Leading Channel Attributions Across the Fixed-Cost Range

Channel	$500	$1000	$2000	$2349
Full-model propensity (%)	45.4	42.9	38.8	37.7
Pricing frictions	+22.4 [1]	+24.4 [1]	+27.3 [1]	+28.1 [1]
Pre-existing annuitization	+10.6 [3]	+12.3 [2]	+14.5 [2]	+15.1 [2]
Medical+R-S	+11.6 [2]	+10.7 [3]	+8.9 [3]	+8.4 [3]
Pessimism	+6.9 [4]	+6.4 [4]	+6.3 [4]	+6.3 [4]
Bequests	+1.8 [5]	+2.0 [5]	+2.5 [5]	+2.7 [5]
Inflation	+0.9 [6]	+1.0 [6]	+1.3 [6]	+1.3 [6]
Public-care aversion	+0.5 [7]	+0.4 [7]	+0.4 [7]	+0.4 [7]
State utility	-0.1 [8]	-0.1 [8]	+0.1 [8]	+0.1 [8]
Age needs	+0.0 [—]	+0.0 [—]	+0.0 [—]	+0.0 [—]

Cells report signed Shapley contribution in percentage points, with competition rank in brackets. Exact null players display an em dash rather than an arbitrary rank. Values are computed over all $2^9 = 512$ subsets of the nine-channel game at each fixed cost. The proportional MWR wedge and minimum-purchase requirement in the pricing-frictions player are held at production values; only the fixed cost varies across $500, $1,000, $2,000, and the CPI-restated production value of $2,349. Positive values suppress demand; the table reports rather than presumes rank stability.

D.3 HRS weighting and asset-definition robustness

The baseline HRS evaluation pools eligible person-waves without survey weights and maps nonhousing net worth into model wealth. Before band aggregation, each observed-age pension-income proxy is grossed to the model's age-65 nominal base at the maintained 2% erosion rate, and the untransformed pooled mean is kept as a diagnostic. Interview age supplies the model-state timing convention because RAND reports prior-calendar-year income without exact within-year payment timing. Table D.6 separates weighting/repeated-observation choices from asset-definition and income-aggregation choices. Every row is a full-model level sensitivity. The coalition game is not recomputed under these constructions, so their effect on the allocation is unknown. A narrow liquid-wealth row reselects eligibility and changes model wealth while retaining the baseline Social Security and pension-income-proxy assignment. It excludes IRAs, so it is a deliberately narrow sensitivity. Separately, a Social Security current-payment row replaces the focal positive-recipient entitlement proxy with the all-respondent band mean, including zero current payments. A further row replaces each band's unconditional pension-income mean with a mixture of zero-pension and positive-recipient cells, using the observed recipient rate and conditional mean. That mixture is not linked at the person level and does not validate life contingency, but it exposes the nonlinearity hidden by an unconditional band mean.

Table D.5: Shapley Attribution Under Alternative Channel Partitions

Channel	Shapley (pp)	Rank
Panel A: medical risk and R-S correlation unbundled		
Loads	+28.06	1
SS+pension-income proxy	+14.85	2
Medical	+8.20	3
Pessimism	+6.20	4
Bequests	+2.61	5
Inflation	+1.24	6
R-S corr	+0.90	7
Public-care aversion	+0.27	8
State util	-0.05	9
Age needs	+0.00	—
Panel B: Social Security and employer-pension-income proxy split		
Loads	+20.29	1
Pension-income proxy	+15.41	2
SS income	+11.11	3
Med+R-S	+6.75	4
Pessimism	+5.24	5
Bequests	+2.08	6
Inflation	+0.99	7
Public-care aversion	+0.41	8
State util	-0.01	9
Age needs	+0.00	—

Exact Shapley values over all $2^{10} = 1024$ subsets of each 10-channel game, production grid and calibration. Panel A separates medical-expense risk from the Reichling–Smetters health–mortality correlation; Panel B separates Social Security income from the RAND employer-pension-income proxy. The model treats that proxy as a nominal life-contingent flow; RAND does not identify plan type or life contingency. A positive value is demand-suppressing. Exact null players receive no substantive rank; numerical ties use competition ranks.

Table D.6: HRS Comparator and Model-Population Robustness

Specification	Rate (%)	N
Panel A: HRS q286 life-contingent-payment stock		
Unweighted person-wave	3.85	1896
Survey-weighted person-wave	3.86	1896
Person level: at least one eligible observation	4.96	1269
Panel B: Model-predicted purchase propensity		
Unweighted person-wave	37.71	1896
Survey-weighted person-wave	40.56	1896
Person-balanced (earliest wave)	35.22	1269
Financial+IRA wealth and eligibility	37.76	1507
Liquid band; baseline wealth	40.35	1896
Narrow liquid wealth and eligibility	28.91	1387
SS current-payment mean (including zeros)	38.92	1896
Pension-recipient mixture	36.45	1896
Panel C: Pension-recipient mixture by baseline wealth band		
\$5–30k	0.00	574
\$30–120k	8.65	532
\$120–350k	66.69	420
Above \$350k	98.65	370
All bands	36.45	1896

Panel A reports the loose q286 stock comparator as unweighted and survey-weighted person-wave rates and as the share of persons with a qualifying response in at least one eligible observation. The last statistic is not lifetime cumulative purchase incidence. A separate disclosure-safe robustness CSV retains $r\{w\}iann > 0$ as a secondary proxy and applies alternative asset-eligibility definitions. Panel B reports model purchase propensities; "financial+IRA" retains retirement accounts, while "narrow liquid wealth" excludes IRAs, and neither is preferred. Model rows retain the focal positive-recipient Social Security entitlement proxy and pension-income assignment except the liquid-band reassignment, the pension-recipient mixture, and the labeled all-respondent current-payment sensitivity. The current-payment row includes zero current Social Security payments and is a level check, not a complete-game recomputation. Panel C splits each wealth band into zero and positive pension-income-proxy cells and prevalence-weights their propensities. The HRS employer-pension-income measure identifies neither plan type nor life contingency, so this is a nonlinear-aggregation sensitivity, not person-linked validation. Production grid.

E State-Split Medical-Shock Quadrature Check

Medical expenditure shocks satisfy $\ln m \sim N(\mu, \sigma^2)$ with base dispersion $\sigma = 1.4$. Direct Gauss–Hermite integration is poorly behaved here because its global nodes move across the state-dependent income and public-floor kinks as the order changes. The implementation instead transforms the shock to $u = F_m(m)$, splits the unit interval at every accounting kink induced by the current state, and applies Gauss–Legendre quadrature separately on each resulting interval. Inside an interval, the interpolated continuation value can still have kinks at grid knots. Once medical spending exhausts the relevant resources, every larger shock maps to the same floor-supported continuation state; that saturated upper tail is therefore evaluated as one exact probability mass rather than approximated with extreme nodes.

Table B.1 (Panel B) documents convergence in the Gauss–Legendre order used on each probability segment. The focal rule uses order 32; at H0, order 48 is the direct higher-order same-rule check. Panel C adds the H2 order-48-to-64 ladder, with H2/order 64 serving as the combined reference. Because purchase is the discrete event $\alpha^* > 0$, small expectation errors can still move marginal households across the extensive margin, so the table reports the continuous mean purchased share alongside purchase propensity.

The convergence target is the state-dependent Bellman expectation. State splitting aligns the numerical rule with the model's nonlinear resource accounting, while the H0 order-32-to-48 and H2 order-48-to-64 exercises assess the sensitivity of purchase propensity and mean annuitization to quadrature order.

F Pairwise Channel Interactions

Table F.1 reports pairwise interaction strengths for the same nine structural players used in the focal Shapley game. The age-needs player remains in the table when its calibrated value is zero, in which case its row and column are null. For a pair (A, B), the reported interaction is the purchase propensity with both players active minus the additive prediction $v(A) + v(B) - v(\varnothing)$. A negative value therefore means that the joint propensity reduction exceeds the sum of the two isolated reductions. This table conditions on the empty predecessor set. The order-independent contributions of these nine players appear in Section 2.11; Appendix A.9 reports the exploratory eleven-player extension.

