Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle*

Victor Duarte
University of Illinois

Julia Fonseca
University of Illinois
and NBER

Aaron Goodman
Vanguard

Jonathan A. Parker
MIT and NBER

September 2026

Abstract

In many areas of economics, relatively simple models developed for insight are used as quantitative guides to behavior. In this paper, we develop a machine-learning algorithm to solve for optimal portfolio choice in a complex, calibrated lifecycle model that includes many features of reality modeled only separately in previous work. The average optimal portfolio share invested in stocks declines with age, consistent both with age-dependent rules derived from simpler models and with implicit advice embedded in current lifecycle investment products. But optimal portfolios decline more evenly, slowly, and by less than in current lifecycle products, and have substantial heterogeneity, particularly by wealth level and expected equity premium, implying significant gains to further customization in these dimensions. Together these two factors imply gains from moving to optimal portfolios of up to 3.5 to 3.7 percent of average consumption.

---

*First Draft December 2021. We thank Gaon Kim and Xin Xiong for excellent research assistance. We also thank our discussants James Choi, David Laibson, Peter Maxted, and Alex Michaelides, as well as Hui Chen, Matthias Doepke, Eric French, Francisco Gomes, Fatih Guvenen, Gaon Kim, Rory McGee, Maarten Meeuwis, Mariacristina De Nardi, Lawrence Schmidt, Chester Spatt, Michele Tertilt, Motohiro Yogo and Yao Zeng for useful discussions on model components. We also thank participants at the 2021 NBER Summer Institute Big Data and High-Performance Computing for Financial Economics Conference, the 2021 European Winter Meetings of the Econometric Society, the 2022 NBER Summer Institute Household Finance Meeting, the April 2023 Micro and Macro Implications of Household Behavior and Financial Decision-Making Virtual Seminar Series, the 2023 NBER Summer Institute Asset Pricing Meeting, the 2025 SED Meetings, and seminars at Amsterdam, Bocconi, Imperial, INSEAD, Northwestern, MIT, McGill, the Swedish House of Finance, and Vanguard. The opinions expressed here are those of the authors and do not represent the views of Vanguard.

Duarte: Gies College of Business, 1206 South 6th Street, Champaign IL 61820, vduarte@illinois.edu; Fonseca: Gies College of Business, 1206 South 6th Street, Champaign IL 61820, juliaf@illinois.edu; Goodman: Investment Strategy Group, The Vanguard Group, Inc. PO Box 2600, Valley Forge PA 19482, protectaaron.goodman.15@gmail.com; Parker: Sloan School of Management, MIT, 100 Main Street, Cambridge, MA 02142, JAParker@mit.edu.

Economics uses parsimonious models to gain insight into observed behaviors or counterfactual outcomes. Even complex quantitative models typically deliver relatively simple implications for economic behavior. However, in some areas, the simple lessons from our models have been adopted by households, firms, or policymakers and used to guide actual behavior, which then raises the question of whether, in such areas, these simple rules are close to optimal in real-world environments. Over the past decade, big data and advances in machine-learning algorithms and hardware have led to a leap forward in the complexity and realism of economic problems that can be calibrated and solved. We develop these tools and solve for optimal behavior in one particular complex environment, household portfolio choice, and show by example how one can measure the costs of current simple rules, evaluate their robustness, and provide evidence on the benefits of more complex rules.

We study the implications for investor well-being of the simple rule that imposes an age-dependent portfolio strategy in retirement accounts, one that mimics the portfolio of existing (low-fee, index) Target Date Funds (TDFs), the most common investment choice in retirement savings plans in the U.S. These rules, derived from quantitative, partial-equilibrium lifecycle models of saving and portfolio choice as pioneered in Samuelson (1969) and Merton (1969, 1971), have been widely adopted and have had a large impact on the allocation of household wealth between stocks and bonds (see e.g. Parker, Schoar, Cole, and Simester, 2025).[1]

We evaluate this rule by developing a machine learning method and using it to solve a relatively realistic lifecycle model that includes a large number of features of investors' environments each of which has been shown to be quantitatively important for optimal investor behavior. More specifically, we study the lifecycle optimization problem of a household that consists of a husband and wife who consume both housing and non-housing consumption and are each endowed with a sex-specific earnings profile that has stochastic, left-skewed, serially-correlated shocks. The household allocates its financial wealth among a stock index, a bond index, and a money market account, all with returns that are both serially correlated and correlated with labor income. It can hold these assets in liquid accounts, or save in (pre-tax) retirement accounts with limited employer matching

[1] Since legal changes allowed TDFs as defaults in retirement plans in 2006, U.S. retail investors have increasingly relied on TDFs to allocate their investments, so that by 2019, TDFs accounted for two-thirds of the financial wealth of young investors and TDFs and related balanced funds accounted for more than $4 trillion and held roughly 5% of the US equity fund market (Parker, Schoar, and Sun, 2023; Investment Company Institute, 2020; Morningstar, 2020, figures 2.2 and 8.20).

and tax penalties for early withdrawal. For housing, the household chooses between renting, owning with a mortgage, or owning outright, and must pay a cost to sell or to refinance. There is a cash-in-advance constraint and a simple tax and benefit system that includes a consumption floor. During retirement, each individual receives a pension that is a function of lifetime labor income, faces mortality risk and stochastic medical expenses, and gets utility from bequests. Three areas we model only simply are retirement (exogenous and fixed), children (a deterministic, age-dependent utility shifter), and divorce and potential remarriage (not happening).

We parameterize our baseline model and two variants using existing large-sample information, such as income processes based on Social Security earnings records. The variants are a model with relatively impatient households who accumulate less wealth, and a model in which equity return are largely unpredictable. Given our purposes, we do not estimate the model on observed behavior, and behaviors other than just portfolio choice deviate from observed household behavior. In particular, in the baseline model, on average households accumulate a lot of wealth, and in all models households hold only a small share of wealth in home equity, as we present in Sections 4.1, 5, and 6.

We have three main substantive results, all applying to the (relatively well-off) households whose lives we model: two-earner couples that accumulate investable wealth throughout life and do not become extremely wealthy either from inter-generational transfers or own business income. We also make contributions to the literature on how to using machine learning techniques to solve dynamic economic models.

First, on average, the share of financial wealth that a household should hold in stocks is hump-shaped over the working life, rising rapidly to 70% before age 35, and declining steadily to just under 50% at and during retirement. However, the average optimal share of *retirement* wealth invested in stocks starts above 80% and declines steadily until dropping just below 50% after retirement. The decline with age is consistent with a typical TDF, but optimal shares decrease more steadily over working life and end higher than the typical TDF glide path.[2]

Second, the average optimal behavior masks substantial variation. Specifically, the 90th percentile of the cross-household distribution of optimal portfolio share in stocks is close to 100% for wealth in retirement accounts at all ages. The 10th percentile starts lower than the average optimal share, and declines rapidly with age falling below 20% by age 30 and

[2]The share in stock in a TDF typically declines by another 10-20% after retirement while the optimal share in our model declines by only about 6% after retirement.

to just a few percent by the early 40s. This variation is driven primarily in differences in wealth levels, the state of the business cycle, and the expected equity premium. These findings quantify the set of strengths and weaknesses of TDFs discussed in Campbell (2016) (Section 5.1) as well as echoing the conclusion of Gomes, Michaelides, and Zhang (2022) that TDFs that take advantage of return predictability would deliver better investor outcomes.

Our third set of results shows that there is the potential to improve household outcomes by the equivalent of three and half percent of consumption in every state and date by moving from a TDF investment strategy to a fully optimal rule.[3] This number is slightly larger when households are impatient and accumulate less wealth and is only one-and-a-quarter percent when expected returns do not vary directly with the payout ratio. To be clear, we focus on evaluation of the welfare benefits measured with a discount factor of one, so that we weight flow expected utility in each year of life equally rather than taking a start-of-life perspective which would imply a lower cost of bad late-in-life outcomes.[4] These numbers do not depend much on whether we allow the household to re-optimize all other behaviors or not.

We are able to evaluate the simple rule in a model with more than 20 state variables and shocks by using new tools from the field of deep reinforcement learning, the field of machine learning that studies sequential decision-making. Specifically, building on Duarte, Duarte, and Silva (2024), we develop a method that uses deep reinforcement learning to solve dynamic stochastic models like this lifecycle model – models with many states and controls, both discrete and continuous controls, highly non-linear policy functions, and high-variance long-horizon simulation objectives – using a policy gradient algorithm (Sutton, McAllester, Singh, and Mansour, 2000).

In brief, we parameterize the policy functions over the high-dimensional state space using a fully connected feedforward neural network that takes age as an input, with parameters shared across all ages, together with a second network that values the discrete housing choices. We use a stochastic gradient descent algorithm to solve for the parameters of the networks that maximize the expected lifetime utility over a large number of simulated sample paths. The challenges of this lifecycle problem lead us to employ projection-based

[3] These welfare losses are similar to those calculated by Dahlquist, Setty, and Vestman (2018) in a simpler model moving from age-based to completely optimal rules.

[4] The standard approach, while not inconsistent with the revealed-preference approach to welfare, is to put a Pareto weight of one on the value function of the household at age 25 rather than give some weight to the preferences of households when older Caplin and Leahy (2004).

feasibility restrictions, performance-guided moving-average optimization, and separation of discrete and continuous controls. We conjecture that our method may readily be applied to evaluate the robustness of simple rules in a wide range of partial-equilibrium dynamic problems faced by managers, employers, small open economies, etc.

The main advantage of this approach over traditional numerical dynamic programming (NDP) is the massive increase in speed that makes it tractable to solve previously infeasible problems. The traditional NDP approach would first characterize the solution to the household's problem by a set of optimality conditions and budget constraints, then construct grids on which choices can be characterized by matrices, and then finally use an optimization algorithm and numerical integration to solve for optimal behavior recursively. In contrast, our method maximizes expected utility using simulated sample paths which avoids computationally-slow numerical integration. And we make use of software (JAX) and hardware developed specifically for machine learning applications, which have been shown considerably reduce computational time even when combined with traditional NDP (Duarte, Duarte, Fonseca, and Montecinos, 2020). These codes will be publicly available.

There are two other advantage. First, this method is far easier to use and program (and so less prone to error) than traditional numerical dynamic programming methods. For example, there is no need to specify the density and scale of high-dimensional grids over which policy functions can then be defined as matrices. Second, and more interestingly, this method arguably captures how investors, practitioners, or data scientists actually determine optimal behavior. Optimal behavior is determined by observing what works best for people living their lives according to different portfolio allocation and consumption rules, that is, from watching and learning from the lives of others.

Related literature Our methodological contribution adds to a recent and rapidly-evolving literature that develops methods to solve dynamic models using deep learning tools (Duarte et al., 2024; Fernandez-Villaverde, Hurtado, and Nuno, 2020; Maliar, Maliar, and Winant, 2021; Azinovic, Gaegauf, and Scheidegger, 2022). Most closely, and following our paper, Druedahl, Huleux, and Røpke (2026) provides an overlapping but distinct approach to solving lifecycle models.[5] Relative to the more general literature applying machine learning to economic models, the tools we use are designed to handle our lifecycle problem with the combination of a long time horizon, a large number of state variables and shocks, highly

[5]Also, see Druedahl and Røpke 2026 for nice work benchmarking the use of different algorithms for solving lifecycle models.

non-linear policy functions, both discrete and continuous actions, and occasionally binding constraints.

Our quantitative model of lifecycle portfolio choice also advances the large literature on portfolio choice (see the surveys Curcuru, Heaton, Lucas, and Moore, 2010; Wachter, 2010) and provides a baseline for future quantitative studies of optimal portfolio choice. While our model omits some features of reality, and while each individual ingredient is not new, the model contains the features of the lifecycle problem that we judge to be both most relevant and most important for lifecycle portfolio choice. Further, our model is calibrated using state-of-the art estimates of the deterministic and stochastic processes facing the investor (e.g. DeNardi, French, and Jones, 2010; Wachter, 2010; Guvenen, Ozkan, and Song, 2014). In the conclusion we discuss missing elements and weaknesses that present avenues for future improvements.

While our focus in this paper is on a particular subset of the population and the evaluation of sub-optimal decision rules, our paper contributes to a long line of quantitative analyses of optimal portfolio behavior in lifecycle models. These analyses can be grouped into two different types: those that focuses on inferring features and parameters of the model from observed behavior, and those that offer prescriptive analysis of how investors should allocate their portfolios. Bodie, Merton, and Samuelson (1991), Gakidis (1998), Campbell and Viceira (2002) (Chapter 6), Gomes and Michaelides (2003), Storesletten, Telmer, and Yaron (2004), Cocco, Gomes, and Maenhout (2005), Davis, Kubler, and Willen (2006), Catherine (2021), and Shen (2021) all have advanced the analysis of lifecycle portfolio choice in the presence income risk that cannot be fully hedged. Gomes, Michaelides, and Polkovnichenko (2006) and Dammon, Spatt, and Zhang (2004) considered the role of taxes and tax deferred retirement accounts, Cocco (2005), Hu (2005), and Yao and Zhang (2005) incorporate housing, and Li and Smetters (2011) and Yogo (2016) focuses on the role of Social Security and of medical expenses during retirement respectively.[6]

We cite the main papers that provide the ingredients for our model – both the mathematical structure and the calibrated parameters – as we describe each element of the model. Similarly, we place our method in the machine learning literature when we describe our method in Section 3.1.

[6]More recently Calvet, Campbell, Gomes, and Sodini (2021) include both housing and returns although perfectly correlated. Also of note, a subsection of Cocco et al. (2005) considers the role of medical expenses while Koijen, Nieuwerburgh, and Yogo (2016) considers allocations across annuity and insurance products rather than stocks and bonds.

1 The Lifecycle Model

Time is discrete and t indexes age in years. At each age, the household chooses how much to save and how much to consume of non-housing goods and housing. Households receive labor income and earn investment returns that follow known stochastic processes. Optimal consumption and portfolio choices are a set of decision rules for every age that maximize the household's (mathematically) expected present discounted value of utility flows from consumption given knowledge of the structure and parameters of the problem and the history of realizations of stochastic processes up to that age.

Overview The complexity and the realism of the model come from the budget constraint which has the following features. The income process for each spouse consists of a deterministic lifecycle profile subject to negatively skewed persistent and transitory shocks during working life. During retirement, individuals receive pension income based on lifetime earnings and face medical expense shocks. Households face an approximation of the progressive US income tax and benefit system including a consumption floor during retirement. Households can allocate their non-housing wealth among 3 financial assets representing stocks, long-term bonds, and liquid savings. Each returns process consists of an idiosyncratic shock and a loading on a common autoregressive process which we refer to as the payout ratio of our model's stock market. Correlation among labor income and asset returns is captured by having innovations in the payout ratio, returns, and the labor income process all depend on a stochastic business cycle state. Households can save into liquid financial accounts or retirement accounts, and the model contains a detailed representation of the legal and institutional structure of retirement saving. The household can rent or own housing, and there are adjustment costs associated with changing the size of an owned home or a mortgage. The household faces constraints on borrowing, on short selling, on consumption from a liquid-wealth-in-advance requirement, and on the loan-to-value ratio of its mortgage when purchasing or refinancing. Finally, the household experiences mortality shocks during retirement, and values leaving a bequest at death. All stochastic processes are fixed and known to households. The most notable omission from our model is that we study only a traditional family consisting of a man and a woman with fixed ages of retirement and who remain married until death do them part.

We set out the model structure and budget constraints below, starting with the household's working life and then describing retirement. We conclude with the objective function and statement of the complete problem. The calibration of the model is presented in Section 2.

1.1 Working Life

The household consists of two individuals, a man and a woman, indexed by $i \in \{1, 2\}$, who both work from age $t = T_0 = 25$ until the exogenous retirement age $t = T_R - 1 = 66$ and who both survive until age T_R with probability one. Households are also differentiated by the deterministic component of each spouse's labor income which depends on their rank in the initial distribution of permanent income, which we denote by q_0^1 and q_0^2.

All random variables sub-scripted by t are realized before the households makes any time-t decisions.

1.1.1 Common Risks

There are two 'aggregate' risks in our partial-equilibrium framework which create correlations among the stochastic processes faced by a given household. A stochastic process representing the business cycle generates correlation between labor income and returns on different assets. Second, a stochastic process representing aggregate risk premia generates correlation among asset returns. These two aggregate processes are correlated.

First, the economy can either be in a recession or expansion, with the state variable $e_t = 1$ indicating a recession and $e_t = 0$ indicating an expansion. The economy's state e_t evolves according to the 2x2 transition matrix P_e.

Second, we capture fluctuations in expected returns from serially-correlated fluctuations in the log of the payout yield (dividend price ratio) of the stock market, v_t.

$$v_t = \theta_v^1 + \theta_v^2 \Delta e_t^+ + \theta_v^3 \Delta e_t^- + \theta_v^4 v_{t-1} + \epsilon_{v,t}, \tag{1}$$

$$\epsilon_{v,t} \sim N(0, \sigma_v), \tag{2}$$

where $\Delta e_t^+ = \max\{0, e_t - e_{t-1}\}$ and $\Delta e_t^- = \max\{0, e_{t-1} - e_t\}$. The dividend yield follows an AR(1) process, with an intercept that shifts when recessions start or end.

1.1.2 Labor Income

To model the household's labor income, we follow the work done with Social Security earnings records by Guvenen et al. (2014) (GOS) and Guvenen, Kaplan, Song, and Weidner

(2017) (GKSW). Each individual's log labor income, $y_t^i = \ln(Y_t^i)$, depends on two factors: a deterministic age profile $\bar{y}_t^i$, and a stochastic process of shocks around this profile that follows GOS's main parametric model. The deterministic profile is a function of gender i, age, and the initial income quantile for each individual, q_0^i, so that $\bar{y}_t^i = \bar{y}(t; i, q_0^i)$. To ease notation, we suppress the dependence of labor income on q_0^i for the rest of Section 1.1.

In each period, log labor income for individual i is equal to the age profile, plus a persistent shock, x_t^i, and a transitory shock, ϵ_t^i.

$$y_t^i = \bar{y}_t^i + x_t^i + \epsilon_t^i, \tag{3}$$

$$x_t^i = \rho_x x_{t-1}^i + \eta_t^i, \tag{4}$$

$$\epsilon_t^i \sim N(0, \sigma_\epsilon), \tag{5}$$

$$\eta_t^i = \begin{cases} \eta_{1,t}^i \sim N(\mu_{\eta_1, e_t}, \sigma_{\eta_1}) \text{ with prob. } p_1 \\ \eta_{2,t}^i \sim N(\mu_{\eta_2, e_t}, \sigma_{\eta_2}) \text{ with prob. } 1 - p_1. \end{cases} \tag{6}$$

The innovation η_t^i in the AR(1) process for x_t^i follows a mixture distribution, which allows for the non-normalities (in particular, negative skewness and excess kurtosis) that GOS document in earnings growth rates.[7] Cyclicality (in particular, countercyclical negative skewness) enters the income process through the dependence of the mean AR(1) innovations, μ_{1,e_t} and μ_{2,e_t}, on the economy's state e_t.