The symmetric interaction is printed once, in the table's upper triangle, and dashes in the lower triangle do not denote zero. These finite differences describe one two-player projection, and their interpretation can change when other players are present. In particular, the pre-existing-annuitization player is the equal-PV resource-form experiment and does not

Table F.1: Pairwise Interaction Strengths, Rational and Preference Channels (pp)

	Pre-ann.	Bequests	Med+R–S	Loads	Inflation	Pessim.	State util.	Age needs	Public care
Pre-ann.	—	+0.0	-2.5	-34.7	+0.0	+0.0	+0.0	+0.0	+0.0
Bequests	—	—	+0.0	+0.0	+0.0	+0.0	+0.0	+0.0	+0.0
Med+R–S	—	—	—	-37.7	+0.0	-12.9	+0.0	+0.0	+0.0
Loads	—	—	—	—	+0.0	+0.0	+0.0	+0.0	+0.0
Inflation	—	—	—	—	—	+0.0	+0.0	+0.0	+0.0
Pessim.	—	—	—	—	—	—	+0.0	+0.0	+0.0
State util.	—	—	—	—	—	—	—	+0.0	+0.0
Age needs	—	—	—	—	—	—	—	—	+0.0
Public care	—	—	—	—	—	—	—	—	—

Each cell shows the interaction: purchase propensity with both channels minus the additive prediction from individual effects. Negative values indicate super-additive demand reduction (channels reinforce each other). Pre-ann. denotes pre-existing annuitization through Social Security and the employer-pension-income proxy.

represent an income addition.

The predecessor-averaged pair interaction uses the same random-order logic as the Shapley values. For the demand-drop game $g(S) = v(\varnothing) - v(S)$ and $J = 9$, define

$$\mathcal{I}_{AB} = \sum_{S \subseteq \mathcal{N} \setminus \{A,B\}} \frac{|S|!\,(J - |S| - 2)!}{(J-1)!} \left[g(S \cup \{A,B\}) - g(S \cup \{A\}) - g(S \cup \{B\}) + g(S) \right]. \tag{F.1}$$

The weights sum to one over the 2^{J-2} predecessor sets. Positive values indicate that the pair suppresses demand more strongly together on average. Bequests and pricing frictions have $\mathcal{I}_{AB} = 3.4$ percentage points for participation and 5.3 points when v is the mean purchased share. Their empty-coalition entry in Table F.1 is 0.0 because the benchmark, either isolated player, and the two-player coalition all remain at the participation ceiling. Section 4.2 discusses the economic interpretation.

G Why No Scalar Factorization Is Used

The smooth problem has a useful local diagnostic. With no fixed purchase cost, minimum-premium discontinuity, or discrete annuitization grid, shifting an initial share α of wealth from the liquid asset to an annuity with payout rate ρ gives

$$\left. \frac{dV}{d\alpha} \right|_{\alpha=0} = W_0 \{ \rho V_A(W_0, 0, H, 1) - V_W(W_0, 0, H, 1) \}, \tag{G.1}$$

where V_W and V_A are marginal values of liquid wealth and annuity income. For the age-65 real-product version of this diagnostic, define $S_1(0) = 1$, $S_1(k) = \prod_{j=1}^{k} s_{\text{base}}(j)$, and

$$a_1^R(r) = \sum_{k=0}^{T-1} \frac{S_1(k)}{(1+r)^k}, \qquad \rho = \frac{\text{MWR}}{a_1^R(r)}.$$

The focal nominal product uses the age-specific nominal factor in Section A.11, and this real-product expression is only a local mnemonic. Define

$$R \equiv \rho \frac{V_A}{V_W} = \text{MWR} \frac{V_A}{a_1^R(r) V_W}. \tag{G.2}$$

The condition $R > 1$ favors an infinitesimal purchase from zero in this smooth problem. In the solved model, the fixed cost, minimum premium, and discrete α grid make participation a global value comparison, so $R > 1$ is not an if-and-only-if participation rule. Nor does the Yaari limiting case imply $R = 1$ at the no-annuity state. With no bequest motive or uninsured background risk and with fair pricing, mortality credits make the annuity weakly dominate the bond for financing consumption conditional on survival, so the household moves toward the annuitization boundary.

The solved model does not factor R into channel-specific scalars. Medical risk, survival, the public floor, resource form, bequests, and product terms change state-contingent marginal values. A channel may suppress demand, support it, be null, or change sign across coalitions and wealth cells. In particular, the focal age-needs player is null because its calibrated parameter is on the zero boundary.

For the same reason, the model does not impose super-additivity or any common interaction sign. Pairwise finite differences in Table F.1 are descriptive projections of the computed game, and higher-order interactions remain embedded in the coalition values. The exact enumeration and Shapley calculation, given the stated player definitions, provide the accounting, and no scalar mnemonic is used to rank channels.

H Joint Parameter Stress Test: Draw Design

The joint stress test evaluates 1,000 fixed-seed design points generated from independent uniform ranges: the Poor-health hazard multiplier $\mu_P \sim U(2.0, 3.5)$, inflation $\pi \sim U(0.015, 0.025)$, MWR $\sim U(0.78, 0.90)$, survival beliefs $\psi \sim U(0.96, 1.00)$, and the nonnegative consumption-needs slope $\delta_c \sim U(0, 0.015)$. Risk aversion is held fixed, and the two exploratory behavioral channels are off. Each point replaces the focal age-varying hazard matrix with the constant vector $[0.49, 1, \mu_P]$ at every age. The design therefore stresses constant hazard ratios and does

not nest the full focal schedule, whose Poor-health multiplier varies with age. Other ranges contain their corresponding focal scalar values. All ranges are researcher-specified. Every reported design point is solved on the nested H2 $195 \times 77 \times 101$ grid. Before those solves, a deterministic 13-point center-and-endpoint design compares H2 with the next nested core refinement, H3, and requires the maximum absolute difference to remain within the frozen 1.00 percentage-point purchase-propensity budget and 0.005 mean-annuitization budget. The realized maximum differences are 0.4747 percentage points and 0.002110, respectively. That comparison covers only the tested points, so it gives no uniform error bound inside the design box, and H3 is itself a finer approximation.

Across the 1,000 H2 design points, median purchase propensity is 26.2%, the central 90% runs from 2.7% to 43.0%, and the full empirical range is 0.3% to 50.6%. Table H.1 reports these empirical design outcomes. Their dispersion reflects the researcher-specified calibration ranges.

I Welfare CEV and the Consumption Floor

The welfare tables define subjective ex ante CEV as a proportional scaling of every realized unit of effective consumption in the fixed no-access allocation, including consumption financed by the public floor. Scaling floor-financed consumption is a unit convention and does not describe a policy that raises the floor or changes transfer eligibility. The grids are evaluated at age 65, and the bequest columns correspond to no bequest, the focal DFJ mapping, and a strong-bequest sensitivity that raises θ from 158.68 to 200 with κ held at \$396,496. Under the nine-player welfare specification, the exploratory behavioral channels are off: $\lambda_W = 1$ and $\psi_{\text{purchase}} = 0$. For $\gamma \neq 1$, consumption felicity, including the public-care transformation in a medical-on, floor-binding Poor-health state, is therefore homogeneous of degree $1 - \gamma$. This homogeneity and the closed form below are conditional on the zero-intercept state-dependent-utility convention in Section A.7. A nonzero additive state term would have to be separated from scalable consumption felicity, as the bequest term is.

Let V^1 denote value with annuity-market access, V^0 value without access, and B^0 the additively separable bequest component of the fixed no-access allocation. Thus $U_c^0 \equiv V^0 - B^0$ is its consumption-felicity component. Bequeathed wealth is a stock and is not scaled. Homogeneity makes the indifference condition

$$(1 + \lambda)^{1-\gamma} U_c^0 + B^0 = V^1,$$

Table H.1: Fixed-Seed Joint Calibration Stress Design at $\gamma = 2.5$

Statistic	Value
Number of draws	1000
Median purchase propensity	26.2%
Central 90% of design outcomes	[2.7%, 43.0%]
Interquartile range of design outcomes	[16.2%, 33.2%]
Mean	24.8%
Min / Max	0.3% / 50.6%
Fraction in [1%, 10%]	12%
Fraction in [3%, 6%] low-ownership reference band	4%

Risk aversion fixed at $\gamma = 2.5$. Fixed-seed design points vary five researcher-specified rational/preference stress-test parameters: $\mu_P \sim U(2.0, 3.5)$, $\pi \sim U(0.015, 0.025)$, MWR $\sim U(0.78, 0.90)$, $\psi \sim U(0.96, 1.00)$, $\delta_c \sim U(0.00, 0.015)$. Behavioral channels (SDU λ_w, PED ψ_{purchase}) are held off ($\lambda_w = 1$, $\psi_{\text{purchase}} = 0$), so the summaries reflect only the stated rational and preference sensitivity design. All draws use the 195×77×101 piecewise-cosine-v1 H2 nested-refinement grid (core 163×61 through the 3-million-dollar wealth splice), wealth support from zero through 13 million dollars, and probability-legendre-floor-split integration at order 32 per probability segment. A predeclared 13-point stress audit compares it with the 357×137×101 piecewise-cosine-v1 H3 core-refined reference grid (core 325×121) on the same support, annuitization-fraction grid, and quadrature conventions. The citable audit's maximum differences are 0.4747 percentage points in purchase propensity and 0.002110 in mean annuitization, against frozen limits of 1.00 and 0.005. The 3–6% row is an illustrative low-ownership band, not an interval between the distinct q286 and any-annuity-income HRS stock comparators.

which gives

$$\lambda = \left(\frac{V^1 - B^0}{V^0 - B^0} \right)^{1/(1-\gamma)} - 1. \tag{I.1}$$

When $\theta = 0$, $B^0 = 0$ and equation (I.1) reduces to the value-ratio expression. With bequests, backward induction propagates B^0 separately. If the ratio in equation (I.1) falls outside its positive domain, the run stops instead of reverting to a total-value approximation. Table I.1 reports the wealth–health grid. The population summary in the main paper uses the same minimum nonhousing-wealth threshold, $W \geq \$5{,}000$, as the ownership analysis. A citable run requires every eligible initial state to lie within the solution domain and age support, so no observation is assigned an artificial zero CEV because its state is unsupported.

The replication code also evaluates an intentionally conservative alternative convention in which all consumption in medical-on, public-floor-binding periods is held fixed while only consumption felicity in non-binding periods is scaled. It does not attempt to separate CRRA utility within a period into household- and transfer-financed dollars. For each listed wealth–health cell, simulated no-access paths estimate the binding-period share of discounted consumption felicity under the same subjective-survival measure as the Bellman value, and an algebraic root calculation reports the resulting alternative equivalent. The table's "floor periods" column is the simulated frequency of binding periods; the algebraic root instead uses the discounted floor-felicity share, which is retained separately in the CSV. All cells use 12,000 paths and fixed seed 20260713. The exercise is a diagnostic on the listed cells and does not re-solve the lifecycle problem under the alternative convention. Its reported column is solved on the piecewise production grid. The welfare grid in Table I.1 has the same size but single-power wealth spacing, so matching cells differ by less than 0.01 percentage point. Table I.2 reports every evaluated cell; the machine-readable results are in `tables/csv/cev_floor_scaling_diagnostic.csv`. The diagnostic shows how much the normative convention matters, and the welfare tables keep the focal convention.

J Secondary Counterfactuals and Distributional Diagnostics

These exhibits probe the extensive margin, the distribution of within-model access value, and alternative product assumptions. They do not enter the focal attribution and are not fitted to the HRS ownership stocks.

Table I.1: Consumption-Equivalent Variation by Household Type

Wealth	Health	No bequest	Moderate (DFJ)	Strong bequest
\$10,000	Good	0.00%	0.00%	0.00%
	Fair	0.00%	0.00%	0.00%
	Poor	0.00%	0.00%	0.00%
\$25,000	Good	0.00%	0.00%	0.00%
	Fair	0.00%	0.00%	0.00%
	Poor	0.00%	0.00%	0.00%
\$50,000	Good	0.00%	0.00%	0.00%
	Fair	0.00%	0.00%	0.00%
	Poor	0.00%	0.00%	0.00%
\$100,000	Good	0.00%	0.00%	0.00%
	Fair	0.00%	0.00%	0.00%
	Poor	0.00%	0.00%	0.00%
\$200,000	Good	1.35%	0.48%	0.33%
	Fair	0.83%	0.13%	0.00%
	Poor	0.09%	0.00%	0.00%
\$500,000	Good	5.80%	2.68%	2.15%
	Fair	5.07%	2.17%	1.64%
	Poor	3.87%	1.30%	0.87%
\$1,000,000	Good	10.07%	3.75%	2.87%
	Fair	9.39%	3.20%	2.41%
	Poor	8.12%	2.33%	1.64%

CEV: consumption-equivalent variation (%). Positive values indicate welfare gain from annuity market access. Nine-channel structural welfare specification with medical costs, health-mortality correlation, MWR = 0.87, inflation = 2%, $\gamma = 2.5$.
Welfare model band-dispatches Social Security plus the employer-pension-income proxy: one value function per wealth band at that band's combined flow (\$20,089/\$23,826/\$27,685/\$27,804/yr over breaks \$30k/\$120k/\$350k); each household is evaluated against the band its original wealth falls in.

Table I.2: Sensitivity of CEV to the Consumption-Floor Scaling Convention

Initial state	Floor periods (%)	Floor felicity (%)	Reported (%)	Alternative (%)	$\|\Delta\|$ (pp)
Band 1: \$8,000, Fair	10.07	31.39	0.000	0.000	0.0000
Band 1: \$8,000, Poor	12.02	39.24	0.000	0.000	0.0000
Band 1: \$25,000, Fair	8.93	29.66	0.000	0.000	0.0000
Band 1: \$25,000, Poor	10.49	35.91	0.000	0.000	0.0000
Band 2: \$50,000, Fair	5.44	24.89	0.000	0.000	0.0000
Band 2: \$50,000, Poor	6.37	29.75	0.000	0.000	0.0000
Band 2: \$100,000, Fair	4.34	22.58	0.000	0.000	0.0000
Band 2: \$100,000, Poor	4.95	26.08	0.000	0.000	0.0000
Band 3: \$200,000, Fair	2.23	17.47	0.140	0.169	0.0296
Band 3: \$200,000, Poor	2.52	19.86	0.000	0.000	0.0000
Band 4: \$500,000, Fair	0.96	12.24	2.164	2.476	0.3112
Band 4: \$500,000, Poor	1.05	13.45	1.290	1.495	0.2043

Each cell is dispatched by original wealth to its focal retirement-income band: Social Security is real and the pension-income proxy is nominal. "Floor periods" is the simulated binding-period frequency; "Floor felicity" is the discounted consumption-felicity share used by the algebraic root. The reported CEV proportionally scales all realized effective consumption in the fixed no-access allocation. The alternative holds consumption in every medical-on, public-floor-binding period fixed and scales only non-binding consumption felicity. Death draws use subjective survival. Each cell uses 12,000 paths and fixed seed 20260713. The exercise is an algebraic finite-cell diagnostic, not a re-solution or global error bound.

J.1 Extensive-margin and wealth-band diagnostics

Table J.1 separates minimum-premium feasibility, the fixed-cost margin, heterogeneous-cost smoothing, and dispersion in bequest preferences. It also repeats the Social Security cut under those treatments. The underlying indifference-cost diagnostic classifies nonbuyers as minimum-premium-infeasible, value-destroying even at zero fixed cost, or positive-value but defeated by the focal fixed cost. In the lowest band, where modeled purchase propensity is 0.0%, those shares are 35%, 65%, and zero. The fixed cost defeats a purchase that would otherwise add value for 30% of all households in band 2 and 20% in band 3. Each band is solved three times for the "theta-dispersed" row: with no bequest motive, with half the focal bequest intensity θ, and with the focal θ. In each solution, κ stays at its focal value and purchases follow that solution's own indifference cost. The three purchase rates are weighted, in that order, by the band's shares of model-eligible HRS person-waves whose self-reported chance of leaving an inheritance of \$100,000 or more is zero, between zero and 75 percent, or 75 percent or higher. Table J.2 then removes players one at a time from the full coalition to diagnose zero-fixed-cost purchase value in the interior wealth bands. These removals are not Shapley values. Table J.3 places the model's one-shot retail purchase propensity beside the q286 ownership stock. The comparison is descriptive because purchase timing and product source differ.