To model taxes and convert the household's gross income to net income, we use the parsimonious tax function introduced by Heathcote, Storesletten, and Violante (2017) (HSV). The household's tax liability as a function of combined gross income, $Y_t = Y_t^1 + Y_t^2$ is given by

$$L(Y_t) = \min\{Y_t - \lambda Y_t^{1-\kappa}, 0.5 Y_t\} \tag{7}$$

where the parameter κ controls the progressivity of the tax schedule and the parameter λ shifts the tax function to determine the economy's average tax burden. Heathcote et al. (2017) show that this simple tax function provides a very good approximation to the US tax system (including federal, state, and payroll taxes), and we use the parameter estimates reported in their paper. We cap the tax rate at 50% which occurs for our parameterization just above $500,000.

[7] And that Schmidt (2025) shows can rationalize asset pricing patterns in a non-lifecycle model.

1.1.3 Financial Assets and Accounts

The household's financial wealth is composed of assets held in a liquid accounts, a^L, and a (relatively) illiquid retirement account, a^I. The illiquid retirement accounts capture 401(k) and other tax-deferred retirement saving accounts that receive special tax treatment. Contributions to these tax-deferred accounts reduce the household's current tax burden and accumulate tax-free, but are taxed as income when withdrawn and are subject to a 10% penalty if withdrawn before retirement.[8] The household starts life with no assets.

In each account, the household can save and invest in $J = 3$ financial assets: short-term government bills ($j = 1$), long-term corporate bonds, and equities.[9] To incorporate realistic return predictability in a tractable way, we follow Wachter (2010) and make each return correlated with the stock market's aggregate log dividend yield, v_t (equations (1) and (2)). The log of the gross return on asset j between time $t-1$ and t, $R_{j,t}$, is given by:

$$\log(R_{j,t}) = \theta_j^1 + \theta_j^2 \Delta e_t^+ + \theta_j^3 \Delta e_t^- + \theta_j^4 v_{t-1} + \epsilon_{j,t}, \tag{8}$$

$$\epsilon_{j,t} \sim N(0, \sigma_{r,j}), \tag{9}$$

where again $\Delta e_t^+ = \max\{0, e_t - e_{t-1}\}$ and $\Delta e_t^- = \max\{0, e_{t-1} - e_t\}$. Each asset's return depends on an intercept term, which is allowed to differ at the beginning and end of recessions; the aggregate log dividend yield state variable; and a transitory shock.

We denote the wealth in asset class J in each account type, L and I, after returns are realized at (the start of) age t but before age-t saving or withdrawals by $a_{j,t}^L$ and $a_{j,t}^I$ respectively so that:

$$a_t^L = \sum_{j=1}^{J} a_{j,t}^L, \qquad a_t^I = \sum_{j=1}^{J} a_{j,t}^I \tag{10}$$

Define $s_{j,t}^x$ to be the household's net contribution to its holdings of asset j in account type x during period t (i.e., contributions net of withdrawals, after time t returns and incomes are realized and at the same time as consumption and other decisions are being made),

[8] We model traditional retirement accounts not Roth accounts which are still relatively uncommon. In 2025, while 95% of Vanguard-administered plans offered a Roth option, 82% of participants contributed only to traditional accounts (Vanguard Group, 2026).

[9] Thus we do not consider diversification within stocks or bonds, an issue that can amplify risk (Fagereng, Gottlieb, and Guiso, 2017).

with $s_t^x = \sum_{j=1}^J s_{j,t}^x$. We can then write the state-evolution equation for financial assets as

$$a_{t+1}^L = \sum_{j=1}^J R_{j,t+1}(a_{j,t}^L + s_{j,t}^L), \tag{11}$$

$$a_{t+1}^I = \sum_{j=1}^J R_{j,t+1}(a_{j,t}^I + \Gamma(s_{j,t}^I)), \tag{12}$$

where $\Gamma(\cdot)$ is a function allowing for matching employer 401(k) contributions, with match rate k and limit l until households reach retirement:

$$\Gamma(x) = \begin{cases} x & \text{if } x \leq 0 \text{ or } t \geq T_R \\ (1+\frac{k}{2})x & \text{if } 0 < x \leq lY_t \\ x + \frac{k}{2}lY_t & \text{if } x > lY_t. \end{cases} \tag{13}$$

We divide the match rate k by 2 to save on state variables and yet to capture the idea that not all jobs match contributions. Annual contributions to tax-deferred retirement accounts are capped as in US law:

$$s_t^I \leq s_{max,t}^I. \tag{14}$$

The contribution limit $s_{max,t}^I$ is constant during the working life, then becomes negative (and age-dependent) beginning at age 70 in order to reflect the IRS schedule of required minimum distributions.[10]

Finally, households cannot short any asset (other than with a mortgage, as described subsequently), so holdings in each asset class and account type must be nonnegative:

$$a_{j,t}^L \geq 0, \tag{15}$$

$$a_{j,t}^I \geq 0. \tag{16}$$

Now define $B_{a,t}$ to be the net contribution to the household's time t budget constraint

[10]We maintain the assumption that households behave optimally other than the constraints we explicitly introduce, so that the primitives of the problem contain no references to default retirement saving rates or portfolios, features that are relevant for actual behavior and welfare (Beshears, Choi, Laibson, and Madrian, 2008; Choukhmane, 2025).

of all decisions regarding portfolio choice and saving in financial assets. We have:

$$B_{a,t} = -s_t^L - s_t^I - L'(Y_t) \sum_{j=1}^{J} \left(R_{j,t} - 1\right) a_{j,t}^L + \xi(s_t^I, t) \tag{17}$$

where $L'(Y_t)$ is the marginal tax rate from equation (7), and where ξ is a function capturing the tax benefits of contributions to ($s^I \geq 0$) retirement accounts less the tax and possibly early-withdrawal penalty costs associated with withdrawal from retirement accounts ($s^I < 0$):[11]

$$\xi(s^I, t) = \begin{cases} L'(Y_t)s^I & \text{if } s^I \geq 0 \\ (L'(Y_t) + 0.1)s^I & \text{if } s^I < 0 \text{ and } t < 60 \\ L'(Y_t)s^I & \text{if } s^I < 0 \text{ and } t \geq 60 \end{cases} \tag{18}$$

Net contributions to asset holdings (i.e., savings) enter negatively into the budget constraint, as do taxes incurred on the realized returns on taxable asset holdings. The benefit of contributing to tax-deferred retirement accounts is the reduction in current taxes, while the cost is the taxes on withdrawals plus a 10% penalty associated with withdrawals from the illiquid account prior to age 60.[12]

Finally, we impose a cash-in-advance constraint that reflects the fact that people hold liquidity for transactions related to consumption. We do this in order to capture the transactions demand for liquidity, and to prevent the household from (unrealistically) financing normal consumption expenditures by making repeated withdrawals from its illiquid retirement savings accounts. In particular, as long as the household has any wealth in its illiquid retirement account, it cannot consume more than a fixed multiple of its liquid

[11]Note that for simplicity, and partly motivated by the fact that households effectively face constant marginal tax rates with the piecewise linearity of the US tax schedule, we take the marginal tax rate on withdrawals during working life to be the derivative of the tax function evaluated at the household's current labor income (and do not account for second-order effects where withdrawals change income and thereby change the marginal tax rate). Note also that we simply tax realized returns in taxable accounts at the marginal income tax rate rather than including separate dividends or capital gains rates. The preferential tax treatment of equity can make a different to optimal equity share in retirement accounts (Dammon and Spatt, 2026). Since withdrawals from the illiquid retirement account are likely to be a substantial share of income during retirement, we account for second-order effects during retirement by treating withdrawals during retirement equivalently to other income (see equation (39)).

[12]Note that, unrealistically, we allow a tax benefit for contributions to retirement plans from age 65 to 69. In our numerical solution however, the average household has a large negative saving rate into retirement wealth immediately before retirement and is dis-saving significantly in the years immediately following retirement.

wealth held in short-term debt:

$$c_t \leq Ma^L_{1,t} \;\text{ if } a^I_t > 0. \tag{19}$$

where c_t denotes non-housing consumption.

1.1.4 Housing

Our model of housing choices is an expanded version of the structure of Berger, Guerrieri, Lorenzoni, and Vavra (2018) (BGLV). Households consume non-housing consumption c_t and housing h_t. At any point in time, the household can choose to either rent or own a home, with $o_t = 0$ indicating renting and $o_t = 1$ indicating ownership. Households start life as renters. The stock of housing owned by the household is given by $\tilde{h}_t$, with $\tilde{h}_t = 0$ when the household is a renter, and ownership conveys an additional utility benefit (not in BGLV), discussed subsequently as part of the objective function. Thus $h_t = \tilde{h}_t$ if $o_t = 1$ and it will turn out that households choose $h_t > \tilde{h}_t = 0$ if $o_t = 0$. Households who own housing stock may take on mortgage debt, d_{t+1} (mortgage debt chosen in period t is carried into $t+1$ and so we denote it with $t+1$ like we do for assets chosen in t). We define by $B_{h,t}(o_{t-1})$ the net contribution to the household's time t budget constraint of all housing decisions which depends on whether the household entered the period as a renter or owner. If the household is a renter and consumes housing h_t, it pays a rental cost of $\phi p_t h_t$ where ϕ is the constant rent-to-house-price ratio and p_t is the price of housing (relative to the numeraire non-housing consumption good c_t). Following Berger et al. (2018), p_t is a geometric random walk with drift:

$$p_{t+1} = \nu_t p_t, \tag{20}$$

$$\log(\nu_t) \sim N(\mu_\nu, \sigma_\nu). \tag{21}$$

Renters who become homeowners make down payments equal to the cost of the house less the amount of the new mortgage which is subject to a loan-to-value constraint. They must also pay for maintenance on their new house, $\delta\tilde{h}_t$, to cover depreciation.

Thus, for households that enter the period as renters ($o_{t-1} = 0$), we have:

$$B_{h,t}(0) = -(1 - o_t) [\phi p_t h_t]$$

$$-o_t \left[p_t \tilde{h}_t - d_{t+1} + \delta p_t \tilde{h}_t + (1 - L'(Y_t))(R_{1,t} - 1 + \Delta r_m) d_{t+1} \right] \tag{22}$$

$$d_{t+1} \begin{cases} = 0 & \text{if } o_t = 0 \\ \leq (1 - \iota) p_t h_t & \text{if } o_t = 1. \end{cases} \tag{23}$$

The mortgage, d_{t+1}, must stay at zero if the household remains a renter. If the household becomes a homeowner, its mortgage choice is restricted by a loan-to-value constraint: households buying a new house (or refinancing, below) must satisfy $d_{t+1} \leq (1 - \iota) p_t h_t$, so that the loan-to-value ratio is at most $1 - \iota$. Note that, in order to match the real-world housing market, the loan-to-value constraint is only imposed when a house is purchased or a mortgage refinanced. An existing homeowner whose house declines in value is never forced to refinance or accelerate payment on its mortgage.

If the household enters the period as an owner and either adjusts its housing stock or becomes a renter, it pays a transaction cost of $f^h p_t h_{t-1}$ (i.e., a fraction f^h of the housing stock it inherited from last period and is now selling, valued at the current house price). Households that remain homeowners must pay maintenance costs. Continuing homeowners must also make mortgage payments if they do not refinance. The amortization schedule of mortgages is described by the parameter χ: a household with current mortgage debt d_t that does not refinance must service its debt by making a payment of at least χd_t. Alternatively, continuing homeowners can adjust their mortgage debt up (refinance) by paying a transaction cost of $f^d p_t h_t$ (a fraction f^d of their house's current value). Finally, owners can walk away from their houses and their (non-recourse) mortgages.[13] We also prohibit refinancing an underwater house (when $(1 - i - f^d) p_t h_{t-1} < \chi d_t$).

[13] In addition to reality, we make this choice to ensure that households can always choose positive consumption. BGLV do not require this feature. It appears that this is because of approximations in traditional dynamic programming: "Shocks to house prices in the model are not large enough..." to ever force consumption negative (BGLV, Online Appendix A.4 footnote 60). However, this assumption is critical for us, both because our method draws from normal distributions and because an exploration device overrides discrete decisions at random early in training, as we describe in Section 3.1. Thus a very poor household may try selling an underwater house and make the objective negative infinity.

For households entering the period as homeowners, $o_{t-1} = 1$, we have:

$$B_{h,t}(1) = -o_t[\delta p_t \tilde{h}_t + (1 - L'(Y_t))(R_{1,t} - 1 + \Delta r_m)d_{t+1} + \mathbf{1}_{\{\tilde{h}_t = \tilde{h}_{t-1},\, d_{t+1} > d_t\}}(f^d p_t \tilde{h}_{t-1} + d_t - d_{t+1}) \tag{24}$$

$$+\mathbf{1}_{\{\tilde{h}_t \neq \tilde{h}_{t-1}\}}(p_t \tilde{h}_t - d_{t+1} - h^{equity}] - (1 - o_t)[\phi p_t h_t - h^{equity}] \tag{25}$$

$$d_{t+1} = \begin{cases} = 0 & \text{if } o_t = 0 \\ \leq (1 - \iota) p_t h_t & \text{if } o_t = 1 \text{ and } (\tilde{h}_t \neq \tilde{h}_{t-1} \text{ or } d_{t+1} > \chi d_t) \\ \leq \chi d_t & \text{otherwise.} \end{cases} \tag{26}$$

where $h^{equity} = max\{(1 - f^h) p_t \tilde{h}_{t-1} - d_t, 0\}$ is home equity. Depreciation costs enter negatively into the budget constraint, as does tax-deductible mortgage interest (paid at the short-term interest rate plus a mortgage spread of Δr_m) and principal repayment. If the household adjusts its owned housing stock $\tilde{h}_t$ (either by choosing a new nonzero value or choosing $\tilde{h}_t = 0$ and becoming a renter) it pays the price of its new house (if any) and receives the selling price of its old house net of transaction costs. Finally, if it maintains its current housing stock but refinances its mortgage, it pays refinancing costs and gains the additional funds borrowed in liquid wealth (since the last term, $d_t - d_{t+1}$, is negative). Equation (26) states that if the household becomes a renter then next period's mortgage must be zero. If the household adjusts its housing stock or refinances, then the household can choose any mortgage amount allowable under the loan-to-value constraint. Otherwise, the household can choose any mortgage amount below that required by the amortization rule.

1.1.5 Working-life Objective Function and Unified Flow Budget Constraint

Total household consumption at each age is a Cobb-Douglas aggregate of non-housing consumption c_t and the service flow from housing with share parameter α. The service flow is a constant fraction of h, which we take to be the rent-to-price ratio.[14] We assume constant relative risk aversion utility of total consumption in each year, with risk aversion coefficient γ. To account for the fact that consumption needs change with family size, consumption enters utility in per-effective-householder form, where the effective size of the household is w_t, the square root of the average family size at each age. Family size rises early in life

[14]The exact fraction of h does not affect the objective function, only the measurement of the consumption aggregate.

and then falls to 2 at retirement (as detailed in Section 2). Thus, the consumption aggregate and utility of consumption at each are given by

$$C_t = c_t^{\alpha}(\phi h_t)^{1-\alpha} \qquad \text{and} \qquad U(w_t, C_t) = \frac{(C_t/w_t)^{1-\gamma}}{1-\gamma}. \tag{27}$$

Define the discount factor as β and the retirement value function, which returns the maximized expected discounted utility during retirement, as $V_R^*(\cdot)$. Additionally, as a state variable summarizing the household's liquid financial resources, define cash on hand, X_t, as the liquid wealth available to the household at the beginning of time t, net of taxes and mortgage interest payments:

$$X_t = a_t^L + Y_t - L(Y_t) - L'(Y_t) \sum_{j=1}^{J} \left(\frac{R_{j,t} - 1}{R_{j,t}} \right) a_{j,t}^L - (1 - L'(Y_t))(R_{1,t} - 1 + \Delta r_m) d_t \tag{28}$$

Policy functions are functions of the complete set of state variables at the beginning of each age t: $\Xi_t = \left\{ e_t, v_t, \{x_t^i\}_{i=1}^2, Y_t, \{\sum_{t'=T_0}^{t} y_{t'}^i\}_{i=1}^2, X_t, a_{1,t}^L, a_t^I, p_t, o_{t-1}, h_{t-1}, d_t \right\}$. The sum of each household member's income up to time t is a state variable because several stochastic processes during retirement, discussed below in Section 1.2, depend on individuals' lifetime labor earnings. The cash-in-advance constraint in (19) makes the liquid holdings of short-term debt (in addition to cash on hand) a state variable.

The household chooses age-specific policy functions for consumption, portfolio allocations, home ownership, and mortgage debt, $\mathcal{A}_t = \left\{ c_t, o_t, h_t, \{a_{j,t+1}^L\}_{j=1}^J, \{a_{j,t+1}^I\}_{j=1}^J, d_{t+1} \right\}$ as a function of the state variables Ξ_t at each age, which we denote as $\pi_t(\Xi_t)$, to maximize the expected discounted sum of time-separable flow utility:

$$V^{\pi}(\Xi_{T_0}) = \mathbb{E}\left[\sum_{t=T_0}^{T_R-1} \beta^{t-T_0} U(w_t, o_t, C_t) + \beta^{T_R - T_0} V_R^*(\Xi_{T_R}) \;\middle|\, \Xi_{T_0} \right] \tag{29}$$

subject to the comprehensive flow budget constraint

$$Y_t - L(Y_t) - c_t + B_{a,t} + B_{h,t}(o_{t-1}) = 0, \tag{30}$$

the constraints (15)-(14) and (19), and where the income process is defined in (3)-(6), the tax function is defined in (7), $B_{a,t}$ defined in (17)-(18), and $B_{h,t}(o_{t-1})$ is defined in (22)-(26). We now turn to the retirement value function, $V_R^*(\Xi_{T_R})$.

1.2 Retirement

When the household reaches retirement, it continues to makes consumption, saving, and portfolio-choice decisions to maximize expected discounted utility over the remainder of its life. However, following DeNardi et al. (2010) (DFJ), we make several changes to reflect the changing financial risks that households confront as they age: income becomes deterministic pension payouts rather than stochastic labor earnings, stochastic health costs become a significant household expenditure, and individuals face mortality risk. Let ν_t^i equal 1 if individual i is still alive at time t and 0 otherwise (the exact mortality process for ν_t^i is described below).

1.2.1 Pension Income, Medical Expenditures, and Mortality

Instead of labor income during retirement, each individual receives non-stochastic income representing Social Security and other pension income and faces stochastic expenditure shocks representing medical expenses. Following DFJ, each individual's pension, health, medical expenditure, and mortality processes depend on the individual's permanent income rank. We denote this rank by q_R^i and calculate it as the quantile into which the individual's realized lifetime labor earnings, $\sum_{T_0}^{T_R-1} y_t^i$, fall in the distribution of life-time labor earnings for the distribution including all q_0^i (as described in subsection 2.5).

Pension income is given by

$$Y_t^i = \nu_t^i Y(t; q_R^i), \tag{31}$$

so that income Y_t^i is a deterministic function of age during retirement, although stochastic from the perspective of working life since Y_t^i depends on lifetime earnings. We postpone the treatment of taxes until section 1.2.2 because withdrawals from retirement accounts are taxed as regular income. We set pension income Y_t^i to zero when individual i is no longer alive.

Health shocks and stochastic medical expenditures follow DFJ. The state variables g_t^i represent each individual i's current health status: $g_t^i = 0$ indicates "healthy," while $g_t^i = 1$ indicates "sick." Health status transitions depend on previous health status, gender, and age, with

$$\Pr(g_t^i = 1) = P_g(g_{t-1}^i, i, t; q_R^i). \tag{32}$$

Each individual's out-of-pocket medical expenditures m_t^i (i.e., net of Medicare and other insurance) depends on their health status g_t^i. Family medical expenses are the sum of the individual expenses but are capped at $\bar{m}$. The process for medical expenses is given by:

$$m_t = \min[m_t^1 + m_t^2, \bar{m}] \tag{33}$$

$$\log(m_t^i) = \nu_t^i m(g_t^i, i, t; q_R^i) + \nu_t^i \sigma(g_t^i, i, t; q_R^i)\psi_t^i, \tag{34}$$

$$\psi_t^i = \zeta_t^i + \epsilon_{\psi,t}^i, \tag{35}$$

$$\zeta_t^i = \rho_\zeta \zeta_{t-1}^i + \epsilon_{\zeta,t}^i, \tag{36}$$

$$\epsilon_{\psi,t}^i \sim N(0, \sigma_\psi), \tag{37}$$

$$\epsilon_{\zeta,t}^i \sim N(0, \sigma_\zeta). \tag{38}$$

Log medical expenditures are the sum of a health, sex, and age specific mean cost $m(g_t^i, i, t)$ and a shock term ψ_t^i, magnified by a health, sex, and age specific volatility term $\sigma(g_t^i, i, t)$. The medical expenditure state variable ζ_t^i follows an AR(1) process. Medical expenditures are zero once an individual dies, $m_t^i = 0$ if $\nu_t^i = 0$.

Turning to mortality, each person's probability of death depends on their age, sex, income rank at retirement, and health status. If individual i is alive, $v_t^i = 1$, else $v_t^i = 0$. Conditional on having survived through age t, each individual i survives, $v_{t+1}^i = v_t^i$, with the health-, sex-, and age-specific probability $P_d(g_t^i, i, t; q_R^i)$. Following DFJ, anyone still alive at age $T_{max} = 102$ dies with probability 1.

We continue to assume that the household pays taxes during the year, but we add an additional tax payment or refund based on the actual tax due calculated from a more accurate version of the true non-linear tax system. As during working life, taxes paid on income in a given year are a non-linear function of Y plus the marginal tax rate (based only on Y) applied to non-retirement capital income, any retirement saving withdrawals, and (as a benefit) mortgage interest payments. However, during retirement, withdrawals of retirement saving may be substantial and they also affect the tax rate on Social Security income. Thus we calculate the complete tax bill that results from including withdrawals from illiquid accounts in the nonlinear tax function. The household must then pay as taxes (or receive as a refund) in the following year any difference between the linear approximation that it pays in t and the actual tax owed based on the non-linear calculation.

The (nonlinear) tax on income and withdrawals from retirement accounts in retirement,

$\tilde{L}$, is given by:

$$\tilde{L}(Y_t, s_t^I) = \begin{cases} L\left(\max\{0, -s_t^I\}\right) & \text{if } 0.5Y_t + \max\{0, -s_t^I\} < I_{0.5} \\ L\left(0.5Y_t + \max\{0, -s_t^I\}\right) & \text{if } I_{0.5} \leq 0.5Y_t + \max\{0, -s_t^I\} < I_{0.85} \\ L\left(0.85Y_t + \max\{0, -s_t^I\}\right) & \text{if } 0.5Y_t + \max\{0, -s_t^I\} \geq I_{0.85} \end{cases} \tag{39}$$

which reflects the tax treatment of Social Security benefits: either 0%, 50%, or 85% of Social Security income is taxable, depending on the level of "combined income" (half of Social Security income plus illiquid asset withdrawals).[15] The difference between this and what is paid during the year is the tax bill (or refund) that the household must pay in $t+1$:

$$\tau_{t+1} = \tilde{L}(Y_t, s_t^I) - L(Y_t) + \xi(s^I, t). \tag{40}$$

Finally, To reflect means-tested government programs such as Medicaid that support the elderly, a poor household can receive a transfer that provides a minimum level of consumption provided that it does not save or own a house. This government transfer, n_t, is given by

$$n_t = \begin{cases} 0 & \text{if } \; a_{t+1}^L + a_{t+1}^I > 0 \; \text{ or } \; o_t = 1 \\ \max\left\{0, \frac{\underline{c} w_t}{\alpha} + m_t + \tau_t - Y_t + L(Y_t) - B_{a,t} - h^{equity}\right\} & \text{otherwise} \end{cases} \tag{41}$$

When the household's budget constraint is tight enough to push consumption below the consumption floor, $\underline{c}$ adjusted for the effective consumption shifter (w_t) and the budget share of c (α), then a government transfer makes up the difference.[16]

[15] This is only approximately correct because non-asset income, Y_t, from DFJ includes not only Social Security benefits but also defined-benefit pensions and annuities which are actually taxed as ordinary income (in the same way as illiquid asset withdrawals $-s_t^I$ in our model). This approximation is close however because Social Security is all or most of Y for households over 65, averaging 3.5 times the income from private pensions and annuities (Social Security Administration, 2016, pages 222-223).

[16] When solving the model, we also enforce the guarantee directly within the budget: whenever chosen non-housing consumption would fall below $\underline{c}\,\omega_t$, where ω_t is the square root of the number of surviving members, it is set to that floor, and a renter's housing services are floored at $((1-\alpha)/\alpha)\,\underline{c}\,\omega_t$ of rent, with the shortfalls financed by reducing liquid saving. These floors keep expected utility bounded under any feasible policy, including the imposed portfolio rules of Section 4.3.

1.2.2 Financial Assets

During retirement the household does not face withdrawal penalties on its illiquid assets nor does it receive employer matching on retirement contributions (as stated for $t \geq T_R$ in equations 13 and 18). However, it must pay income taxes on withdrawals from the retirement account (as given in equation 18) and is subject to the IRS's rules regarding required minimum distributions from retirement accounts after age 70 (equation 14). Thus, $B_{a,t}$, as defined in equation 17, summarizes the contribution of financial assets to the budget constraint.

1.2.3 Housing

During retirement, the part of the budget constraint relating to housing and mortgages, $B_{h,t}(o_{t-1})$, as well as the constraints on these choices are exactly the same as during working life.

1.2.4 Objective Function and Unified Flow Budget Constraint

We can now write down a complete statement of the household's problem during retirement. Following DFJ, we define utility over bequests, b_t, as

$$U_b(b_t) = \tilde{b}\frac{(b_t + \underline{b})^{1-\gamma_b}}{1-\gamma_b}, \tag{42}$$

which the household receives when the longer-lived spouse dies at age t and where $\tilde{b}$ captures the intensity of the bequest motive, $\underline{b}$ shifts the curvature of utility function and allows bequests to be a luxury good, and γ_b is risk aversion over bequests. Mortality is realized at the beginning of the period, so the bequest amount is given by

$$b_t = a_t^L + a_t^I - L\left(\frac{a_t^I}{5}\right) - \tau_t + \max\{(1-\delta-f^h)p_t\tilde{h}_{t-1} - d_t, 0\}, \tag{43}$$

where $a_t^L + a_t^I - L\left(\frac{a_t^I}{5}\right)$ gives financial assets at the beginning of the period net of taxes paid by the recipient on inherited pre-tax wealth,[17] τ_t gives the carry-forward tax bill determined in the previous period, and the final term gives the current value of the

[17] Under current US tax law, inherited 401(k) or IRA wealth must be withdrawn over 10 years and is taxed as income when withdrawn. Because generally the inheritor will also have labor income, we suppose a tax bill as if the inheritor withdrew over only 5 years.

owned housing stock $\tilde{h}_{t-1}$ carried from the previous period (net of the beginning-of-period mortgage debt, depreciation, and the transaction costs associated with selling the house). The maximum operator around the last term means that underwater mortgages (i.e., $(1 - \delta - f^h) p_t \tilde{h}_{t-1} - d_t < 0$) cannot cause the household's bequest to be less than its financial wealth net of taxes. Additionally, if the carry-forward tax bill is large enough to make the bequest value in (43) negative, we set it to zero.

Utility during retirement is the same as during working life except that the utility from housing is lower if the household owns its own home and has members that are in the "sick" state, to reflect the difficulty of maintaining a home and the increased likelihood of moving to assisted living environments when in poor health.[18] We implement this utility reduction by multiplying homeowners' housing consumption h_t by a scaling factor $\tilde{u}_s < 1$, raised to the number of household members who are sick if $o_t = 1$, so that utility at any age is given by

$$U(w_t, g_t^1, g_t^2, o_t, C_t) = \frac{(w_t^{-1} (\tilde{u}_s^{(g_t^1 + g_t^2)})^{o_t(1-\alpha)} C_t)^{1-\gamma}}{1-\gamma} \tag{44}$$

The set of state variables during retirement is different from that during working life in two ways. First, in place of the state variables for forecasting labor income, x_t^i, and pension income, $\sum_{t_0}^{t} y_t^i$, are state variables for summarizing pension income and forecasting medical expenses, q_R^i and ζ_t^i. Second, the cash-on-hand state variable must account for medical expenses m_t and the previous year's owed taxes τ_t. Denoting the modified retirement cash-on-hand variable as X_t^R, this gives

$$X_t^R = X_t - m_t - \tau_t, \tag{45}$$

where X_t is defined in (28). Finally, the health status (g_t^i) and survival (ν_t^i) states of the household's members are now state variables.[19]

Given Ξ_{T_R} at retirement, the household chooses the set of policy functions for consumption, portfolio allocations, home ownership, and mortgage debt as a function of Ξ_t, which,

[18]For household size, when both members of the household are still alive ($v_t^1 = v_t^2 = 1$), $w_t = \sqrt{2}$. When the shorter-lived spouse dies, $w_t = 1$, reflecting the one-member household's diminished consumption needs. When the second spouse dies, $w_t = \sqrt{v_t^1 + v_t^2} = 0$.

[19]Thus, during retirement the state at age t is given by $\Xi_t = \left\{ e_t, v_t, Y_t, X_t^R, a_{1,t}^L, a_t^I, \{g_t^i\}_{i=1}^2, \{q_R^i\}_{i=1}^2, \{\zeta_t^i\}_{i=1}^2, \{\nu_t^i\}_{i=1}^2, p_t, o_{t-1}, h_{t-1}, d_t \right\}$.

again, we denote by $\{\pi_t(\Xi_t)\}$to maximize:

$$V_R^{\pi}(\Xi_{T_R}) = \mathbb{E}\left[\sum_{t=T_R}^{T_{max}} \beta^{t-T_R}\left(\mathbf{1}_{\{w_t>0\}}U(w_t, g_t^1, g_t^2, o_t, C_t) + \mathbf{1}_{\{w_t=0,\, w_{t-1}>0\}}U_b(b_t)\right) \;\middle|\; \Xi_{T_R}\right] \tag{46}$$

where equations (44) and (42) give the utility and bequest functions. The household is subject to a modified version of the budget constraint during the working life (30) which accounts for the fact that the household faces medical expenditures, may receive government transfers, and pays any additional tax liabilities from the previous year:

$$Y_t - L(Y_t) - c_t - m_t + n_t - \tau_t + B_{a,t} + B_{h,t}(o_{t-1}) = 0. \tag{47}$$

The household is subject to the contribution and withdrawal constraints (15)-(14) and (19). Income is given by (31), the tax function is defined in (39), the process for medical expenditures is defined in (33)-(38), the benefit n_t is defined by (41), b is given by (43), $B_{a,t}$ and $B_{h,t}(o_{t-1})$ are defined as during working life by (17)-(18) and (22)-(26), and finally consumption is given by (44).

The retirement value function referenced in the working-life objective (29), $V_R^*(\Xi_{T_R})$, is given by the maximized value of (46).

2 Model Parameters

The complexity of our setup creates a large number of parameters that we must calibrate or estimate before proceeding to solve the model. Wherever possible, we take parameter values directly from the same literature that we use to construct the model itself. We discuss our parameterization decisions here and provide the list of parameter values in Table I.[20] All dollar amounts in the paper are given in 2013 dollars and are inflation-adjusted using the PCE deflator (following GKSW).

We parameterize the model to match a married couple both born in 1924 (the youngest cohort included in DFJ's data) and their working lives run from age 25 (in 1949) to age 65 (in 1989). We adjust all dollar amounts to 2013 dollars. Households begin life as renters with no financial assets, no housing wealth, no mortgage debt, and the persistent component of

[20]The online appendix contains the inputs that cannot be neatly listed in Table I because they are vectors or matrices.

labor income variance $x^i_{T_0}$ set to 0.

2.1 Common Risks and Asset Returns

We use official NBER recession-dating between 1915 and 2015 to estimate the recession transition matrix P_e.

To estimate the processes for the dividend yield and returns, we use the period 1915-2015. We take the dividend yield, the return on short-term government bills, and the return on equities from Jorda, Kroll, Kuvshinov, Schularick, and Taylor (2019), and the return on bonds from the Dow Jones Total Corporate Bond index.[21] As a robustness check, Section 6 also estimates a version of the return processes that omits the dividend yield as a predictor, re-estimating the intercepts, the business-cycle terms, and the covariance matrix of the return innovations.

2.2 Labor Income and Taxes

We take all parameter values for the labor income process in (3)-(6) directly from GOS.[22]

We build our deterministic age profiles from the summary statistics that GKSW report in their appendix. In particular, GKSW compute a series of quantiles (the 25th, 50th, 75th, and 90th) of the log income distribution, by gender, age (between 25 and 55) and cohort (those that turned 25 between 1957 and 1983).[23] Because the oldest cohort for which GKSW report

[21]The Jorda et al. (2019) dividend yield does not include share buybacks which became substantial during the 1980's. Boudoukh, Michaely, Richardson, and Roberts (2007) shows that as a result, the predictive power of the dividend-price ratio is lower in the last quarter of our sample, while the predictive power of payout yield does not decline. Thus we understate slightly the predictability of U.S. equity returns relative to the unbiased in-sample estimate. However, there are many reasons one might want to use slightly less predictability than estimated from an ex-post, in-sample estimate.

[22]Because the GOS model is estimated to match observed income growth rates rather than deviations from a pre-specified lifetime profile, the income growth shocks it produces are not mean zero. Since our deterministic age profiles already account for expected real earnings growth, we subtract a constant amount from GOS's μ_η values so that the earnings growth shocks produced by (4)-(6) are mean zero.

[23]We thank Fatih Guvenen for providing detailed data on cohort- and percentile-specific lifetime income paths, contained in the file "gksw2017_figuredata_v1.xlsx." In their baseline sample, GKSW impose minimum-participation requirements and include only individuals that earn sufficiently high income in a sufficient number of years. We use the male income profile derived from GKSW's baseline sample. However, given that our targeted cohorts worked during a period of rising female labor force participation, we prefer not to assume that all of our simulated households have a second earner working full-time during the entire working life. We therefore use the female income profile derived from an expanded sample in GKSW that does not impose minimum-participation requirements. This allows us to capture empirical patterns of female labor-force participation during the period of study.

Table I: Parameter Values

Parameter	Value	Description	Source
T_0	25	beginning of working life	calibrated
T_R	66	retirement age	calibrated
T_{max}	102	maximum attainable age	DFJ
P_e	see appendix	business cycle transition matrix	estimated
$\theta_v^1 - \theta_v^4$	see appendix	dividend yield process	estimated
$\theta_j^1 - \theta_j^4$	see appendix	return processes for $j \in \{\text{long bond, short bond, equity}\}$	estimated
ρ_x	0.979	labor income persistence	GOS
$\mu_{\eta_1,0}$	0.091	first labor income innovation mean, expansion	GOS
$\mu_{\eta_1,1}$	-0.130	first labor income innovation mean, recession	GOS
$\mu_{\eta_2,0}$	-0.054	second labor income innovation mean, expansion	GOS
$\mu_{\eta_2,1}$	0.066	second labor income innovation mean, recession	GOS
σ_{η_1}	0.325	first labor income innovation std. dev.	GOS
σ_{η_2}	0.001	second labor income innovation std. dev.	GOS
σ_ϵ	0.186	labor income transitory shock std. dev.	GOS
p_1	0.490	labor income mixture probability	GOS
λ	5.5	income tax level	HSV
κ	0.181	income tax progressivity	HSV
k	0.5	employer 401(k) match rate	calibrated
l	0.06	employer 401(k) match limit	calibrated
$s^I_{max,t}$	see appendix	401(k) contribution limit & req. min. distributions	US tax code
M	12	cash-in-advance constraint on liquid wealth	calibrated
δ	0.022	housing stock depreciation rate	BGLV
ϕ	0.060	rent-to-house-price ratio	BGLV
f^h	0.050	housing stock adjustment transaction cost	BGLV
f^d	0.012	mortgage refinancing transaction cost	BGLV
χ	0.969	mortgage amortization speed	BGLV
ι	0.100	mortgage down payment requirement	BGLV
μ_ν	0.012	mean house price innovation	estimated
σ_ν	0.039	house price innovation std. dev.	estimated
Δr_m	0.03	mortgage interest rate spread	calibrated
ρ_ζ	0.922	medical expenditure shock persistence	DFJ
σ_ζ	0.224	medical expenditure persistent shock std. dev.	DFJ
σ_ψ	0.815	medical expenditure transitory shock std. dev.	DFJ
$\bar{m}$	200,000	family medical expenditure cap	calibrated
$\underline{c}$	\$3,957	retirement consumption floor	DFJ
$\tilde{b}$	23.6	bequest utility intensity	DFJ (2016)
$\underline{b}$	369,000	bequest utility intercept	DFJ
$\tilde{u}_s$	0.85	homeownership scaling factor for poor health	calibrated
$I^m_{0.5}$	\$32,000	Social Security 50% tax cutoff for married households	US tax code
$I^m_{0.85}$	\$44,000	Social Security 85% tax cutoff for married households	US tax code
$I^s_{0.5}$	\$25,000	Social Security 50% tax cutoff for singles	US tax code
$I^s_{0.85}$	\$34,000	Social Security 85% tax cutoff for singles	US tax code
α	0.75	Cobb-Douglas utility share of non-housing consumption	BGLV
γ	3.84	relative risk aversion	DFJ
β	0.96	discount factor	calibrated
w_t	see appendix	effective family size scaling factors	calibrated

Notes: The following process from DFJ are presented in the appendix: deterministic retirement pension income profiles, $Y(t;q^i_R)$; health status transition probabilities, $P_g(g^i_{t-1}, i, t; q^i_R)$; mean medical expenditure profiles, $m(g^i_t, i, t; q^i_R)$; medical expenditure volatility profiles, $\sigma(g^i_t, i, t; q^i_R)$; survival probabilities, $\pi(g^i_t, i, t; q^i_R)$. Additional parameter values are as listed in the text of Section 2.

income summary statistics turned 25 in 1957, and because GKSW only report summary statistics through age 55, we must perform some interpolation to construct $\bar{y}_t^i = \bar{y}(t; i, q_0^i, k)$. For the 25-32 age range, we simply use the data that GKSW report for the 1957 cohort, adjusted by the average annual rate of real wage growth during the 1950s.[24] For the 56-65 age range, we use the predicted values obtained from regressing the GKSW income data between ages 25 and 55 on a second-order polynomial of age. Finally, because the GKSW figures represent the experiences of a single cohort (and therefore incorporate common business-cycle-related shocks that we already capture with the stochastic processes in ((3)-(6)), we smooth the age 25-55 values using a fourth-order polynomial of age.