The threshold model's band propensities are 0.0%, 4.9%, 76.0%, and 100.0%, a much steeper rise across bands than the single-digit q286 stocks show in the same table. In particular, the modeled propensity jumps from 4.9% in the \$30–120k band to 76.0% in the \$120–350k band. The one-shot decision, fixed cost, minimum premium, discrete purchase

Table J.1: Social Security Cut Response: Robustness to the Extensive-Margin Treatment

			Induced response by wealth band (pp)			
Extensive-margin treatment	Base (%)	Cut (%)	$5–30k	$30–120k	$120–350k	>$350k
hard	37.7	43.2	+0.0	+8.5	+14.3	+0.0
feasibility-only	50.5	56.4	+0.5	+17.3	+4.0	+0.0
cost-smoothed	37.3	42.6	+0.0	+7.8	+13.8	+0.0
theta-dispersed	39.9	44.6	+0.0	+8.8	+10.5	+0.0

Each row models the evaluated-age participation margin differently; each derives from a per-household indifference fixed cost, which the theta-dispersed row recomputes under each bequest cell's preferences, so no band rate is a calibration target. "hard" is the production model; "feasibility-only" drops the fixed cost (minimum purchase and consumption floor remain); "cost-smoothed" applies heterogeneous fixed costs (fixed dispersion, not fitted); "theta-dispersed" applies the observed within-band bequest-intention distribution (HRS beq100), frozen before the cut. Entries are percentage-point changes in purchase propensity; the table is the robustness result, and no common sign or concentration pattern is imposed in its note.

Table J.2: Value-Destruction Diagnostic: Would-Annuitize Share at Zero Fixed Purchase Cost, by Wealth Band

Specification	Band 2 ($30–120k)	Band 3 ($120–350k)
Full structural	34.6%	96.0%
less Bequest	44.2%	100.0%
less Pre-existing annuitization	80.1%	100.0%
less Med+R-S	67.3%	100.0%
less Survival pessimism	66.9%	100.0%
no MWR wedge or minimum premium	85.0%	100.0%
less Public-care aversion	52.8%	100.0%
less Inflation	39.5%	97.1%
less Age needs	34.6%	96.0%
less State utility	34.4%	95.7%

Share of households in each band who would purchase the immediate annuity at zero *fixed purchase cost* (indifference fixed cost $F^\star > 0$), under the full structural model and under each one-channel-removed variant. The focal MWR wedge and minimum premium remain except in the explicitly labeled row, which removes both. The pre-existing-annuitization omission replaces real Social Security and the nominal pension-income proxy with respondent-age-specific equal-PV liquid wealth, with the original wealth band pinned. Bands 2 and 3 isolate the interior-band zero-fixed-cost purchase margin. The omission rows are ordered by their reported Band 3 share; these are grand-coalition removals, not Shapley values. Production grid.

Table J.3: Predicted Purchase Propensity and the q286 Ownership Stock by Wealth Band

Wealth band	Observed (HRS q286 stock)	Model (threshold)	Model (cost-smoothed)
\$5–30k	1.4%	0.0%	0.0%
\$30–120k	4.3%	4.9%	5.5%
\$120–350k	4.8%	76.0%	73.5%
Above \$350k	5.9%	100.0%	100.0%

Observed: unweighted q286 life-contingent-payment ownership stock in the HRS analysis sample (73 qualifying observations among 1896 pooled person-wave observations, from 1269 distinct persons of whom 63 report a qualifying payment in at least one eligible observation).
Model (threshold): the single-product structural model's predicted purchase propensity by band. Model (cost-smoothed): the same prediction under lognormal heterogeneous transaction costs with median at the CPI-restated Lockwood fixed cost (\$2,349) and log-standard-deviation 0.5 — the non-fitted dispersion of Section J.1; no band ownership rate enters its construction.
Inference is at the person level: each repeatedly observed person counts once, with wealth band assigned at the first eligible wave and the q286 flag set to one if a qualifying response appears in at least one eligible observation. This is not lifetime cumulative purchase incidence. A Cochran–Armitage trend test across the four bands gives $z = 4.97$, $p < 0.001$. The pooled person-level q286 share in the three non-top bands is 3.9% (Wilson 95% CI [2.86%, 5.23%]). The displayed person-wave and person-level rates use different repeated-observation conventions. At the person level, the third- and top-band rates are 6.6% and 9.7%, with the top-two-band difference not statistically distinguishable (Fisher exact $p = 0.25$).

grid, and banded retirement-income assignment all act at this margin, so the table does not attribute the jump to any one mechanism.

J.2 Price, indexation, belief, and access-value sensitivities

Table J.4 varies the modeled payout schedule, payment indexation, survival beliefs, and Social Security benefits in partial equilibrium. Its TIPS-backed row applies an assumed MWR of 0.78 to real payments, so price and indexation change together. The nominal-equivalent row keeps the MWR at 0.87 and changes indexation alone. Within-model access values for these scenarios appear in Table J.5, where the real-annuity column uses the nominal-equivalent contract. Fiscal costs and equilibrium responses remain outside the welfare measure.

Table J.4: Predicted Annuity Purchase Propensity Under Policy Counterfactuals

Policy Scenario	MWR	Inflation	Propensity (%)	Δ (pp)
Baseline	0.87	2%	37.7	—
Lower-price schedule (MWR=0.90)	0.90	2%	42.9	+5.2
Alternative price schedule (MWR=0.95)	0.95	2%	49.9	+12.2
MWR=1 payout schedule	1.00	2%	56.0	+18.3
Real annuity, TIPS-backed	0.78	0 (real)	24.5	-13.2
Real annuity, nominal-equiv	0.87	0 (real)	40.2	+2.5
MWR=1 + real payments	1.00	0 (real)	56.0	+18.3
SS cut 22%	0.87	2%	43.2	+5.5
Objective beliefs (psi=1.0)	0.87	2%	47.4	+9.7
Lower price + objective beliefs	0.90	2%	51.5	+13.8
Public consumption floor doubled	0.87	2%	42.7	+5.0
Combined price/indexation/beliefs	0.90	0 (real)	54.3	+16.6
HRS q286 ownership stock			3.85	
HRS any-annuity-income proxy			6.12	

Baseline: $\gamma = 2.5$, $\beta = 0.97$, DFJ bequests ($\theta = 158.68$, $\kappa = \$396{,}496$), $\psi = 0.981$. Population: model-eligible HRS person-waves ages 65–69 with $W \geq \$5{,}000$ ($N = 1{,}896$); see the manuscript for the labor-force and prior-year-earnings screens. The lower-price schedule is a hypothetical price counterfactual, not observed group access. SS cut: permanent, fully anticipated 22% reduction in Social Security benefits only; the nominal pension proxy, prices, claiming, legislation, and equilibrium responses are held fixed. This is a partial-equilibrium comparative static, not a transition forecast. The two HRS rows are distinct, nonnested stock comparators, not an interval.

Table J.5: Welfare Gain from Annuity Access Under Policy Counterfactuals (CEV, %)

Wealth	Baseline	Lower price	Real annuity	Combined scenario
Panel A: Good Health, DFJ Bequests				
\$50K	0.00	0.00	0.00	0.11
\$100K	0.00	0.08	0.00	1.38
\$200K	0.48	1.13	0.85	3.46
\$500K	2.68	4.00	3.45	8.36
\$1000K	3.75	5.35	4.76	11.13
Panel B: Fair Health, DFJ Bequests				
\$50K	0.00	0.00	0.00	0.00
\$100K	0.00	0.00	0.00	0.72
\$200K	0.13	0.60	0.39	2.55
\$500K	2.17	3.30	2.91	7.27
\$1000K	3.20	4.68	4.17	10.09
Panel C: Good Health, No Bequests				
\$50K	0.00	0.00	0.00	0.20
\$100K	0.00	0.42	0.17	1.58
\$200K	1.35	2.25	1.76	4.21
\$500K	5.80	7.55	6.64	10.82
\$1000K	10.07	12.64	11.32	16.74

CEV reports 100λ, where every effective-consumption outcome in the fixed no-access allocation is scaled by $1 + \lambda$ to equate its value to the with-access allocation. Welfare cells dispatch the Social Security and modeled nominal pension-proxy schedule by original wealth band. Baseline: MWR=0.87, 2% inflation, $\psi = 0.981$; behavioral channels off. Lower-price scenario: MWR=0.90. Real annuity: 0% inflation, MWR=0.87. The combined scenario uses the lower price, real payments, and objective survival beliefs ($\psi = 1.0$).

J.3 Illustrative lower-price mixture

The lower-price exercise solves the same contract at its stated alternative MWR and combines that prediction with the retail prediction using each wealth band's employer-pension receipt rate. Reported pension income in the HRS does not establish plan type, life contingency, or access to a new group annuity. Pension receipt is therefore only an illustrative mixture weight and does not measure or bound access. Table J.6 and Figure J.1 report the exercise.

Table J.6: Illustrative Two-Price Pension-Linkage Mixture

Panel A: Model propensities and HRS stock by band (%)					
Band	Mixture wt.	Retail prop.	Lower-price prop.	Mixture prop.	HRS stock
\$5–30k	46%	0.00	0.00	0.00	1.39
\$30–120k	55%	4.89	33.08	20.36	4.32
\$120–350k	63%	75.95	95.48	88.27	4.76
>\$350k	45%	100.00	100.00	100.00	5.95
Panel B: Model mean purchased share of wealth, α (%)					
Band		Retail	Lower-price	Mixture	
\$5–30k		0.00	0.00	0.00	
\$30–120k		2.56	24.43	14.56	
\$120–350k		38.17	69.32	57.82	
>\$350k		52.15	67.93	59.31	
Panel C: 22% Social Security cut (mixture, %)					
Band		Baseline	Under cut	Response (pp)	Abs. share
\$5–30k		0.00	0.00	+0.00	0%
\$30–120k		20.36	29.54	+9.18	59%
\$120–350k		88.27	96.25	+7.98	41%
>\$350k		100.00	100.00	+0.00	0%

The mixture weight is each band's RAND HRS employer-pension-income receipt rate. It is not observed access to a new group-annuity purchase: The HRS employer-pension-income measure does not identify plan type, life contingency, or a current purchase option. The exercise is therefore illustrative. Retail prices at MWR = 0.87; the lower-price contract at MWR = 0.95. The contracts otherwise are identical, so product choice is dominance-ordered. Observed columns are unweighted q286 life-contingent-payment indicators in the restricted person-wave sample. Panel C applies a permanent, anticipated 22% cut to Social Security only; it is a partial-equilibrium model comparative static, not a forecast. "Abs. share" allocates the total absolute population-weighted band response; signs remain in the response column, so offsetting changes cannot create negative or above-100% incidence shares.

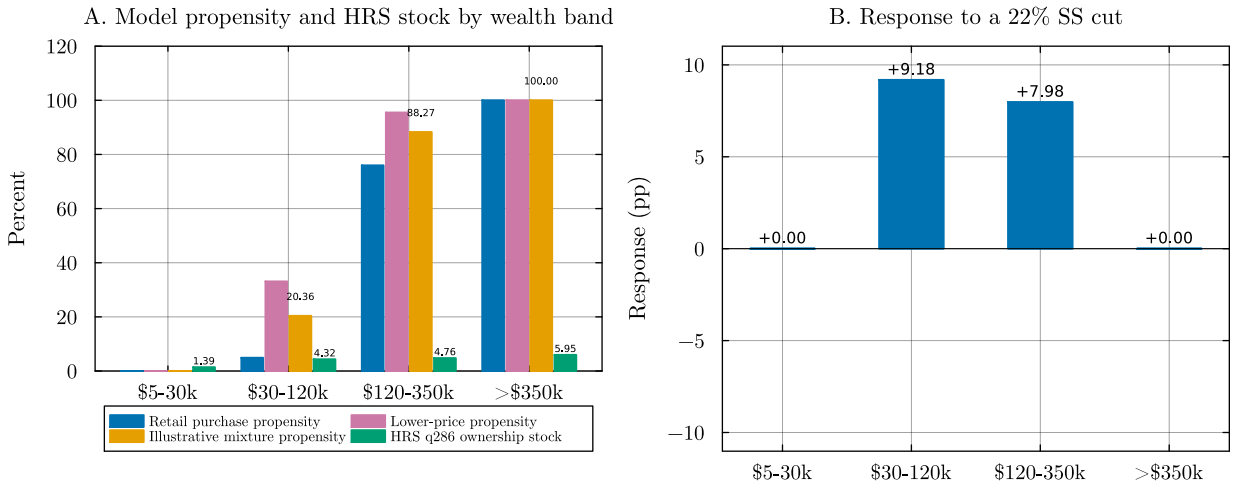


Figure J.1: Illustrative two-price pension-linkage mixture. Panel A compares the retail prediction, the assumed lower-price schedule, the illustrative mixture, and the q286 ownership stock. RAND employer-pension-income receipt rates are used only as illustrative mixture weights and do not measure current group-annuity access. Panel B reports the mixture's partial-equilibrium response to a permanent, anticipated 22% Social Security cut.

K Solution Algorithm

K.1 Backward induction

The model is solved by backward induction from the terminal period $T = 46$ (age 110) to $t = 1$ (age 65). Terminal medical need remains stochastic. Define total post-medical resources, inclusive of the public consumption floor, as

$$R_T(m) = \max\{\underline{c},\ W + A_T + Y_T^{\text{pre}} - m\}\,,$$

where Y_T^{pre} is real Social Security plus the modeled real value of the nominal pension-income proxy. Let $\widetilde{u}_T(c, H, m)$ denote the implemented terminal flow utility: it includes the age and health felicity weights, the source-dependent split between income and portfolio drawdown when that exploratory channel is active, and the public-care consumption transformation when the floor binds in Poor health. The code-consistent terminal condition is

$$V(W, A, H, T) = \mathbb{E}_m\left[\max_{\underline{c}\le c\le R_T(m)} \{\widetilde{u}_T(c, H, m) + \beta v(R_T(m) - c)\}\right].$$

Medical expenditure first absorbs the period's income flows and then liquid wealth, and that source accounting affects $\widetilde{u}_T$ only when source-dependent utility is active. Absent medical risk, the expectation degenerates. Without a bequest motive, the terminal household

consumes all $R_T(m)$.

For $t = T-1, \ldots, 1$, the value function was computed at each grid point (W_i, A_j, H_k). After the medical draw, resources include private-annuity income A_t and pre-existing annuitized income Y_t^{pre}. If post-medical resources fall below $\underline{c}$, the public-floor branch consumes $\underline{c}$ and saves zero. Otherwise the search considered consumption from $\underline{c}$ to post-medical resources, subject to the certified continuation domain. The continuation value $V(W', A', H', t+1)$ was evaluated by piecewise-linear interpolation over the wealth grid.

The search partitioned this feasible interval at the consumption levels corresponding to continuation-wealth knots and, when source-dependent utility was active, the income-financed-consumption threshold. Within each smooth segment, the objective is concave and its derivative decreases. The solver evaluated every boundary and used derivative bisection when the derivative changed from positive to negative across a segment. It then selected the candidate with the greatest represented Bellman value. CPU and CUDA implementations used the same scalar search in double precision, without fast-math transformations. This procedure globally maximizes the represented objective up to Float64 arithmetic and the root-bracket tolerance, but it does not eliminate interpolation or quadrature error. A Brent search in the code serves only the regression tests and produced no reported value.