The tax progressivity parameter κ comes directly from HSV, and we set the shift parameter λ so that tax liability becomes positive at the same real income level as in HSV.

2.3 Retirement Accounts and Cash-in-Advance

We set $k = 0.5$ and $l = 0.06$ in the employer-matching rule (13) to reflect the common employer policy of matching 50% of contributions up to 6% of income. As reflected in (13), we divide k by 2 on the assumption that only one spouse has access to employer matches at any given time, yielding an effective match rate of 0.25. We set the 401(k) contribution limit $s^I_{max,t}$ for $t < T_R$ in equation (14) to $\$20,603$ reflecting the highest IRS limit that our target cohort experienced during their working lives (in 1989), adjusted for inflation and multiplied by 2 to account for the two-member household.[25]

We calibrate the cash-in-advance parameter M to 12 so that the household must carry the equivalent of one month's worth of consumption in safe, liquid wealth.

2.4 Housing

We take parameter values for the housing part of the model directly from BGLV, except for two small changes. First, we estimate the μ_ν and σ_ν parameters governing the house price process using the housing capital gain series in Jorda et al. (2019) since 1985. Second, we add the mortgage interest spread Δr_m and set it to 0.03.

[24] We obtain real wage growth rates from Peake and Vandenbroucke (2020). Specifically, we use the composition-adjusted annual growth rates for white men and white women for the 1950-1960 period, reported in their Figure 5.

[25] In reality, the maximum contribution rises each year but we have a real not a nominal model.

2.5 Retirement Period

We take all parameter values for the pension income, health status, medical expenditures, and mortality processes directly from DFJ. The tax liability parameters $I_{0.5}$ and $I_{0.85}$ in (39) and the required minimum distributions ($s^I_{max,t}$ for $t \geq T_R$ in equation (14)) are set based on IRS regulations.

We determine q^i_R as follows. For each initial quantile q^i_0, we simulate paths for income during the working life from the earnings process (described by equations (3)-(6)) and combine to obtain a simulated distribution of cumulative, working-life incomes. At retirement, we compare each individual's lifetime earnings, $\sum_{t=T_0}^{T_R-1} y^i_t$, to this distribution to determine their lifetime earnings quantile, q^i_R.

We take our process for pension income from DFJ, who estimate the age path of the combined value of "Social Security benefits, defined-benefit pension benefits, annuities, veteran's benefits, welfare, and food stamps." We also use their estimated processes for health shocks and medical costs. We set $Y(t; q^i_R)$ and the medical expense processes based on the permanent income quantile closest to q^i_R for each individual.[26]

2.6 Objective Functions

Following DFJ, we set both coefficients of relative risk aversion (γ for the consumption utility function and γ_b for the bequest utility function) to 3.84. We also use DFJ's estimated values for the bequest utility parameters $\tilde{b}$ and $\underline{b}$ (adjusting the latter for inflation).[27] Finally, we set the discount factor β to 0.96.

To construct our age profile of family size, we follow Salcedo, Schoellman, and Tertilt (2012), who use decennial census data to compute average family sizes for 5-year age buckets (see their Figure 3). We set our family size equal to 2 at age 65, then apply the average growth rates implied by the estimates in Salcedo et al. (2012) to fill in our profile back to age 25. w_t is the square root of this average family size.

For household utility benefits of home ownership and costs during retirement, the former is calibrated to make owning more attractive in a world with attractive retirement

[26]While DFJ uses permanent income quintiles, DFJ provided us with their processes at each fifth percentile, and we choose the DFJ quantile in $\{0, .05, .1, ..., .95, 1\}$ that is closest to each q^i_R. These processes are in our supplemental materials.

[27]De Nardi, French, and Jones (2016a) find somewhat different parameter values for the bequest function for risk aversion closer to 3. De Nardi, French, and Jones (2016b) discuss how the parameters of the bequest function are not well identified due to precautionary saving motives, at least when identification comes primarily from the structural model and data on saving choices.

saving options (not in BGLV) and the second is our choice for realism. We set the homeownership scaling factor $\tilde{u}_s$ to 0.85, so that effective housing consumption is reduced by 15% if the household is a homeowner and one member is sick, and by 28% if the household is a homeowner and both members are sick.

3 Solution Method

We solve the model using a forward-simulation method that avoids grid-based value function iteration in a high-dimensional state space with both continuous and discrete controls. We approximate the household's decision rules with neural networks and evaluate expected lifetime utility by Monte Carlo simulation (sample averaging) over simulated life histories. All expectation objects in the algorithm, including the policy objective and the continuation-value targets, are computed from simulated paths rather than analytical integration or quadrature. We report evidence that our solution is accurate and robust in Appendix C.

Part of our computational method relies on standard machine-learning components: (i) feedforward neural networks for the continuous policy and continuation-value networks; (ii) policy optimization by separable natural evolution strategies (SNES) (Schaul, Glasmachers, and Schmidhuber, 2011; Wierstra, Schaul, Glasmachers, Sun, Peters, and Schmidhuber, 2014), which performs gradient ascent in parameter space using only simulated objective evaluations; and (iii) exploration and regularization devices (entropy bonuses for portfolio weights and ϵ-greedy exploration over discrete housing actions).

To make our economic environment tractable, our method also incorporates three algorithmic devices that we develop. The first is discrete/continuous action separation, which isolates the non-smooth housing/tenure choice from the smooth continuous-allocation problem. The second is projection-based feasibility, which maps the network's softmax outputs, shares that are positive and sum to one, into an economically admissible action set that respects within-period constraints, such as the cash-in-advance constraint, non-negative mortgage balances, and regime-specific accounting identities. Finally, performance-guided moving-average optimization stabilizes learning by selecting the best convex combination of pre- and post-update parameters based on out-of-sample simulated lifetime utility. Together, these three strategies address the key obstacles in solving this model—mixed discrete/continuous controls, tight within-period feasibility restrictions, and high-variance long-horizon simulation objectives —and make it feasible to solve lifecycle problems that would be computationally prohibitive with grid-based methods.

3.1 Algorithm

Discrete/continuous action separation. Our first algorithmic device separates the household's discrete and continuous choices. Recall that Ξ_t denotes the period-t state vector (continuous states such as liquid wealth and return predictors together with discrete indicators such as ownership, the business cycle, retirement, health, and survival). We write the household's control $\mathcal{A}_t = (\mathcal{A}_{d,t}, \mathcal{A}_{c,t})$, where $\mathcal{A}_{d,t}$ is a discrete housing/tenure decision and $\mathcal{A}_{c,t}$ collects continuous controls (consumption, saving into liquid and retirement accounts, housing expenditure, and portfolio shares). The household's objective is expected discounted lifetime utility $V^{\pi}(\Xi_{T_0})$ as defined in Section 1.

We approximate the policy functions for the continuous controls by a neural network with parameters θ, denoted π_θ. Along a simulated life history in which the household follows π_θ, let $G_t^{\pi_\theta}$ denote the discounted sum of utilities realized from age t onward,

$$G_t^{\pi_\theta} \equiv \sum_{s=t}^{T_{max}} \beta^{s-t} \left(\mathbf{1}_{\{w_s>0\}}\, U_s + \mathbf{1}_{\{w_s=0,\, w_{s-1}>0\}}\, U_b(b_s) \right), \tag{48}$$

where U_s is flow utility at age s and U_b is utility from bequests, (42). Both spouses survive to retirement with probability one, so the indicators only matter after T_R.

We approximate the value of each discrete choice by a second network with parameters θ',

$$Q_{\theta'}(\Xi_t, \mathcal{A}_{d,t}) \approx \mathbb{E}\left[G_t^{\pi_\theta} \mid \Xi_t, \mathcal{A}_{d,t} \right], \tag{49}$$

the expected value of making that choice at Ξ_t, with continuous choices in every period given by π_θ (conditioned on the period's discrete choice) and discrete choices from $t+1$ on made by maximizing $Q_{\theta'}$, subject to the feasibility screening described below. Each period proceeds as follows:

1. Compute $Q_{\theta'}(\Xi_t, \mathcal{A}_{d,t})$ for each candidate discrete option $\mathcal{A}_{d,t}$ and select the maximizer, subject to a feasibility screening. If the selected option would leave the household without positive resources, we replace it with one that frees up resource, so an owner sells their house and a renter who would have bought continues to rent. The discrete decision takes four values for owners (keep, refinance, adjust the housing stock, and sell and become a renter) while for renters only the buy–rent margin is

relevant.

2. Condition the continuous policy on the realized discrete decision by appending a one-hot encoding of $\mathcal{A}_{d,t}$—a vector of indicator variables for the realized discrete choice—to the policy input, and compute continuous allocations and portfolio shares via π_θ.

3. Map policy outputs to economically meaningful actions and enforce within-period constraints (budget, cash-in-advance, non-negative mortgage balance) via a projection operator described below and detailed in Appendix A.

This structure isolates non-smooth discrete decisions from the smooth continuous policy while still allowing continuous choices to depend on the discrete regime.

Projection-based feasibility. Our second device enforces within-period feasibility. The policy network outputs a softmax-normalized share vector over five uses of within-period resources, but regime-specific accounting restrictions and occasionally-binding constraints require an additional feasibility correction. Let $z_t \in \mathbb{R}^5$ denote the tentative share vector after discarding inactive components and renormalizing within the realized discrete regime, and let W_t denote the resources available for allocation within the period: cash on hand, the retirement-account balance, and the cash flows implied by the housing decision. We impose feasibility by projecting z_t onto a regime- and state-dependent bounded simplex,

$$x_t \;=\; \arg\min_{x \in \mathbb{R}^5} \; \tfrac{1}{2}\|x - z_t\|_2^2 \quad \text{s.t.} \quad \mathbf{1}^\top x = 1, \;\; \underline{x}_t \le x \le \overline{x}_t. \tag{50}$$

Under feasibility ($\mathbf{1}^\top \underline{x}_t \le 1 \le \mathbf{1}^\top \overline{x}_t$), the projection has the closed form

$$x_{t,i} = \min\big\{ \max\{z_{t,i} - \lambda_t,\ \underline{x}_{t,i}\},\ \overline{x}_{t,i} \big\}, \qquad i = 1, \ldots, 5,$$

where the scalar λ_t is chosen so that $\sum_i x_{t,i} = 1$. The bounds $\underline{x}_t, \overline{x}_t \in [0,1]^5$ encode component activity ($\overline{x}_{t,i} = 0$ for inactive controls) and upper bounds implied by within-period constraints. In particular, for a household holding a positive retirement-account balance, the cash-in-advance constraint (19) is implemented as an upper bound on the consumption share,

$$x_t^c \le \min\left\{1,\ \frac{M\, a^L_{1,t}}{W_t}\right\}, \tag{51}$$

and the mortgage-principal share is capped so that the implied next-period mortgage balance is non-negative. Appendix A provides computational details and the additional downpayment/mortgage coupling used in purchase regimes.

Policy optimization by separable natural evolution strategies. Our goal is to choose policy parameters θ that maximize expected discounted lifetime utility. For any fixed parameter vector θ, we evaluate the objective by forward simulation: we draw shocks, simulate life histories under the induced continuous policy and discrete-choice rule, and average realized lifetime utility across paths. Expected lifetime utility is the expectation of this realization at the start of life, $V^{\pi_\theta}(\Xi_{T_0}) = \mathbb{E}\left[G_{T_0}^{\pi_\theta}\right]$.

Standard backpropagation-based policy gradients can be fragile in our environment because (i) discrete actions and states make the objective non-differentiable and (ii) gradients through long stochastic horizons have high variance. We therefore optimize θ using an alternative commonly used in machine learning, a separable natural evolution strategy (SNES), which can be interpreted as a policy-gradient method in the parameter space. SNES maintains a Gaussian search distribution over policy parameters,

$$\theta \sim \mathcal{N}(\mu_\theta, \mathrm{diag}(\sigma_\theta^2)), \tag{52}$$

and chooses μ_θ and σ_θ to maximize $\mathbb{E}_\theta\left[V^{\pi_\theta}(\Xi_{T_0})\right]$. Although $V^{\pi_\theta}(\Xi_{T_0})$ is not differentiable in θ, this objective is differentiable in μ_θ and σ_θ, with

$$\nabla_{\mu_\theta}\mathbb{E}_\theta\left[V^{\pi_\theta}(\Xi_{T_0})\right] = \mathbb{E}_\theta\left[\frac{\theta-\mu_\theta}{\sigma_\theta^2}\,V^{\pi_\theta}(\Xi_{T_0})\right], \tag{53}$$

where the division is elementwise. Evaluating this gradient therefore requires no derivatives of V^{π_θ} with respect to θ, only unbiased estimates of $V^{\pi_\theta}(\Xi_{T_0})$ at sampled parameters, which we obtain by averaging realized lifetime utility $G_{T_0}^{\pi_\theta}$ across simulated life histories. The resulting rank-based updates of μ_θ and σ_θ are standard and are reproduced in Appendix B; we use paired (antithetic) parameter perturbations and evaluate all population members on the same shock realizations to reduce simulation noise (e.g. Salimans, Ho, Chen, Sidor, and Sutskever, 2017; Brockhoff, Auger, Hansen, Arnold, and Hohm, 2010).

Discrete-choice valuation by Monte Carlo regression. We fit $Q_{\theta'}$ by standard supervised regression to $(\Xi_t, \mathcal{A}_{d,t}, G_t^{\pi_\theta})$ tuples drawn from simulated trajectories under the current policy and discrete-choice rule, using a robust Huber loss.

For many households and states, the values of owning and renting are nearly equal, so that the gaps between discrete choices are small relative to the noise in the single-path continuation values $G_t^{\pi_\theta}$, and the estimated $Q_{\theta'}$ can select the housing tenure inconsistently. We therefore perform an additional step on $Q_{\theta'}$ to make our numerical procedure more robust. For a sample of simulated states, we compute the realized continuation value of each discrete choice by imposing that choice in the current period and following the current policy functions thereafter, using the same shock realizations for every choice. We then refit $Q_{\theta'}$ to these values. We repeat this step periodically during training and more intensively near convergence, and we retain each refit only if it does not lower expected lifetime utility in a held-out simulation.

Performance-guided moving-average optimization. Our third device stabilizes the parameter updates to improve convergence. Training alternates between (i) policy improvement steps (SNES updates of μ_θ and σ_θ) holding the Q-function fixed and (ii) Q-function regression steps (gradient-based updates of θ') holding the policy fixed. After each update phase, we stabilize training using an explicit *performance-guided moving-average* step: we evaluate a small grid of convex combinations of the pre-update and post-update parameter vectors,

$$\theta^{(\omega)} = \omega\, \theta^{\text{old}} + (1-\omega)\, \theta^{\text{new}}, \qquad \omega \in \Omega, \tag{54}$$

and retain the weight ω that maximizes out-of-sample simulated lifetime utility.[28] We apply the same mixing rule to the Q-function parameters and select the mix that maximizes lifetime utility under the current policy. This reward-based damping is an efficient alternative to fixed-coefficient exponential moving averages or target networks: it preserves a single policy representation but smooths unstable parameter updates in long-horizon environments. As an additional safeguard, if an update lowers the evaluation objective, we revert to the previous parameter state and reduce the learning rate.

Exploration and regularization. We use two standard exploration mechanisms during training. First, we add an entropy bonus for portfolio allocations to the period utility to encourage diversified holdings early in training (Williams and Peng, 1991; Mnih, Badia, Mirza, Graves, Lillicrap, Harley, Silver, and Kavukcuoglu, 2016). Second, we employ

[28]We use $\Omega = \{0, 0.5, 0.9, 0.99, 0.999\}$, so the evaluation selects between the raw post-update parameters and increasingly damped updates.

ϵ-greedy exploration in discrete housing choices: with probability ϵ we override the greedy Q-maximizing discrete decision with a random feasible action (Mnih, Kavukcuoglu, Silver, Graves, Antonoglou, Wierstra, and Riedmiller, 2013). Both the entropy coefficient and ϵ are decayed over training, so that evaluation is conducted under the deterministic greedy rule without exploration incentives. Finally, to reduce churning in tenure status, we include a (decaying) penalty for changing ownership status between periods.

3.2 Network Architecture

Inputs and normalization. Both π_θ and $Q_{\theta'}$ are fully-connected feedforward neural networks. Rather than estimating a separate network for each age, we treat time t as an explicit input (normalized by $T_{\max}$) and share parameters across ages. This parameter-sharing structure, motivated by multitask learning (Caruana, 1993) and approximate dynamic programming ideas (Tsitsiklis and Van Roy, 2001), reduces the number of free parameters and improves the signal-to-noise ratio of simulation-based policy updates: a small perturbation to a network parameter affects decisions at many ages, increasing the magnitude of fitness differences used by SNES without proportionally increasing Monte Carlo noise.

The policy input contains (i) a vector of continuous state variables and (ii) one-hot encodings for discrete state components and for the realized discrete choice. We normalize time by $t/T_{\max}$. The remaining continuous variables are normalized using running first and second moments computed from simulated state realizations (time is excluded from these running moments). Discrete indicators are not normalized. These moments are refreshed from the current simulation batch during an initial burn-in and then held fixed thereafter.

Policy network for continuous controls. The policy network outputs three softmax-normalized vectors:

1. Wealth allocation shares over five within-period uses of W_t: consumption, liquid saving, retirement saving, housing expenditure, and mortgage principal changes.[29]

2. Portfolio weights for liquid accounts over the three traded assets.

[29] When the household chooses to rent, the optimal Cobb-Douglas share is imposed between housing and non-housing expenditures.

3. Portfolio weights for retirement accounts over the same three traded assets.

Because the relevant set of active controls depends on the discrete choice (e.g., no mortgage adjustment when renting), the raw shares are re-normalized by regime and then projected via (50) onto a regime-specific feasible set.

Q-network for discrete choices. The Q-network takes as input the state Ξ_t together with a one-hot representation of a candidate discrete decision $\mathcal{A}_{d,t}$ and returns a scalar approximation to this value. Discrete decisions are chosen by evaluating $Q_{\theta'}(\Xi_t, \mathcal{A}_{d,t})$ for all candidate actions and selecting the maximizer, subject to feasibility screening.

3.3 Implementation and Convergence

The entire simulation–learning loop is implemented in JAX and compiled end-to-end. For a given training configuration, the full forward simulation together with the parameter updates is compiled into a single program that runs on the accelerator, as in Hessel, Kroiss, Clark, Kemaev, Quan, Keck, Viola, and van Hasselt (2021). In the baseline configuration, each SNES generation evaluates a population of policy parameter candidates on a large Monte Carlo batch, and Q-function updates use large simulated batches and an Adam optimizer with gradient clipping. We monitor the estimated lifetime utility on a fixed evaluation batch after each update and stop training when this evaluation criterion fails to improve beyond a small threshold for several consecutive updates.

After convergence, we briefly re-optimize the policy functions with an added penalty on violations of the consumption Euler equation, applied after age 30 but weighted most heavily at the transition into retirement, where income and the tax treatment of retirement account withdrawals change discontinuously. We adjust the penalty weight until the residual at the retirement transition is in line with the residuals at other ages, subject to preserving expected lifetime utility in held-out simulations. If the refined policy functions nonetheless deliver lower utility than the original ones, we discard the refinement.

For each permanent-income percentile and discount factor, we solve the model four times, varying the seed that controls the network initializations, the optimizer's random draws, and the simulated histories used in training. Each figure displays, at each age, the mean across these independent solutions together with a band of plus and minus one standard deviation across solutions. The welfare losses are computed from 100,000 simulated lifecycles, using the same shock realizations for the policies being compared.

4 Lifecycle portfolio behavior

This section characterizes the optimal portfolio choices over the life implied by the model and shows two main results. First, the average optimal share of retirement wealth invested in equity starts above 80% and declines steadily over the working life. The typical glide path of a TDF conforms broadly to this pattern, but starts higher, declines by less during the first half of working life, and finishes below the optimal share at around 30% several years after retirement. Second, however, optimal portfolios vary significantly depending on other state variables, while TDFs portfolios depend only only on age.

As described in Section 3, our algorithm involves simulating the shocks, state variables, and choices of a large number of households. We use the final simulations at the optimal parameter values (optimal policy functions) to characterize optimal behavior and outcomes. We approximate the population of interest by combining the lifecycles of state variables and actions of households from each of the four percentiles (q) of average income profiles ($\bar{y}^i$). We take a random subset of households from the $q = 25$ group to represent the lowest 37.5% of the permanent income distribution, and appropriately lower number of households for each other percentile so that the $q = 50$ represents the next 25% of the permanent income distribution (between the 37.5th and 62.5th percentiles), the $q = 75$ group represents the 20% of the distribution between the 62.5th and 82.5th percentiles, and the $q = 90$ group represents the top 17.5% of earners.

We characterize optimal behavior with lifecycle figures. Each figure characterized the numerical accuracy of the solution by plotting one standard deviation bands based on the standard deviation across four independent solutions. To interpret dollar amounts, we note again that our model applies to a married, two-earner, couple born in 1924 and amounts are 2013 dollars. We plot choices and outcomes from age 27 (so that initial conditions to not drive some figures) to the end of working life or until 85 (a household can survive until age 102 with very low probability).

4.1 Optimal Behavior

Figure I.a plots the average level of non-housing consumption (c which also excludes medical cost shocks), total consumption expenditures (c plus mortgage and home-buying transaction costs, medical expenses, and rent or rental equivalent (ϕh)), and the con-

Figure I: Average Consumption, Saving Rates, and Wealth by Age

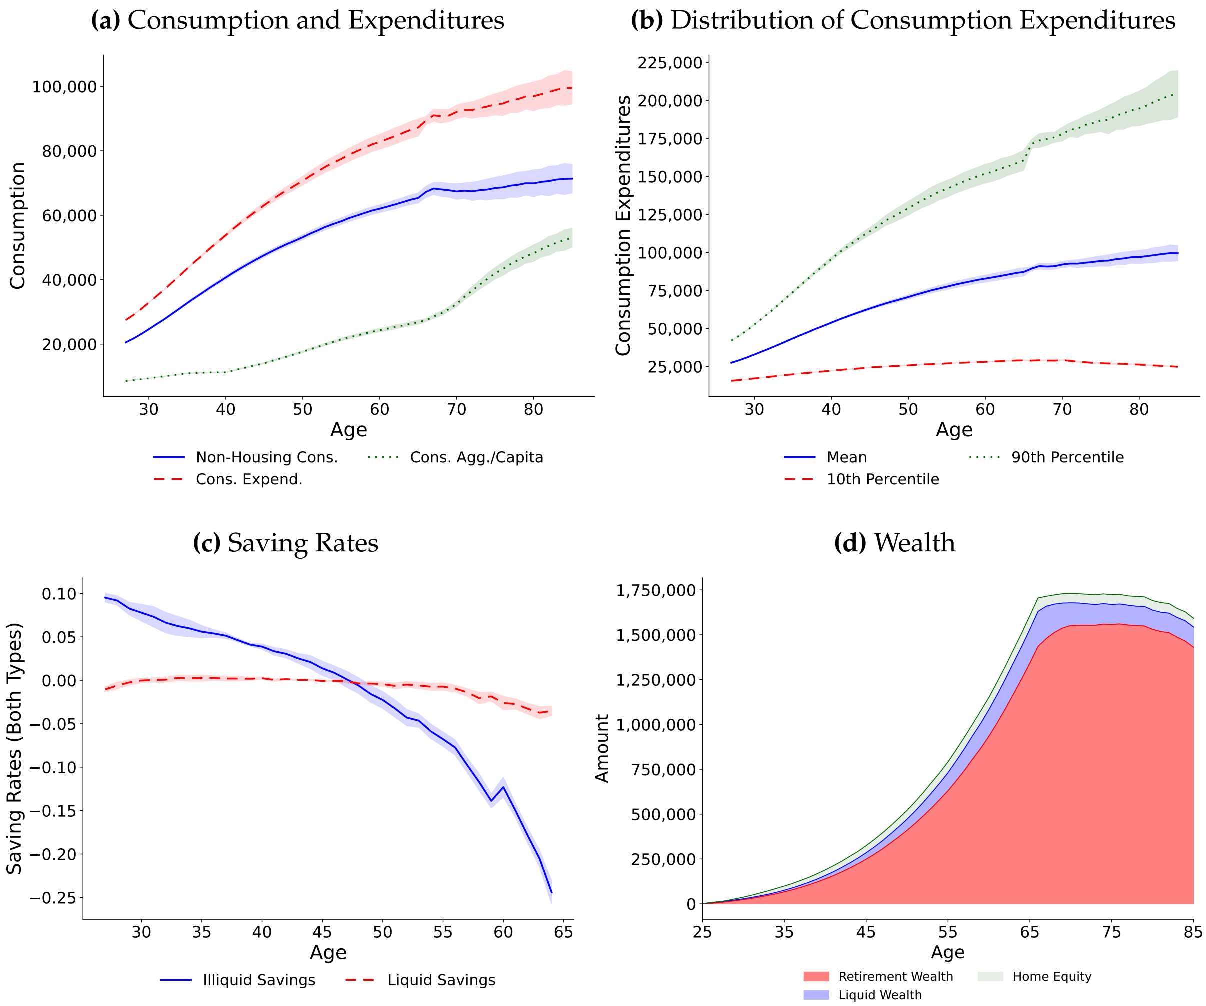

Notes: Consumption expenditures are defined as: non-housing consumption, health expenditures, house and mortgage transaction costs, and rent or rental equivalence (ϕh). The consumption aggregate per capita is the argument of the CRRA period utility function. Saving rates are averages of x^L/Y and x^I/Y dropping family labor incomes below \$10,000.

sumption aggregate per effective family member, $c_t^{\alpha}(\tilde{u}^{o_t(g_t^1+g_t^2)}\phi h_t)^{(1-\alpha)}/w_t$.[30] Because of precautionary saving, increasing family size, high returns relative to discounting, and matching of retirement saving, the first two measures rise rapidly until the mid-40's. Non-housing consumption, c, rises rapidly from an average of about \$20,000 to around \$48,000 at age 45, rises more slowly thereafter to peak at about \$65,000 right before retirement. Total expenditures follow a similar pattern but do not decline in retirement due to medical

[30] We exclude maintenance costs from consumption expenditures since they are part of the rental equivalence. When a household hits the consumption floor, their consumption expenditures are $\underline{c}$.

cost shocks. Again, these figures pertain to people retiring in 1989 and are in 2013 dollars. The measure most closely related to utility, consumption per capita, rises steadily but more slowly throughout working life and then accelerates in old age largely due to medical expenses and in response to the death of one spouse.

Figure I.b shows that cross-household dispersion in consumption expenditures increases throughout the lifecycle. This dispersion represents not just dispersion as one would typically measure it – dispersion across one cohort at a given age – but also across different cohorts because we draw different 'aggregate' shocks for each lifetime that we simulate. In mid-life, the 10th percentile of expenditures is around $25,000 in a year while the 90th percentile is about $120,000. After retirement inequality grows due to medical expenses (as well as mortality).

Turning to saving rates, Figure I.c shows that before age 50, households mainly save in retirement accounts, with retirement saving rates of over 5% until nearly age 40. The average saving rate in retirement accounts turns negative for three reasons: i) some people with high wealth smooth (nonlinear) taxes by starting to withdraw retirement saving before retirement; ii) the retirement contributions of high income households are capped so decrease with income; and iii) high-wealth households who realize low labor income start withdrawing substantial amounts because they have little liquid wealth. Figure I.d shows the resulting average wealth accumulated at each age. On average, households build liquid and illiquid wealth steadily and rapidly during their working lives until age 65, after which they on average consume their returns and pension incomes. At 65, average wealth is around $1.6 million, with most of that in retirement accounts. As noted in the introduction, households accumulate only a small share of wealth in home equity, largely because of the advantageous returns provided by retirement saving.[31]

More importantly, underlying these averages are substantial differences in financial wealth across households at every age. Wealth inequality rises rapidly during working life so that by and during retirement, the 90th percentile of the wealth distribution is more than double the average, at $3.6 million, and the 10th percentile is about one fifth of the average, at around $300,000 (Figure II.a). There is also substantial heterogeneity across households in the fraction of their financial wealth held in retirement accounts. Figure II.b shows that young households first build a little liquidity, and then primarily accumulate

[31]Not shown, the share of households that are homeowners rises from under 10% at age 30 to about 30% by age 45 and remains near that level through retirement. Average mortgage debt conditional on home ownership is roughly $250,000 to $280,000 in the years before retirement.

Figure II: Financial Wealth by Age

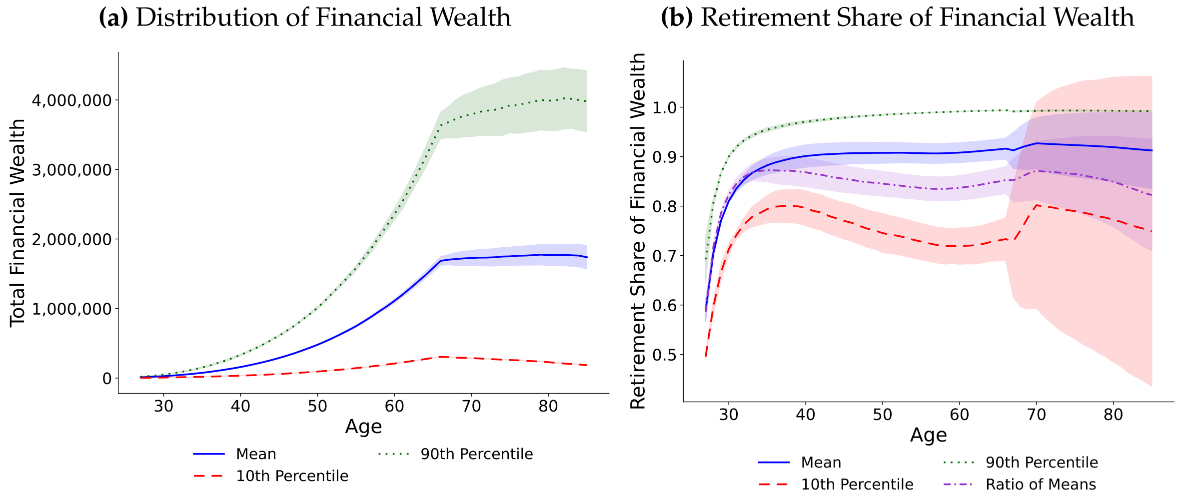

Notes: Financial wealth is liquid wealth plus retirement wealth. Shaded bands span plus and minus one standard deviation across four independent solutions.

retirement wealth so that from ages 35 to 75, households average more than 80% of their financial wealth in retirement accounts. However, the top 10 percent of households hold nearly all their wealth in retirement accounts at all ages above 35, while the bottom ten percent hold roughly 80% of their financial wealth in retirement accounts.

Turning to our main focus, portfolios, the average optimal equity share in total financial wealth is hump-shaped over the lifecycle, but the average optimal equity share in retirement accounts echoes the pattern of typical advice that is delivered by Target Date Funds. Figure III.a shows that the share of financial wealth optimally invested in the stock market at the very beginning of life is about 60%, but rises within 5 years to 70%. The initial low level is driven entirely by the need for liquidity. After reaching 70%, the average equity share declines slowly and steadily going below 50% around retirement age.

In retirement accounts, the average optimal share of stock market investment starts higher – above 80% – and declines steadily with age during the working life, to just under 50% after retirement. This declining pattern is broadly similar to that prescribed by TDFs, but the decline in TDF stock holding is slower in the first half of working life and more rapid thereafter. Current TDFs typically hold roughly 90 percent of their assets in stocks until roughly age 40 (or 25 years before retirement), at which point the share in stocks declines smoothly to roughly 30 to 40 percent by age 75 (ten years after the retirement

Figure III: Distribution of Equity Shares of Wealth

(a) Equity Share of Financial Wealth **(b)** Equity Share of Retirement Wealth

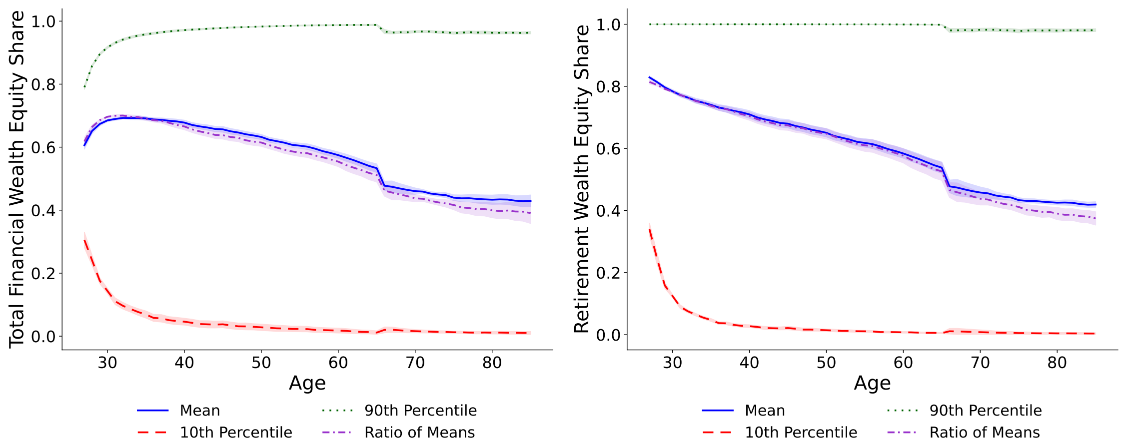

target date). Figure III.b shows that the average optimal share declines by less over the working life and very little after retirement.

However, as there is for other household financial choices, there is a lot of variation in the optimal equity share across households at each age, variation that is far from the homogeneous allocations provided by current TDFs. Figure III.a shows that the 90th percentile of the optimal equity share in financial wealth rises rapidly early in life to reach over 95% in equity at the start of the households' 40's, and remains between 90% and 100% for the remainder of life. For retirement wealth, the 90th percentile of the optimal equity share remains at 100% until just prior to retirement, when it declines slightly to 98%. The 10th percentiles are also far from the average optimal equity share. For retirement wealth, the optimal share invested in equity begins at about a third but declines rapidly in early life, remaining around under 5% from the early 40's onward.

These large differences in optimal portfolio shares arise from differences in the economic circumstances. We now turn to the reasons why optimal equity shares vary across people.

4.2 Optimal Equity Shares of Retirement Wealth by State

This subsection shows how the substantial variation in optimal equity share at each age shown in Figure III relates to differences in economic circumstances. We focus on equity shares in retirement accounts, but Appendix D shows that the differences in equity shares across values of state variables for financial wealth are similar to those for retirement wealth shown in this subsection, just with average lifecycle profiles reflecting Figure III.a

Figure IV: Optimal Equity Shares of Retirement Wealth Over the Business Cycle

(a) Equity Share by State of the Cycle **(b)** Equity Share by Dividend-Price Ratio

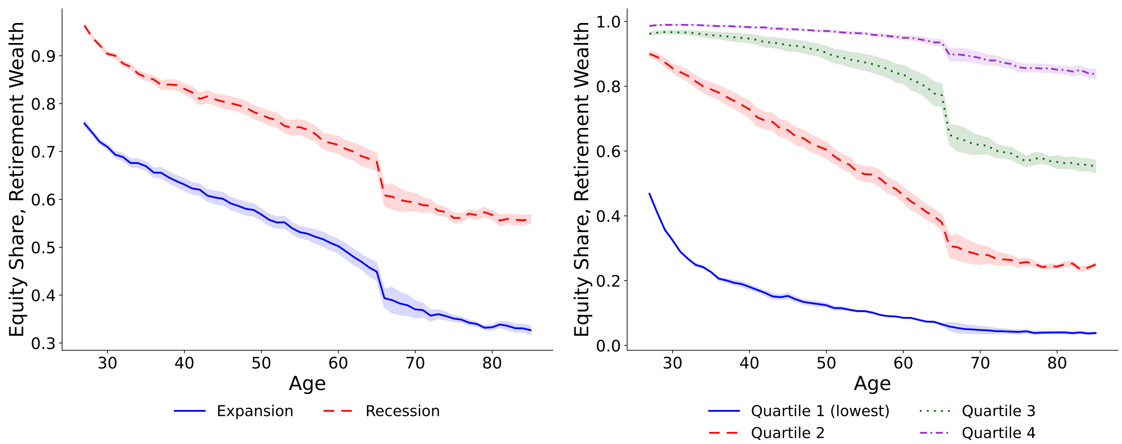

instead of Figure III.b.