Production solves constructed the admissible domain backward. At each annuity–health–age state, the search was restricted to the largest contiguous wealth prefix whose positive-probability next-health interpolation stencils were already certified at $t + 1$. A current node was certified only when every positive-mass medical node produced an optimal W' strictly inside that clean continuation domain, using an optimizer-scale boundary tolerance. Flat extrapolation remains an interpolation-library convention, but no certified production optimization was permitted to query or select it.

K.2 Expectation computation

The Bellman expectation combines:

1. An outer expectation over the current medical draw m_t, observed before the consumption choice, evaluated with state-split probability-domain Gauss–Legendre sums at order 32 on each nonempty accounting subinterval plus an exact saturated-tail mass.

2. Conditional on survival, an inner sum over next-period health H' weighted by $Q_t(H_t, H')$.

The number of medical-shock nodes is state dependent: each nonempty probability segment receives an order-32 Gauss–Legendre rule, while the saturated upper tail enters as one exact mass. Consequently there is no valid fixed nodes-per-grid-point arithmetic. The transition

diagnostic records the realized node counts and probability mass; the run validator checks positive, internally consistent counts and that the recorded probability mass aggregates to the state count, while unit mass for each individual state is enforced inside the solver itself.

The pre-draw Bellman value integrates shock-contingent consumption optima over the current medical draw and, conditional on survival, next-period health; the object carried backward is the value function $V(\cdot, t+1)$. Forward simulations are shock-contingent: at each simulated period the consumption choice is re-optimized against the realized health state and medical-expenditure draw using the stored continuation value, rather than applying the quadrature-weighted average consumption stored for diagnostics. This matters because the medical-expenditure distribution is heavy-tailed and interacts with the public consumption floor, so applying an averaged policy to realized tail draws would misstate spend-down behavior near the floor.

K.3 Annuitization optimization

At purchase period τ, the global search evaluates zero, every feasible node of the 101-point uniform grid on $[0, 1]$, and the exact minimum-premium and maximum-feasible-share endpoints whenever those endpoints are feasible. Let $P_{\min}$ denote the minimum premium, set to \$10,000. Positive candidates satisfy $\alpha W_0 \geq P_{\min}$ and $(1-\alpha)W_0 - F \geq 0$. Each respondent is priced at the observed age. For a feasible candidate, post-purchase wealth is $(1-\alpha)W_0 - F$ and the fixed nominal-payment state increases by $\alpha W_0(1+\pi)^{\tau-1}\rho_\tau$, where ρ_τ is the age-specific first-year real payout rate, and F is omitted for $\alpha = 0$. This gross-up makes the real payment in period τ equal to $\alpha W_0 \rho_\tau$. Bilinear interpolation over the (W, A) grid evaluates the value function at the resulting state for the relevant health state and period.

K.4 Population evaluation

The purchase propensity is evaluated on 1,896 eligible HRS person-waves with observed initial wealth, age, and health. Retirement income is assigned at the wealth-band level, while the equal-PV off-state endowment is priced at each respondent's age. Purchase is $\alpha^* > 0$. For each solve, the code first constructs the exact set of initial and purchase-state interpolation stencils reachable by the population under the configured age-specific pricing, annuitization candidates, equal-PV endowments, and any fixed-cost wedge. Every required stencil must lie in the backward-certified domain described above; the same certificate is checked again by welfare and forward-simulation consumers. The disclosure-safe population-grid aggregate independently verifies the corresponding exogenous support. A supplemental endogenous-boundary artifact summarizes realized production transitions, but certification

does not rely on a rectangular $W_{\max} \times A_{\max}$ corner that no household can jointly reach.

K.5 Computational performance

The reported GPU computation used one NVIDIA H100. Its CUDA kernels batched the shock-contingent consumption searches, with one scalar optimization assigned to each GPU thread. Medical-quadrature construction, continuation-value assembly, probability-weighted aggregation, and support certification remained on the CPU. Both backends preserved the same numerical rules and summation order. Runtime depends on grid size, the number of state-dependent probability segments, and reuse of solved value functions; stage timings are retained with the run record. All focal solves used the common order-32 state-split rule and the grid reported above.

Numerical-diagnostic scope. Independent scalar tests compared the interval optimizer with separately implemented piecewise and high-precision references, including separated local maxima and source-income kinks. CPU–CUDA comparisons checked value functions, support certificates, and purchase decisions on public test populations. Numerical approximation was assessed separately through wealth–annuity and purchase-grid refinement, the H2-versus-H3 Monte Carlo-parameter stress design, and the H0 order-32-to-48 and H2 order-48-to-64 medical-integration checks. Limiting-case tests, the backward clean-domain certificate, and the endogenous next-period-wealth audit supplied additional checks. No single Euler-residual scalar is reported. The consumption problem has corner KKT inequalities and a piecewise-linear continuation value with kinks, so a smooth finite-difference residual has no uniform interpretation across states. These checks do not give a uniform error bound for the continuous-state model. Exact Shapley enumeration establishes accounting conditional on the computed coalition values and supplies no separate lifecycle-solver validation.

L Pricing-Frictions-Split Shapley Decomposition

The nine-player game bundles the money's-worth wedge, fixed purchase cost, and minimum-purchase requirement. Table L.1 reruns an eleven-player game with the components separated; its full coalition reproduces the focal configuration (37.7%). Pre-existing annuitization leads at 16.2 pp, followed by the MWR wedge (14.9 pp), fixed cost (9.7 pp), and medical/mortality (9.2 pp). The minimum purchase contributes 1.3 pp and bequests 2.5 pp. Shapley values depend on the player partition, so the pricing components need not sum to the bundled attribution.

Table L.1: Pricing-Frictions-Split Shapley Decomposition (11 Players, 2,048 Subsets)

Channel	Shapley (pp)
Pre-existing annuitization (SS+pension-income proxy)	+16.2
MWR wedge	+14.9
Fixed purchase cost	+9.7
Medical process and health-dependent mortality	+9.2
Subjective survival pessimism	+7.0
Bequests	+2.5
Minimum-purchase requirement	+1.3
Nominal-payment erosion	+1.1
Public-care aversion	+0.2
State-dependent utility	+0.1
Age-varying consumption needs	+0.0

Exact Shapley values over all $2^{11} = 2{,}048$ subsets of the game that replaces the bundled pricing-frictions channel with its three components as separate players: the money's-worth wedge (MWR = baseline vs. 1.0; the only component that changes the payout rate), the fixed purchase cost, and the minimum-purchase requirement. All other channels and the production grid match the nine-channel headline game; the full coalition reproduces the nine-channel full model. Shapley values are not additive across game partitions, so the three components need not sum to the bundled channel's value. Positive values are demand-suppressing.

M Mortality-Convention Robustness

The focal calibration uses the sex-blended, aggregate-anchored life table (Section 3). Table M.1 recomputes the full nine-channel game under the prior convention in this literature: the male SSA 2003 table from the Lockwood (2012) replication files, with unnormalized health hazards, used for both survival and pricing. Purchase propensity becomes 21.9%. Pricing frictions remain first at 26.5 pp, medical/mortality rises to second at 23.9 pp, and pre-existing annuitization moves to third at 17.2 pp. Bequests contribute 4.4 pp. This joint convention changes both the population table and health-hazard normalization.

Table M.1: Nine-Channel Shapley Decomposition under the Prior Mortality Convention (Male Table, Unnormalized Hazards)

Channel	Shapley (pp)
Loads	+26.5
Medical+R-S	+23.9
Pre-existing annuitization (SS+pension-income proxy)	+17.2
Pessimism	+6.4
Bequests	+4.4
Inflation	+0.1
State utility	+0.1
Age needs	+0.0
Public-care aversion	-0.5

Exact Shapley values over all $2^9 = 512$ subsets under the prior mortality convention: the male SSA 2003 table used by Lockwood (2012), with unnormalized health hazards, for both agent survival and annuity pricing. The headline game instead uses the sex-blended, aggregate-anchored table. All other channels and the production grid match the headline game. Rows are ordered by signed Shapley value; positive values are demand-suppressing and negative values are demand-raising.