Our model implies that the optimal portfolios are quite sensitive to time-variation in expected returns and that this sensitivity is increasing in age. Figure IV.a shows that equity shares are about 20% higher in a recession than an expansion, and this difference is roughly constant across the age distribution. Differences in response to variation in the stock market's payout ratio (correlated with the state of the economy) are also large and differ significantly by age. Figure IV.b shows that when the payout ratio is in the bottom quarter of its distribution – so expected stock returns are low – the share of retirement wealth invested in equity should be about 10% on average for households over 45 and below 5% for households over 65. When the dividend-price ratio is in the top quarter of its distribution, the equity share should be above 95% during working life and roughly 85% in retirement for these households. These effects are large we study an alternative model in which expected returns do not depend on the payout ratio in Section 6.

Figure V.a shows smaller but still significant differences in portfolio allocations associated with differences in the level of total financial wealth. Households with less financial wealth – and so more of their total lifetime wealth in labor income – optimally hold more of their retirement wealth in equity at every age, and the difference grows with age, from about 5 percentage points at age 45 to about 10 percentage points at 65 (these patterns are similar for the equity share of total wealth (Appendix Figure A.4)). Reflecting a similar force but slightly smaller, optimal equity shares are larger for households who earn less

Figure V: Optimal Equity Shares of Retirement Wealth

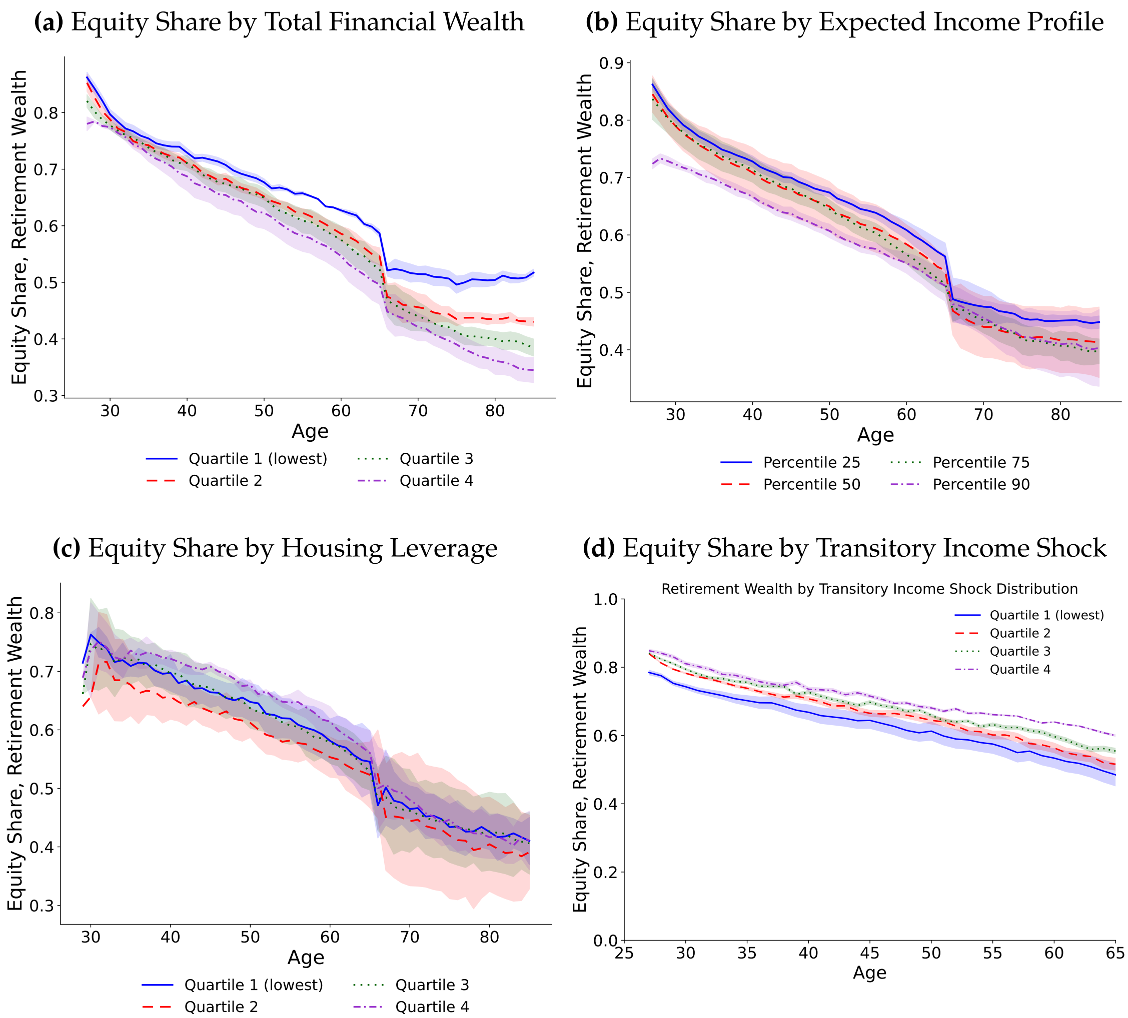

(a) Equity Share by Total Financial Wealth

(b) Equity Share by Expected Income Profile

(c) Equity Share by Housing Leverage

(d) Equity Share by Transitory Income Shock

over their lives, because these low income households also accumulate less wealth and rely more on pensions during retirement (Figure V.b).

In contrast, there are also only small differences in optimal equity share across the housing leverage distribution (Figure V.c). And, as shown in the appendix, optimal equity shares of retirement wealth and of financial wealth do not vary significantly by health or mortality status during retirement (see Appendix Figure A.5). Finally, Figure V.d shows that good transitory shocks to income raise the share invested in stocks. For well-off households, transitory shocks have a small impact on wealth, but have larger effects on liquidity. For households with little wealth, bad shocks raise effective risk aversion and reduce the share of all portfolios invested in equity.

4.3 The Welfare Costs and Benefits of Simple Portfolio Rules

Suppose, for un-modeled reasons, that the household were either not sophisticated enough to make reasonable portfolio choices or not interested in spending the time and effort to make them. How much would the household lose by instead following the simply portfolio rule that depends only on age? Or put differently, how much could be gained by moving to from current TDFs to more sophisticated portfolio strategies?

We begin by considering how much worse the household would do if their retirement saving were allocated to stocks and bonds following the mix prescribed by current TDFs. We consider the following portfolio rule. Up to age 40, retirement wealth is invested 90% in stock, with the remainder allocated between bills and bonds in proportion to the household's unconstrained choices. From age 40, the equity share declines 1.5% per year to reach 60% at age 60, then declines by 2.5% per year to reach 47.5% at age 65, and then declines by 3% per year to reach 32% at age 71. The equity share remains at 30% from age 72 onward.[32]

Table II shows how much a household would hypothetically pay in percent of consumption (in every state at every age) to be able to (costlessly) follow the optimal portfolio rules rather than the TDF rules in retirement wealth.[33] Columns 1 – 5 show these figures from the perspective of the household at age 25, first for the average household and then for households at different points in the distribution of deterministic income profiles. This perspective discounts future flow utility at the discount rate $\beta = 0.96$ per year and so puts more weight on flow utility when young than when old. Columns 6 – 10 report the results of the same calculations but for an equally-weighted average of flow utility across the household's life, as if $\beta = 1$, so that the flow utility of the household when old receives equal weight to the flow utility of the household when young.

The utility loss from following a TDF's prescriptions rather than the optimal state-dependent portfolio rule would cost the average household in expectation the equivalent of losing around 1.1 percent of all their consumption from the perspective of age 25, or between 3.5 and 3.7 percent of consumption in expectation for any year chosen at random.

[32] This pattern is based on the Vanguard Target Retirement series, but other major providers have very similar patterns. That said, some TDFs are more active across asset classes or within asset class and some invest in active underlying funds, so that even different TDFs with the same target date give different returns (Balduzzi and Reuter, 2019) and can have quite different post-fee performance (Shoven and Walton, 2020; Brown and Davies, 2020).

[33] At this level of significant digits, there is essentially no numerical difference across the simulation that we ran.

Table II: Consumption certainty equivalent loss of imposing TDF rules on retirement portfolio

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Age 25 certainty equivalent					Avg. flow utility certainty equivalent				
Other behaviors		Income percentile					Income percentile			
not re-optimized	All hhs	25th	50th	75th	90th	All hhs	25th	50th	75th	90th
None	1.14	0.99	1.27	1.29	1.10	3.50	3.21	3.96	4.02	2.86
Liq. portfolio	1.14	0.99	1.27	1.29	1.10	3.48	3.20	3.93	3.98	2.87
All	1.15	0.99	1.28	1.29	1.12	3.68	3.25	3.96	4.02	3.82

Notes: Average family is a weighted average of the given percentiles as described in the text (the 25th, 50th, 75th, and 90th percentile represent, respectively, 37.5%, 25%, 20%, and 17.5% of the population). The certainty equivalent at age 25 is the percent reduction in consumption at all ages and possible outcomes in the original problem that delivers the same expected present discounted value of utility at age 25. The certainty equivalent for average flow utility is the same calculation with $\beta = 1$. The consumption-equivalent loss is $1 - (V_{\text{rule}}/V_{\text{none}})^{1/(1-\gamma)}$.

We reach this conclusion from several different calculations.

First, we impose the TDF portfolio on retirement wealth and allow the household to re-optimize all other behaviors. The first row of Table II shows that the average household loses the equivalent of only around one percent of consumption from the perspective of age 25. But this under-weights the utility of the household when they are old which is when the effects of an incorrect portfolio choice are the largest. That is, from a welfare perspective, this measures the cost to a 25 year old rather than considering equally the perspectives of the family at other ages. Column 6 shows that when weighting flow utility at different ages equally, the consumption equivalent is a much larger 3.5 percent of consumption at all ages and in all states of the world.

However, this first experiment is in some ways inconsistent. We are assuming that the household is somehow unable or unwilling to optimize its retirement wealth portfolio, but at the same time we allow the household to optimize its portfolio choice for non-retirement wealth conditional on holding a TDF in its retirement account. This assumption may understate the true welfare losses because the household in the model adjusts its portfolio in its liquid accounts to 'unwind' the portfolio imposed by the TDF on their retirement account, while households in reality may not. Thus we experiment with not allowing the household to re-optimize their portfolio behavior in liquid accounts.

Row 2 of Table II shows the welfare costs relative to the unconstrained optimum, of imposing a TDF on retirement account, having the household allocate its liquid wealth among assets as it would in the fully-optimal solution, and allowing it to re-optimize all other behavior given these portfolio rules. Under these assumptions, the cost of imposing the TDF portfolio on retirement wealth is essentially unchanged, at 1.1 percent of consumption from the perspective of age 25 and 3.5 percent under the equally-weighted perspective. There are two reasons for the limited effect of this constraint. First, most households accumulate the majority of their financial wealth in their retirement accounts. Second, the cash in advance constraint keeps low-wealth households from taking a large amount of risk.

An alternative way to impose that the household does not unwind the TDF allocations using liquid wealth is to impose that the household simply uses the optimal unconstrained decision rules for **all** choices other than its portfolio of retirement wealth. That is, we impose the TDF shares in retirement wealth and keep the remaining behaviors of the household for all other decision variables – optimized for the fully-optimal portfolio in the retirement portfolio – the same (conditional on state variables). As shown in the third row of Table II, this leads to a similar total loss of 1.1 percent of consumption from the perspective of a household at age 25, and to a modestly larger 3.7% of consumption for the average household simply averaged across all years of life.

In all three of these measures, the potential welfare gains from moving to optimal portfolio behavior in the retirement account tend to be larger for higher-income households, who are those in the upper middle class. Because high net worth individuals tend to use personal financial advisers and because low-income households have less access to 401k plans (or are offered plans with lower match rates than we have modeled), this is a segment of the population that has leaned heavily on TDFs.

To summarize our findings, while TDFs may lead household to avoid worse mistakes, because they impose the same portfolio on everyone of the same age, there is scope for substantial improvement – 3.5 to 3.7 percent of consumption – from more individualized financial advice or from more customized TDFs.

5 Portfolio Choice with Lower Wealth Accumulation

The previous section evaluates prescriptive portfolio advice and the benefits to further customization in a model in which all household decisions besides portfolio choices are optimal. However, it is reasonable to consider whether households that cannot choose optimal portfolios might make other sub-optimal decisions. In particular there is substantial concern that many households do not save sufficient wealth for retirement. While households in our model consist of two-earner, stable families, they accumulate quite a bit more wealth than the average household in the cohort that we observe one particular draw for in reality. Households in our simulation accumulate on average about $1.6 million in total wealth by age 65 (Figure I.d), while, measured across all types of households, the same cohort in reality accumulated mean net worth of only around $600,000 in 2013 dollars according to the 1989 and 1992 waves of the Survey of Consumer Finances.

In this section, we show that our main conclusions continue to hold in an alternative model in which households make decisions as if they were more impatient and have a discount rate of $\beta = 0.93$, rather than $\beta = 0.96$. While our original specification of time preferences is consistent with the dynamic models that we build upon (e.g. DFJ estimate and use $\beta = 0.97$), there are many other lifecycle models that deliver less saving by many or some households due, for example, to present bias (Laibson, 1997; Laibson, Repetto, and Tobacman, 2007) or some more impatient households (Samwick, 1998; Warner and Pleeter, 2001; Hendricks, 2007).

In this alternative specification, households accumulate about 20% less wealth, but optimal portfolio behavior is largely unchanged, with average optimal equity shares of retirement wealth staying close to those of the baseline model at every age and responding to the same state variables by similar amounts, except for a stronger response to expected income profiles. Thus, when households save less and accumulate less wealth, there are still large benefits to further customization of portfolio shares across broad asset classes.

Figure VI shows that a reduction in the discount factor to $\beta = 0.93$ substantially reduces lifecycle wealth accumulation. Whereas the mean household in the baseline model has about $1.6 million in total financial wealth at the retirement age of 65, the corresponding mean wealth accumulation in Figure VI.a is 20% lower at about $1.2 million. Much of this reduction comes from eliminating a right tail of very high wealth accumulation: while the 90th-percentile line in Figure II.a approaches $4 million during retirement, the 90th percentile in Figure VI.a stays near $3 million.

The equity-share patterns in Figure VII.a show that the average optimal share of retirement wealth invested in the stock market is quantitatively similar to that in the baseline model at every age, declining from 80% at the beginning of working life to 56% at 65 and to about 46% during retirement.

Figure VII.b shows that, as was the case for more patient households, optimal equity

Figure VI: Financial Wealth by Age, with Lower Wealth Accumulation

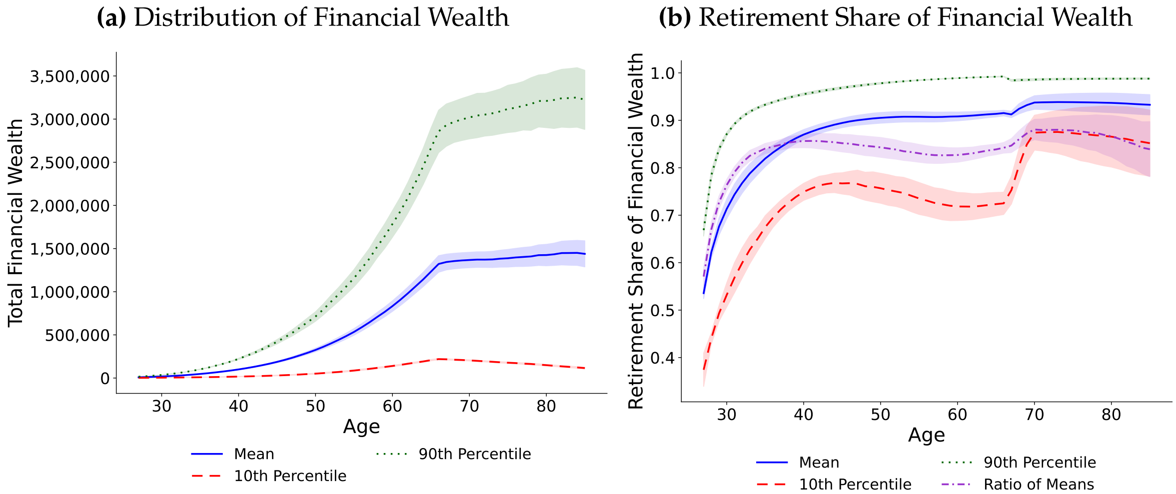

Notes: Financial wealth is liquid wealth plus retirement wealth. Shaded bands span plus and minus one standard deviation across four independent solutions.

shares are higher in recessions. Differences in optimal shares across the wealth distribution are also close to those in the baseline model: households with less wealth optimally hold slightly more equity, with the gap growing from about 1 percentage point at age 36 to about 7 percentage points at 65 (Figure VII.d). Differences in optimal shares across income profiles are larger than in the baseline model during working life, at about 11 percentage points at age 45 versus 6 in the baseline (Figure VII.c).

When households accumulate less wealth, the welfare costs of the imperfect portfolio rules embedded in TDF-type strategies are larger relative to fully-optimal behavior. This difference is driven by two opposing forces. First, since the discount factor β is smaller in the low-wealth model, portfolio-allocation decisions that affect later-in-life consumption levels have a mechanically smaller effect on lifetime utility from the perspective of a discounting household at age 25. The welfare losses shown in columns 1-5 of Table III are therefore generally smaller than the corresponding figures in Table II. However, when average wealth and consumption levels are lower, portfolio choice is more consequential and the constraints imposed by the simple allocation rules have higher flow utility costs in any given time period. The losses to average flow utility in columns 6-10 (where all time periods are treated equally) are thus modestly larger than in the baseline model: while the certainty-equivalent loss of the TDF portfolio for all households is 3.48-3.68% in the

Figure VII: Optimal Equity Shares of Retirement Wealth, with Lower Wealth Accumulation

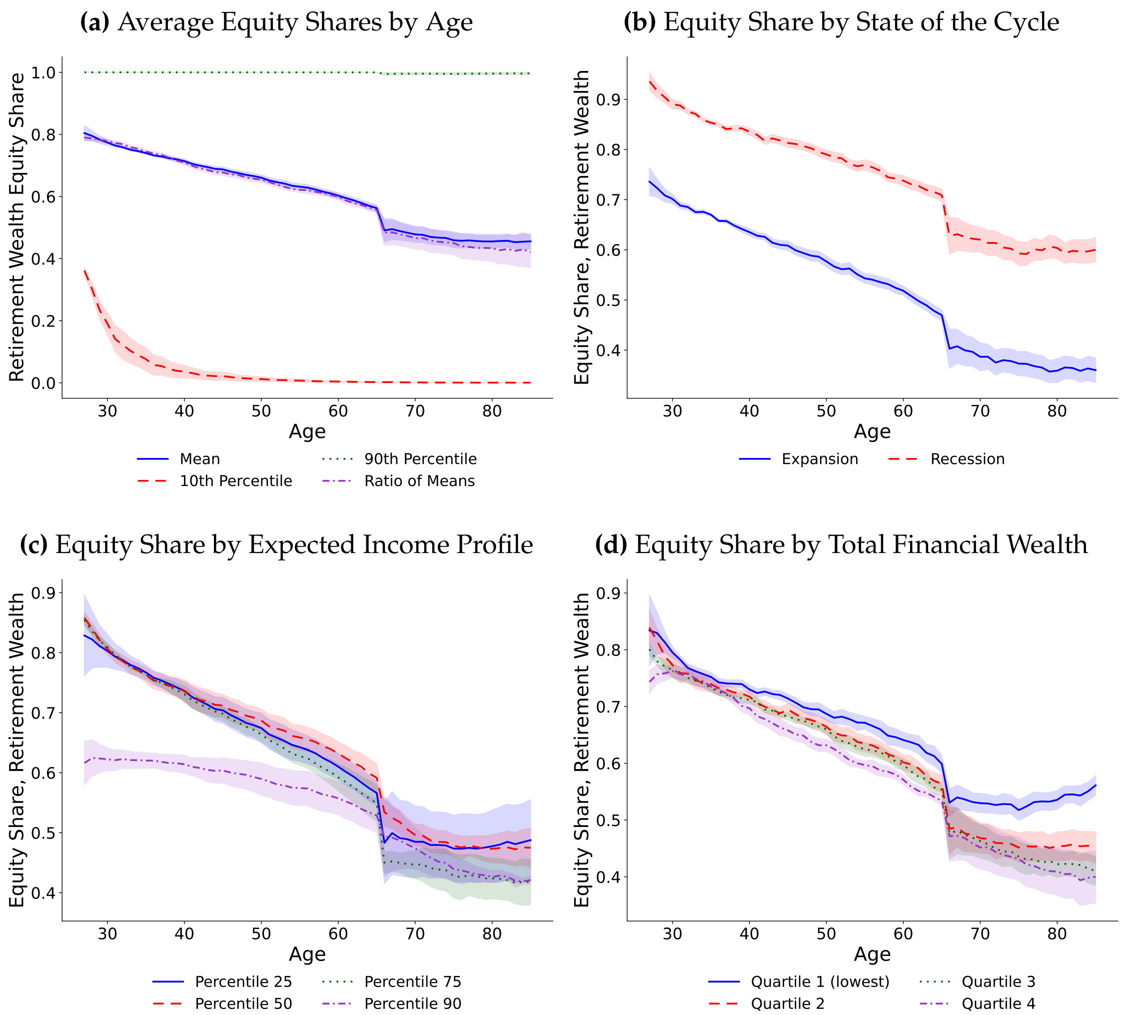

Notes: Data series are only plotted for ages where at least 10% of households have at least $1,000 in retirement wealth. Shaded bands span plus and minus one standard deviation across four independent solutions.

baseline model, the corresponding magnitudes in Table III are slightly larger at 3.67-3.76%.

In sum, when households accumulate less wealth, the average optimal retirement portfolio declines over working life, like a TDF, but still has far more individual-variation, and variation, so that there are still large gains to further individual customization of portfolios across broad asset classes, particularly based on individual investor wealth and income.

Table III: Consumption certainty equivalent loss of imposing TDF rules on retirement portfolio, with lower wealth accumulation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Age 25 certainty equivalent					Avg. flow utility certainty equivalent				
Other behaviors		Income percentile					Income percentile			
not re-optimized	All hhs	25th	50th	75th	90th	All hhs	25th	50th	75th	90th
None	0.47	0.39	0.52	0.59	0.42	3.76	3.53	4.26	4.52	2.66
Liq. portfolio	0.47	0.39	0.52	0.59	0.43	3.75	3.54	4.26	4.52	2.57
All	0.47	0.39	0.52	0.59	0.43	3.67	3.45	4.08	4.52	2.57

Notes: Average family is a weighted average of the given percentiles as described in the text (the 25th, 50th, 75th, and 90th percentile represent, respectively, 37.5%, 25%, 20%, and 17.5% of the population). The certainty equivalent at age 25 is the percent reduction in consumption at all ages and possible outcomes in the original problem that delivers the same expected present discounted value of utility at age 25 (with the discount factor set at $\beta = 0.93$). The certainty equivalent for average flow utility is the same calculation with $\beta = 1$ (and behaviors still re-optimized using $\beta = 0.93$). The consumption-equivalent loss is $1 - (V_{\text{rule}}/V_{\text{none}})^{1/(1-\gamma)}$.

Figure VIII: Financial Wealth by Age, with No Return Predictability

(a) Distribution of Financial Wealth **(b)** Retirement Share of Financial Wealth

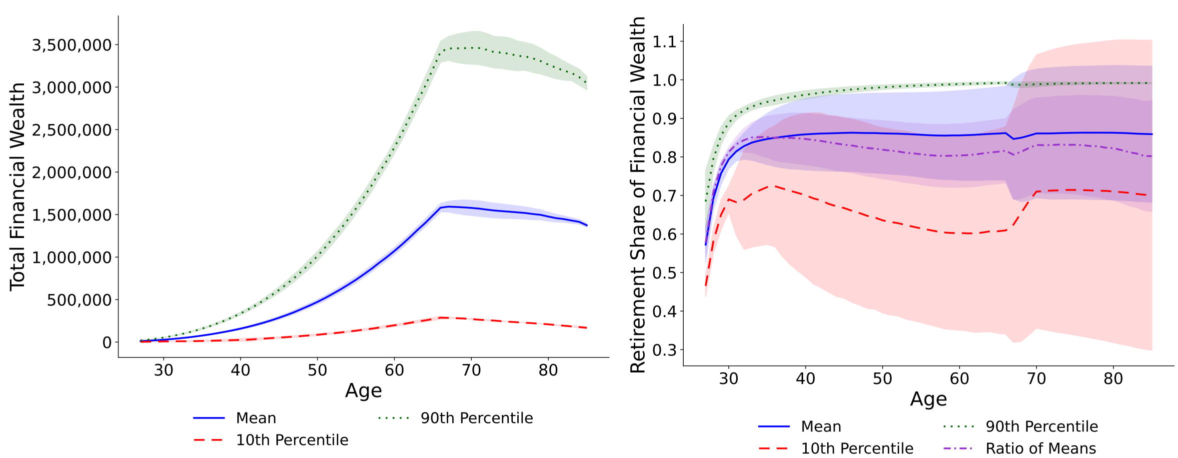

Notes: Financial wealth is liquid wealth plus retirement wealth. Shaded bands span plus and minus one standard deviation across four independent solutions.

6 Portfolio Choice with Little Predictability of Returns

The single largest source of welfare losses in our model is that purely age-dependent rules do not time the market, they do not take advantage of return predictability based on the payout ratio of the stock market. While some predictability is well-established, it's stability is not. In this section, we show that when there is little predictability in returns, TDFs come much closer to the optimum, but also there is more individual-specific variation

Figure IX: Optimal Equity Shares of Retirement Wealth, with No Return Predictability

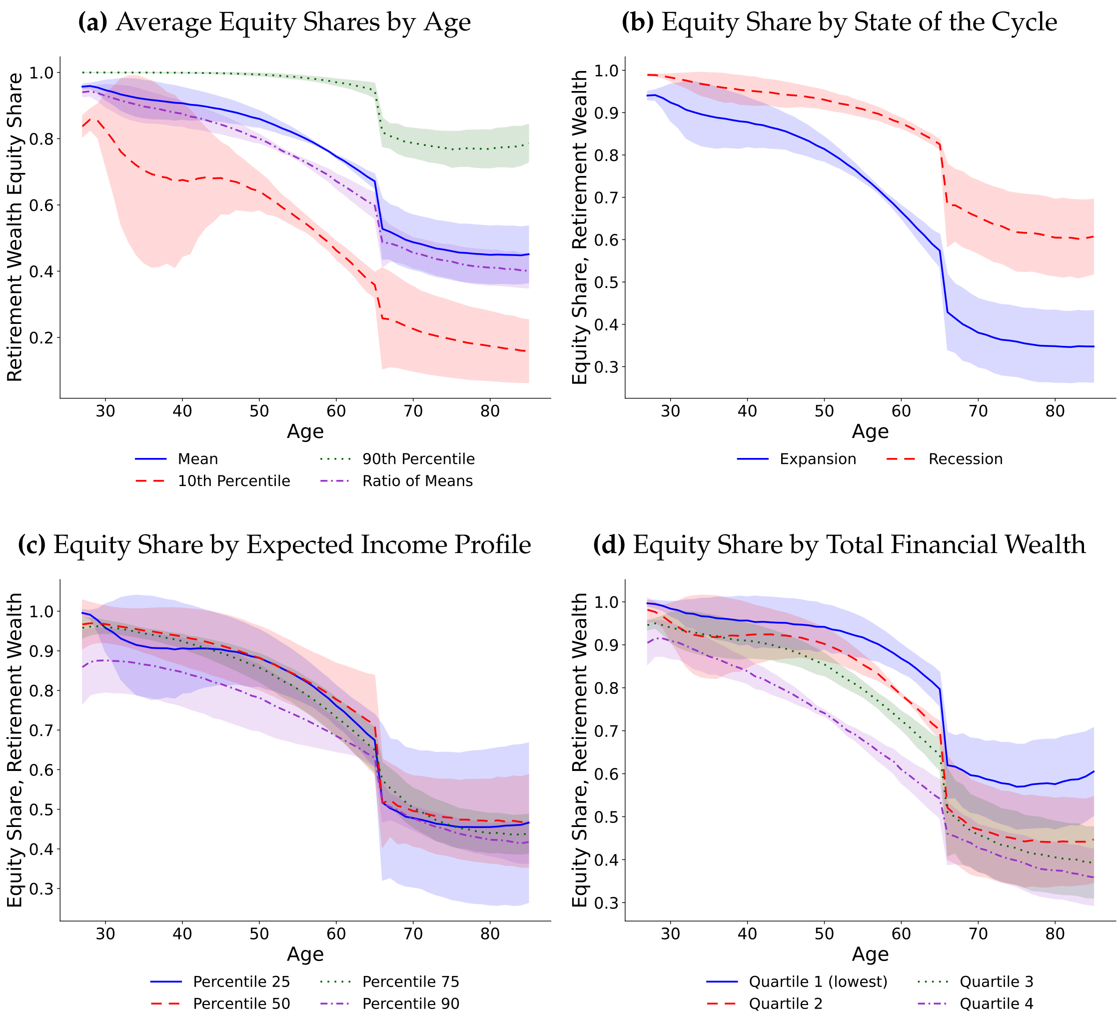

Notes: Data series are only plotted for ages where at least 10% of households have at least $1,000 in retirement wealth. Shaded bands span plus and minus one standard deviation across four independent solutions.

in optimal portfolio shares, particularly by wealth level.

To almost entirely eliminate return predictability in our model economy, we estimate the following process for the log of the gross return on assets

$$\log(R_{j,t}) = \theta_j^1 + \theta_j^2 \Delta e_t^+ + \theta_j^3 \Delta e_t^- + \epsilon_{j,t}. \tag{55}$$

That is, we re-estimate equation (8) with the payout ratio (v) omitted as a predictor, using the same data and sample period as in Section 2. Expected returns still vary with the business cycle, but the payout ratio no longer directly affects expected returns, so there is

Table IV: Consumption certainty equivalent loss of imposing TDF rules on retirement portfolio, with little return predictability

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Age 25 certainty equivalent					Avg. flow utility certainty equivalent				
Other behaviors		Income percentile					Income percentile			
not re-optimized	All hhs	25th	50th	75th	90th	All hhs	25th	50th	75th	90th
None	0.41	0.25	0.51	0.49	0.50	1.23	0.61	1.64	1.55	1.58
Liq. portfolio	0.41	0.25	0.51	0.49	0.50	1.22	0.61	1.62	1.55	1.56
All	0.41	0.26	0.51	0.49	0.50	1.24	0.61	1.63	1.55	1.67

Notes: Average family is a weighted average of the given percentiles as described in the text (the 25th, 50th, 75th, and 90th percentile represent, respectively, 37.5%, 25%, 20%, and 17.5% of the population). The certainty equivalent at age 25 is the percent reduction in consumption at all ages and possible outcomes in the original problem that delivers the same expected present discounted value of utility at age 25. The certainty equivalent for average flow utility is the same calculation with $\beta = 1$. The consumption-equivalent loss is $1 - (V_{\text{rule}}/V_{\text{none}})^{1/(1-\gamma)}$.

little scope for market timing based on valuation ratios.

Figure VIII shows that wealth accumulation is similar to the baseline model: mean financial wealth at retirement is \$1.5 million, and households again hold most of their financial wealth in retirement accounts from their early 30's onward.

Figure IX shows that the average optimal equity share of retirement wealth starts near 96%, remains above 90% until age 40, and declines to 67% at 65 and to just under 50% during retirement, higher in levels than in the baseline model at every age, but displaying a similar pattern. The overall shape of the age profile therefore does not depend on return predictability, but the lack of predictability reduces how much optimal portfolios vary around the average path. The 10th percentile of the optimal equity share is 68% at age 45, compared with 2% in the baseline model. The variation that remains largely comes from the level of wealth and from the business cycle, to which expected returns still respond in this specification. Heterogeneity by wealth is stronger than in the baseline model, with the bottom quartile of the wealth distribution optimally holding 80% of retirement wealth in equity at age 65 and the top quartile 54%.

Table IV repeats the welfare calculations of Section 4.3 for the model without return predictability. Across both welfare measures, the welfare losses from imposing the TDF portfolio on the retirement account are 50 to 65% smaller than in the baseline model, but still meaningful. The cost of imposing the TDF rule falls to 0.41% of consumption from the perspective of age 25 and 1.23% under the equally-weighted perspective.

7 Concluding Remarks

This paper illustrates how recent advances in the algorithms, software, and hardware of machine learning can be used to vastly increase the complexity of problems that economists can solve quantitatively, and therefore enable the field to provide more rigorous quantitative evaluation of the simple rules derived from quantification of parsimonious models. We also illustrate how such an approach can be used to investigate improvements in simple rules by solving for constrained-optimal policy functions over a subset of states.

We apply this idea to a lifecycle problem and study household portfolio choice when the model environment includes many features of reality studied only independently – one or two at a time – in previous work. We find that the simple rules embedded in TDFs approximate the average age-specific optimal portfolio rules of our model, but our results also point to significant benefits of raising the share invested in stock on average later in life and of further customizing the share based on expected returns and individual-level factors.

We see several areas for future work. First, as we noted in the introduction, our model does not allow the household to choose their retirement date and has an unrealistic and exogenous treatment of children, marriage, and divorce. Second, while we model a rich labor income process, one might enrich the model to include more heterogeneity in labor income processes, different exposures of labor income to aggregate returns, or long-run correlation between dividends and labor income (Braxton, Herkenhoff, Rothbaum, and Schmidt, 2021; Guvenen and Smith, 2014; Benzoni, Collin-Dufresne, and Goldstein, 2007).

Third, it would be interesting to also consider a wider choice in financial instruments, such as the choice between traditional and Roth IRA/401k's or the option of various types of annuities or life insurance (as in Yogo, 2016; Koijen et al., 2016).

Finally, our model has no preference heterogeneity. How robust are our results to other realistic assumptions about utility that have been studied in lifecycle models of portfolio choice, such as habits (Gomes and Michaelides, 2003), luxury goods (Wachter and Yogo, 2010), hyperbolic discounting (Angeletos, Laibson, Repetto, Tobacman, and Weinberg, 2001; Love and Phelan, 2015), flow utility from information (Pagel, 2018), risk aversion that declines with age or wealth (Meeuwis, 2019), or incorporating flow disutility from anxiety about future uncertainty as in Epstein-Zin preferences?

References

ANGELETOS, G.-M., D. LAIBSON, A. REPETTO, J. TOBACMAN, AND S. WEINBERG (2001): "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15, 47–68.

AZINOVIC, M., L. GAEGAUF, AND S. SCHEIDEGGER (2022): "Deep Equilibrium Nets," *International Economic Review*, 63, 1471–1525.

BALDUZZI, P. AND J. REUTER (2019): "Heterogeneity in target date funds: strategic risk-taking or risk matching?" *The Review of Financial Studies*, 32, 300–337.

BENZONI, L., P. COLLIN-DUFRESNE, AND R. S. GOLDSTEIN (2007): "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated," *The Journal of Finance*, 62, 2123–2167.

BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2018): "House Prices and Consumer Spending," *Review of Economic Studies*, 85, 1502–1542.

BESHEARS, J., J. CHOI, D. LAIBSON, AND B. MADRIAN (2008): "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States," in *Lessons from Pension Reform in the Americas*, Oxford University Press, 59–87.

BODIE, Z., R. C. MERTON, AND W. F. SAMUELSON (1991): "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model," *Journal of Economic Dynamics and Control*, 16, 427–49.

BOUDOUKH, J., R. MICHAELY, M. RICHARDSON, AND M. R. ROBERTS (2007): "On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing," *The Journal of Finance*, 62, 877–915.

BRAXTON, J. C., K. F. HERKENHOFF, J. L. ROTHBAUM, AND L. SCHMIDT (2021): "Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter," Working Paper 29567, National Bureau of Economic Research.

BROCKHOFF, D., A. AUGER, N. HANSEN, D. V. ARNOLD, AND T. HOHM (2010): "Mirrored Sampling and Sequential Selection for Evolution Strategies," in *Parallel Problem Solving from Nature, PPSN XI*, ed. by R. Schaefer, C. Cotta, J. Kołodziej, and G. Rudolph, Berlin, Heidelberg: Springer Berlin Heidelberg, 11–21.