N Constant Hazard-Ratio Sensitivity

The focal model interpolates the HRS health-hazard ratios across age because mortality after the purchase age enters every continuation value. Table N.1 recomputes the complete game under the alternative convention that holds the entry-age ratio vector fixed through age 110. This sensitivity changes the mortality correlation while preserving the focal sex-blended population survival target through age-specific normalization.

Table N.1: Nine-Channel Shapley Decomposition with Constant Entry-Age Health Hazards

Channel	Shapley (pp)
Pricing frictions	+28.3
Pre-existing annuitization (SS+pension-income proxy)	+15.1
Medical+R-S	+8.7
Pessimism	+6.3
Bequests	+2.8
Inflation	+1.3
Public-care aversion	+0.8
Age needs	+0.0
State utility	-0.0

Exact values over all $2^9 = 512$ coalitions. This sensitivity holds the HRS entry-age hazard-ratio row fixed through age 110; the focal model instead interpolates all three age-band estimates. Equal-resource and full-model propensities are 100.0% and 36.8%. All other primitives and the production grid are unchanged.

O Additional Tables and Figures

This section collects the lifecycle moment diagnostic (Table O.1) and supporting sensitivity figures. Figure O.1 shows the risk-aversion path, Figure O.2 provides the health-hazard sensitivity discussed in the main text, Figure O.3 displays selected purchase policies, and Figure O.4 displays the no-bequest access-value grid.

P HRS Variable Codes

Evaluation sample. Respondent weights are RAND `RwWTRESP`. "Single" requires RAND `RwMSTAT` code 4 (separated), 5 (divorced), 6 (separated/divorced), 7 (widowed), or 8 (never married), together with `RwMPART`=0. Code 3 (partnered) is excluded. For the earnings screen, both each wave's earnings and Lockwood's 2005-dollar reference cutoff are restated with prior-calendar-year CPI-U. Baseline wealth is RAND total nonhousing wealth, `h{w}atotn`, and focal Social Security is built from RAND retirement-benefit income, `RwISRET`. The pension-income proxy uses RAND employer-pension income, `RwIPEN`, which refers to the prior calendar year and excludes annuity income. In Table D.6, the narrow liquid-wealth row uses RAND non-housing financial wealth, `h{w}atotf`: stocks and mutual funds; checking, savings, and money market accounts; CDs, government savings bonds, and Treasury bills; bonds and bond funds; and other savings, net of other debt. That measure excludes IRAs, real estate, vehicles, and businesses. The financial+IRA row adds IRA and Keogh

Table O.1: Simulated Lifecycle Bequest Diagnostics

Moment	Simulated
Mean bequest	\$125627
Median bequest	\$19143
Fraction bequest > \$10K	58.1%

Simulated: 100,000 trajectories, initial wealth and health drawn jointly from the weighted model-eligible HRS sample (single respondents ages 65–69 who pass the labor-force and prior-year-earnings screens and have wealth of at least \$5,000), all initialized at age 65, with band-level Social Security plus the employer-pension-income proxy; no annuity purchased. These are untargeted model implications, not an empirical validation target. Reported-estate measures are not used because their asset definitions and imputation conventions do not match the model wealth state.

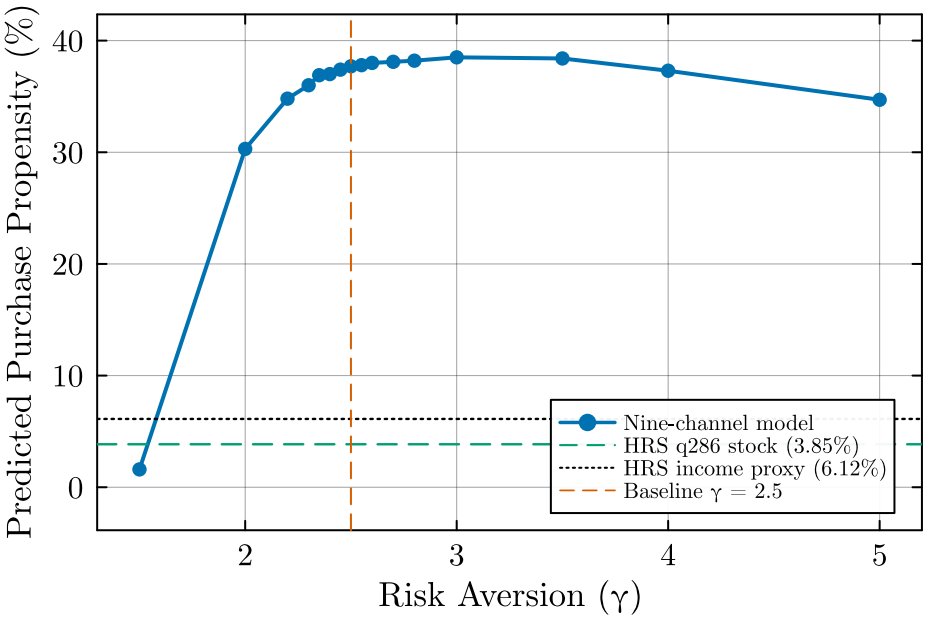


Figure O.1: Predicted purchase propensity by coefficient of relative risk aversion (γ). Inflation and MWR remain at baseline values ($\pi = 2\%$, MWR $= 0.87$); the linked DFJ intensity is recomputed at each γ by the portability rule in Section C.1. Separate horizontal lines show the untargeted HRS q286 ownership stock and any-annuity-income proxy; they are distinct measures and do not bound a model-fit interval.

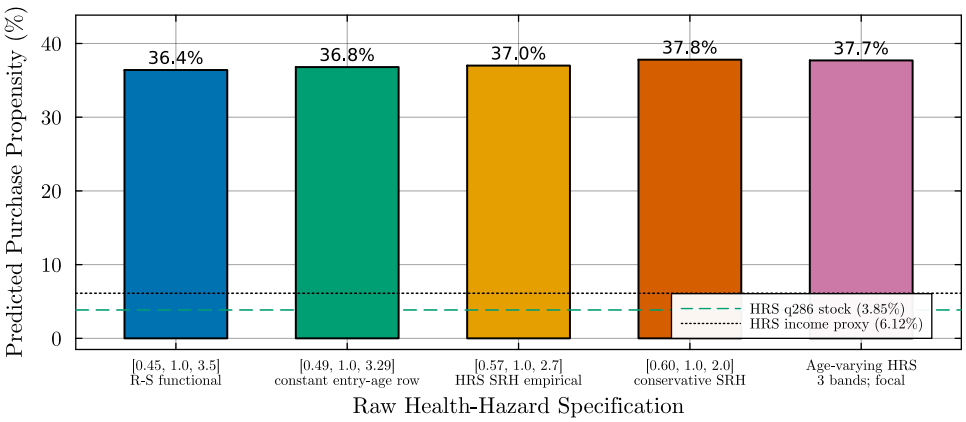


Figure O.2: Predicted purchase propensity by raw health-hazard specification. Constant-ratio rows vary the Good- and Poor-health multipliers while normalizing Fair to 1.0; the focal row interpolates three age-band vectors with the same raw normalization. The age-specific common anchor described in Section A.3 then restores population survival. In the axis labels, the $[0.45, 1.0, 3.5]$ vector ("R-S functional") and the $[0.60, 1.0, 2.0]$ vector ("conservative SRH") are researcher-specified stress mappings, and the other three come from HRS estimates. Separate horizontal lines show the untargeted HRS q286 ownership stock and any-annuity-income proxy; they are distinct measures and do not bound a model-fit interval.

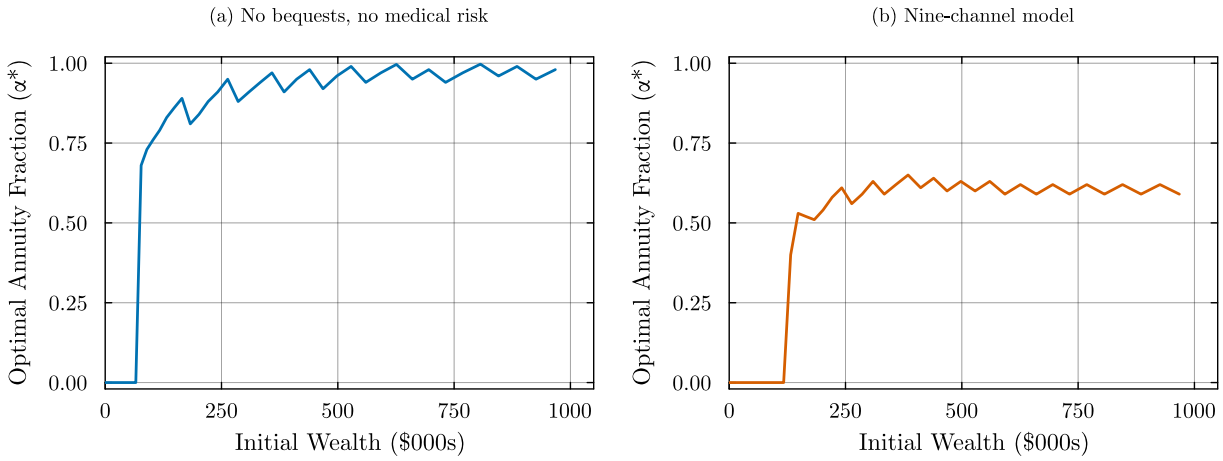


Figure O.3: Optimal annuitization fraction (α^*) at age 65 by initial wealth, for Good health ($H = 1$). Left panel: no bequest motive and zero medical expenditure, with pricing frictions, inflation, and survival pessimism active. Right panel: full nine-channel structural model (DFJ bequests, the medical-expenditure/health-mortality player, survival pessimism, the focal null age-needs player, state-dependent utility, public-care aversion, pricing frictions, and inflation). Both panels use MWR $= 0.87$, $\pi = 2\%$, and a uniform flow equal to the sample mean of real Social Security plus the modeled nominal employer-pension-income proxy.



Figure O.4: Subjective ex ante consumption-equivalent variation (%) from annuity market access at age 65, by initial wealth and health status. No-bequest specification ($\theta = 0$). Nine-channel structural welfare specification with medical costs, health-dependent mortality, MWR = 0.87, and 2% inflation.

balances, `h{w}aira`, and the liquid-band row assigns each observation's retirement-income band by `h{w}atotf` while keeping nonhousing net worth as model wealth.

Mortality and life-table inputs. For the hazard multipliers, next-wave RAND interview-status code 5 identifies death since the preceding interview, codes 1 and 4 identify survival, and indeterminate interval outcomes are excluded. The life-table blend uses RAND's binary `ragender` field, with code 2 mapped to the female schedule and code 1 to the male schedule. Missing or other codes fail validation.

Ownership measures. The q286 series comes from the HRS annuity-income loops in waves 5–9. Wave 5 uses Fat-File items `g5494_1` and `g5494_2`, and waves 6–9 use the two wave-prefixed q286 contract slots. The any-annuity-income proxy is RAND `r{w}iann`. For both measures' descriptive Wilson-type intervals, the effective sample size is

$$n_{\text{eff}} = \frac{\left(\sum_i m_i\right)^2}{\sum_i m_i^2},$$

where m_i is each person's number of eligible observations, under the conservative assumption of within-person correlation one.

References

Aguiar, M., Hurst, E., 2013. Deconstructing life cycle expenditure. Journal of Political Economy 121, 437–492. doi: 10.1086/670740.

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M.D., Tonetti, C., 2020. Long-term-care utility and late-in-life saving. Journal of Political Economy 128, 2375–2451. doi: 10.1086/706686.

Barberis, N., Huang, M., 2009. Preferences with frames: A new utility specification that allows for the framing of risks. Journal of Economic Dynamics and Control 33, 1555–1576. doi: 10.1016/j.jedc.2009.01.009.

Blanchett, D., Finke, M.S., 2024. Guaranteed Income: A License to Spend. Original Research 028-2024. Retirement Income Institute. Washington, DC. URL: https://www.limraconsumer.com/wp-content/uploads/2024/06/RP-28_BlanchettFinke_v2.pdf.

Blanchett, D., Finke, M.S., 2025. Retirees spend lifetime income, not savings. Financial Planning Review 8, e70010. doi: 10.1002/cfp2.70010.

Brown, J.R., Kling, J.R., Mullainathan, S., Wrobel, M.V., 2008. Why don't people insure late-life consumption? A framing explanation of the under-annuitization puzzle. American Economic Review: Papers and Proceedings 98, 304–309. doi: 10.1257/aer.98.2.304.

Chalmers, J., Reuter, J., 2012. How do retirees value life annuities? evidence from public employees. Review of Financial Studies 25, 2601–2634. doi: 10.1093/rfs/hhs057.

De Nardi, M., French, E., Jones, J.B., 2010. Why do the elderly save? The role of medical expenses. Journal of Political Economy 118, 39–75. doi: 10.1086/651674.

Finkelstein, A., Luttmer, E.F.P., Notowidigdo, M.J., 2013. What good is wealth without health? The effect of health on the marginal utility of consumption. Journal of the European Economic Association 11, 221–258. doi: 10.1111/j.1542-4774.2012.01101.x.

Health and Retirement Study, 2025. RAND HRS Longitudinal File 2022 (V1) public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant numbers NIA U01AG009740 and NIA R01AG073289). Ann Arbor, MI. URL: https://hrsdata.isr.umich.edu/data-products/rand-hrs-longitudinal-file-2022-v1.

Heimer, R.Z., Myrseth, K.O.R., Schoenle, R.S., 2019. YOLO: Mortality beliefs and household finance puzzles. The Journal of Finance 74, 2957–2996. doi: 10.1111/jofi.12828.

Hu, W.Y., Scott, J.S., 2007. Behavioral obstacles in the annuity market. Financial Analysts Journal 63, 71–82. doi: 10.2469/faj.v63.n6.4928.

Jones, J.B., De Nardi, M., French, E., McGee, R., Kirschner, J., 2018. The Lifetime Medical Spending of Retirees. Working Paper 24599. National Bureau of Economic Research. doi: 10.3386/w24599.

Lockwood, L.M., 2012. Bequest motives and the annuity puzzle. Review of Economic Dynamics 15, 226–243. doi: 10.1016/j.red.2011.03.001.

Mitchell, O.S., Poterba, J.M., Warshawsky, M.J., Brown, J.R., 1999. New evidence on the money's worth of individual annuities. American Economic Review 89, 1299–1318. doi: 10.1257/aer.89.5.1299.

O'Dea, C., Sturrock, D., 2023. Survival pessimism and the demand for annuities. The Review of Economics and Statistics 105, 442–457. doi: 10.1162/rest_a_01048.

Reichling, F., Smetters, K., 2015. Optimal annuitization with stochastic mortality and correlated medical costs. American Economic Review 105, 3273–3320. doi: 10.1257/aer.20131584.

Social Security Administration, 2007. Period life table, 2003. Annual Statistical Supplement, Table 4.C6. URL: https://www.ssa.gov/policy/docs/statcomps/supplement/2007/4c.html. accessed July 18, 2026.

U.S. Bureau of Labor Statistics, 2017. Historical consumer price index for all urban consumers (CPI-U): U.S. city average, all items. CPI Detailed Report, Table 24. URL: https://www.bls.gov/cpi/tables/historical-cpi-u-201707.pdf. accessed July 18, 2026.

Wettstein, G., Munnell, A.H., Hou, W., Gok, N., 2021. The Value of Annuities. Working Paper CRR WP 2021-5. Center for Retirement Research at Boston College. URL: https://crr.bc.edu/wp-content/uploads/2021/02/wp_2021-5.pdf.

Yaari, M.E., 1965. Uncertain lifetime, life insurance, and the theory of the consumer. The Review of Economic Studies 32, 137–150. doi: 10.2307/2296058.