BROWN, D. C. AND S. DAVIES (2020): "Off target: On the underperformance of target-date funds," *Available at SSRN 3707755*.

CALVET, L. E., J. Y. CAMPBELL, F. GOMES, AND P. SODINI (2021): "The Cross-Section of Household Preferences," *NBER Working Paper, W28788*.

CAMPBELL, J. AND L. VICEIRA (2002): *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.

CAMPBELL, J. Y. (2016): "Restoring Rational Choice: The Challenge of Consumer Financial Regulation," *American Economic Review: Papers & Proceedings*, 106, 1–30.

CAPLIN, A. AND J. LEAHY (2004): "The Social Discount Rate," *Journal of Political Economy*, 112, 1257–1268.

CARUANA, R. (1993): "Multitask Learning: A Knowledge-Based Source of Inductive Bias," in *Proceedings of the Tenth International Conference on Machine Learning*, Morgan Kaufmann, 41–48.

CATHERINE, S. (2021): "Countercyclical Labor Income Risk and Portfolio Choices over the Life-Cycle," *The Review of Financial Studies*, 35, 4016–4054.

CHOUKHMANE, T. (2025): "Default Options and Retirement Saving Dynamics," *American Economic Review*, 115, 3749–3787.

COCCO, J. F. (2005): "Portfolio Choice in the Presence of Housing," *Review of Financial Studies*, 18, 535–567.

COCCO, J. F., F. J. GOMES, AND P. J. MAENHOUT (2005): "Consumption and Portfolio Choice over the Life Cycle," *Review of Financial Studies*, 18, 1502–1542.

CURCURU, S., J. HEATON, D. LUCAS, AND D. MOORE (2010): "Heterogeneity and Portfolio Choice: Theory and Evidence," in *Handbook of Financial Econometrics: Tools and Techniques*, ed. by Y. Ait-Sahalia and L. P. Hansen, London: North-Holland, chap. 6.

DAHLQUIST, M., O. SETTY, AND R. VESTMAN (2018): "On the Asset Allocation of a Default Pension Fund," *The Journal of Finance*, 73, 1893–1936.

DAMMON, R., C. S. SPATT, AND H. H. ZHANG (2004): "Optimal Asset Location and Allocation with Taxable and Tax-Deferred Investing," *Journal of Finance*, LIX, 999–1037.

DAMMON, R. M. AND C. S. SPATT (2026): "Do Target-Date Funds Hit the Target?" Working paper, Carnegie Mellon University.

DAVIS, S. J., F. KUBLER, AND P. WILLEN (2006): "Borrowing Costs and the Demand for Equity over the Life Cycle," *The Review of Economics and Statistics*, 88, 348–362.

DENARDI, M., E. FRENCH, AND J. JONES (2010): "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118, 39–75.

——— (2016a): "Medicaid Insurance in Old Age," *American Economic Review*, 106, 3480–3520.

——— (2016b): "Savings After Retirement: A Survey," *Annual Review of Economics*, 8, 177–204.

DRUEDAHL, J., R. HULEUX, AND J. RØPKE (2026): "Deep Learning Solutions of Large Non-Convex Life-Cycle Models," Working paper, University of Copenhagen.

Druedahl, J. and J. Røpke (2026): "Deep Learning Algorithms for Solving Convex Finite-Horizon Models," Working paper, University of Copenhagen.

Duarte, V., D. Duarte, J. Fonseca, and A. Montecinos (2020): "Benchmarking machine-learning software and hardware for quantitative economics," *Journal of Economic Dynamics and Control*, 111, 103796.

Duarte, V., D. Duarte, and D. H. Silva (2024): "Machine Learning for Continuous-Time Finance," *The Review of Financial Studies*, 37, 3217–3271.

Fagereng, A., C. Gottlieb, and L. Guiso (2017): "Asset Market Participation and Portfolio Choice over the Life-Cycle," *Journal of Finance*, 72, 705–750.

Fernandez-Villaverde, J., S. Hurtado, and G. Nuno (2020): "Financial Frictions and the Wealth Distribution," Working Paper.

Gakidis, H. E. (1998): "Stocks for the Old? Earnings Uncertainty and Life-Cycle Portfolio Choice," Working Paper, M.I.T.

Gomes, F. and A. Michaelides (2003): "Portfolio Choice With Internal Habit Formation: A Life-Cycle Model With Uninsurable Labor Income Risk," *Review of Economic Dynamics*, 6, 729–766.

Gomes, F., A. Michaelides, and V. Polkovnichenko (2006): "Optimal Savings with Taxable and Tax-Deferred Accounts," *Review of Economic Dynamics*, 12, 718–735.

Gomes, F. J., A. Michaelides, and Y. Zhang (2022): "Tactical Target Date Funds," *Management Science*, 68, 3047–3070.

Guvenen, F., G. Kaplan, J. Song, and J. Weidner (2017): "Lifetime Incomes in the United States Over Six Decades," *NBER Working Paper, W23371*.

Guvenen, F., S. Ozkan, and J. Song (2014): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122, 621–660.

Guvenen, F. and A. A. Smith (2014): "Inferring Labor Income Risk and Partial Insurance From Economic Choices," *Econometrica*, 82, 2085–2129.

Heathcote, J., K. Storesletten, and G. Violante (2017): "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 132, 1693–1754.

Hendricks, L. (2007): "How Important is Discount Rate Heterogeneity for Wealth Inequality?" *Journal of Economic Dynamics and Control*, 31, 3042–3068.

Hessel, M., M. Kroiss, A. Clark, I. Kemaev, J. Quan, T. Keck, F. Viola, and H. van Hasselt (2021): "Podracer Architectures for Scalable Reinforcement Learning," *arXiv preprint, arXiv:2104.06272*.

Hu, X. (2005): "Portfolio Choices for Homeowners," *Journal of Urban Economics*, 58, 114–136.

INVESTMENT COMPANY INSTITUTE (2020): *Investment Company Fact Book*, Investment Company Institute, Washington, DC.

JORDA, O., K. KROLL, D. KUVSHINOV, M. SCHULARICK, AND A. M. TAYLOR (2019): "The Rate of Return on Everything, 1870-2015," *Quarterly Journal of Economics*, 134, 1225–1298.

KOIJEN, R. S., S. NIEUWERBURGH, AND M. YOGO (2016): "Health and Mortality Delta: Assessing the Welfare Cost of Household Insurance Choice," *Journal of Finance*, 71, 957–1010.

LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–477.

LAIBSON, D., A. REPETTO, AND J. TOBACMAN (2007): "Estimating Discount Functions with Consumption Choices Over the Lifecycle," National Bureau of Economic Research Working Paper.

LI, J. AND K. SMETTERS (2011): "Optimal Portfolio Choice with Wage-Indexed Social Security," .

LOVE, D. A. AND G. PHELAN (2015): "Hyperbolic discounting and life-cycle portfolio choice," *Journal of Pension Economics and Finance*, 14, 492 – 524.

MALIAR, L., S. MALIAR, AND P. WINANT (2021): "Deep learning for solving dynamic economic models." *Journal of Monetary Economics*, 122, 76–101.

MEEUWIS, M. (2019): "Wealth Fluctuations and Risk Preferences: Evidence from U.S. Investor Portfolios," MIT Sloan Working Paper.

MERTON, R. C. (1969): "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics*, 51, 247–57.

——— (1971): "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.

MNIH, V., A. P. BADIA, M. MIRZA, A. GRAVES, T. P. LILLICRAP, T. HARLEY, D. SILVER, AND K. KAVUKCUOGLU (2016): "Asynchronous Methods for Deep Reinforcement Learning," .

MNIH, V., K. KAVUKCUOGLU, D. SILVER, A. GRAVES, I. ANTONOGLOU, D. WIERSTRA, AND M. RIEDMILLER (2013): "Playing Atari with Deep Reinforcement Learning," .

MORNINGSTAR (2020): "2020 Target-date strategy landscape," *https://www.morningstar.com/lp/tdf-landscape*.

PAGEL, M. (2018): "A News-Utility Theory for Inattention and Delegation in Portfolio Choice," *Econometrica*, 86, 491–522.

PARKER, J. A., A. SCHOAR, A. COLE, AND D. SIMESTER (2025): "Household Portfolios and Retirement Saving over the Life Cycle," *Journal of Finance*, 80, 2739–2787.

PARKER, J. A., A. SCHOAR, AND Y. SUN (2023): "Retail Financial Innovation and Stock Market Dynamics: The Case of Target Date Funds," *The Journal of Finance*, 78, 2673–2723.

PEAKE, M. AND G. VANDENBROUCKE (2020): "Worker Diversity and Wage Growth Since 1940," *Federal Reserve Bank of St. Louis Review*, 102, 1–18.

SALCEDO, A., T. SCHOELLMAN, AND M. TERTILT (2012): "Families as Roommates: Changes in US Household Size from 1850 to 2000," *Quantitative Economics*, 3, 133–175.

SALIMANS, T., J. HO, X. CHEN, S. SIDOR, AND I. SUTSKEVER (2017): "Evolution Strategies as a Scalable Alternative to Reinforcement Learning," .

SAMUELSON, P. A. (1969): "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51, 239–246.

SAMWICK, A. A. (1998): "Discount Rate Heterogeneity and Social Security Reform," *Journal of Development Economics*, 57, 117–146.

SCHAUL, T., T. GLASMACHERS, AND J. SCHMIDHUBER (2011): "High Dimensions and Heavy Tails for Natural Evolution Strategies," in *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation (GECCO '11)*, ACM.

SCHMIDT, L. D. (2025): "Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk," *Journal of Financial Economics*, 172.

SHEN, J. (2021): "Countercyclical Risks, Consumption and Portfolio Choice: Theory and Evidence," Manuscript, University of Missouri.

SHOVEN, J. B. AND D. B. WALTON (2020): "An Analysis of the Performance of Target Date Funds," *NBER Working Paper, W27971*.

SOCIAL SECURITY ADMINISTRATION (2016): "Income of the Population 55 or Older, 2014," Social Security Administration Report: Report Link.

STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): "Consumption and Risk Sharing Over the Life Cycle," *Journal of Monetary Economics*, 51, 609–633.

SUTTON, R. S., D. A. MCALLESTER, S. P. SINGH, AND Y. MANSOUR (2000): "Policy Gradient Methods for Reinforcement Learning with Function Approximation," in *Advances in Neural Information Processing Systems 12*, ed. by S. A. Solla, T. K. Leen, and K. Müller, MIT Press, 1057–1063.

TSITSIKLIS, J. N. AND B. VAN ROY (2001): "Regression methods for pricing complex American-style options," *IEEE Transactions on Neural Networks*, 12, 694–703.

VANGUARD GROUP (2026): "How America Saves 2026: 25th Edition," Industry report, Vanguard Workplace Solutions.

WACHTER, J. (2010): "Asset Allocation," *Annual Reviews of Financial Economics*, 2, 175–206.

WACHTER, J. A. AND M. YOGO (2010): "Why Do Household Portfolio Shares Rise in Wealth?" *Review of Financial Studies*, 23, 3929–3965.

WARNER, J. T. AND S. PLEETER (2001): "The Personal Discount Rate: Evidence from Military Downsizing Programs," *American Economic Review*, 91, 33–53.

WIERSTRA, D., T. SCHAUL, T. GLASMACHERS, Y. SUN, J. PETERS, AND J. SCHMIDHUBER (2014): "Natural Evolution Strategies," *Journal of Machine Learning Research*, 15, 949–980.

WILLIAMS, R. J. AND J. PENG (1991): "Function Optimization using Connectionist Reinforcement Learning Algorithms," *Connection Science*, 3, 241–268.

YAO, R. AND H. H. ZHANG (2005): "Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints," *Review of Financial Studies*, 18, 197–239.

YOGO, M. (2016): "Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets," *Journal of Monetary Economics*, 80, 17–34.

Online Appendix for

"Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle"

Victor Duarte   Julia Fonseca   Aaron Goodman   Jonathan A. Parker

A Projection operator for within-period feasibility

This appendix formalizes the projection step referenced in Section 3. In each period, after selecting the discrete housing/tenure action, the policy network produces a vector of tentative within-period allocation shares. These shares are mapped into feasible economic controls by a regime-specific renormalization followed by a projection onto a bounded simplex.

Pre-projection shares and resources. Let W_t denote the resources available for allocation within the period, net of the transaction costs implied by the discrete housing choice. After regime-specific renormalization, the policy induces a share vector $z_t \in \mathbb{R}^5$ over

$$z_t = \left(z_t^c,\ z_t^L,\ z_t^I,\ z_t^h,\ z_t^{\Delta d}\right),$$

corresponding to non-housing consumption, liquid saving, retirement saving, housing expenditure, and mortgage principal changes. By construction $z_t \geq 0$ and $\mathbf{1}^\top z_t = 1$.

Regime-specific bounds. For each discrete regime we construct lower and upper bounds $\underline{x}_t, \overline{x}_t \in [0,1]^5$ such that feasibility can be expressed as

$$\underline{x}_t \leq x_t \leq \overline{x}_t, \qquad \mathbf{1}^\top x_t = 1.$$

In the baseline code $\underline{x}_t = \mathbf{0}$ and the upper bounds encode: (i) component activity (inactive components have $\overline{x}_{t,i} = 0$); (ii) the cash-in-advance constraint, implemented as an upper bound on the consumption share,

$$x_t^c \leq \min\left\{1,\ \frac{M\, a_t^{L1}}{W_t}\right\},$$

(applied when $a_t^I > 1$ as in the code); and (iii) bounds on the mortgage-principal share that guarantee a non-negative next-period mortgage balance in the keep and refinance regimes.

Boxed-simplex projection. Given $(z_t, \underline{x}_t, \overline{x}_t)$, we compute the feasible share vector x_t as the Euclidean projection onto the bounded simplex,

$$x_t \;=\; \Pi_{[\underline{x}_t, \overline{x}_t]}(z_t) \;\equiv\; \arg\min_{x \in \mathbb{R}^5} \frac{1}{2}\|x - z_t\|_2^2 \quad \text{s.t.} \quad \mathbf{1}^\top x = 1, \;\; \underline{x}_t \le x \le \overline{x}_t. \tag{56}$$

Assuming feasibility ($\mathbf{1}^\top \underline{x}_t \le 1 \le \mathbf{1}^\top \overline{x}_t$), the solution is unique and has the form

$$x_{t,i} = \min\{\, \max\{z_{t,i} - \lambda_t,\ \underline{x}_{t,i}\},\ \overline{x}_{t,i}\}, \qquad i = 1, \ldots, 5, \tag{57}$$

where the scalar λ_t is chosen so that $\sum_i x_{t,i} = 1$. In implementation, we bracket λ_t by $\lambda_t^- = \min_i(z_{t,i} - \overline{x}_{t,i})$ and $\lambda_t^+ = \max_i(z_{t,i} - \underline{x}_{t,i})$ and solve for λ_t by bisection (40 iterations).

Buy-regime coupling. In regimes in which the household acquires housing (moving/adjusting housing stock or buying from renting), feasibility imposes an additional coupling between the housing-expenditure share and the mortgage-principal share. Let $s_t = x_t^h + x_t^{\Delta d}$ and define the mortgage fraction $r_t = x_t^{\Delta d}/s_t$ (with $r_t = 0$ if $s_t = 0$). The code caps r_t at

$$r_{\max,t} = \frac{k_{m,t}}{1 + k_{m,t}}, \qquad k_{m,t} = \frac{(1-\iota)\,\mathrm{coef}_t}{\mathrm{denom\_house}_t},$$

where $\mathrm{coef}_t = 1 - (1 - L'(Y_t))\, r_{m,t}$ and $\mathrm{denom\_house}_t = \delta + \iota + (1-\iota)(1 - L'(Y_t))\, r_{m,t}$ with $r_{m,t}$ the net mortgage interest rate. After capping, $(x_t^h, x_t^{\Delta d})$ is reset to $\big((1 - r_t)s_t,\ r_t s_t\big)$.

Final safety nets. After mapping shares into levels (consumption, savings, housing stock, and next-period mortgage debt), any remaining numerical violations are resolved by transfers to liquid saving: negative mortgage balances are set to zero with the violation amount added to liquid saving, and any remaining cash-in-advance violations are corrected by reducing consumption and increasing liquid saving by the same amount.

B SNES Update Equations

This appendix states the SNES parameter updates used in Section 3.1. Writing the N population draws as $\theta_i = \mu_\theta + \sigma_\theta \odot \varepsilon_i$ with $\varepsilon_i \sim \mathcal{N}(0, I)$, where $\odot$ denotes the elementwise product, the (diagonal) natural-gradient SNES updates take the form

$$\mu_\theta \leftarrow \mu_\theta + \eta_\mu\, \sigma_\theta \odot \sum_{i=1}^{N} w_i\, \varepsilon_{(i)}, \tag{58}$$

$$\sigma_\theta \leftarrow \sigma_\theta \odot \exp\left(\frac{\eta_\sigma}{2} \sum_{i=1}^{N} w_i\, (\varepsilon_{(i)}^2 - 1) \right), \tag{59}$$

where $\varepsilon_{(1)}, \ldots, \varepsilon_{(N)}$ are the perturbations ordered so that candidate (1) has the highest estimated objective, w_i are rank-based recombination weights (with a baseline so that

$\sum_i w_i = 0$), and η_μ, η_σ are learning rates.

C Accuracy and Robustness of the Numerical Solution

This appendix collects evidence that the numerical solution is accurate and robust. We first show that independently trained solutions agree closely. We then measure the direction of any bias from under-optimization, and find that it pushes the reported losses toward zero, thus biasing against our findings. Finally, we show that, on the states where the consumption Euler equation must hold with equality, the final policy functions satisfy it to within one to two percent of one-period consumption at every age, and that residuals of this size can leave at most about 0.1 percent of lifetime consumption unexploited, well below the smallest welfare difference the paper reports.

C.1 Agreement across independently trained solutions

As described in Section 3, for each permanent-income percentile and discount factor we solve the model four times, varying the seed that controls the network initializations, the optimizer's random draws, and the simulated histories used in training. The figures in the paper display the mean across these independent solutions together with bands of plus and minus one standard deviation; the policies and the associated welfare levels agree closely across solutions.

C.2 Optimization slack and the direction of bias

A natural concern is that the reported welfare losses could reflect under-optimization of the benchmark rather than the cost of the rules. To measure the direction and size of any such bias, we deliberately re-solved the median-income $\beta = 0.96$ calibration at reduced optimization budgets – between 12 and 50 percent of the production optimizer updates, and with search populations of 32 and 64 against the production 128 – gave each run the identical downstream treatment, and evaluated the same rules against each run's own baseline. The measured cost of the rules rises monotonically with the quality of the baseline: from 1.00 to 1.34 percent of consumption for the TDF rule across a baseline-quality range of 1.1 percentage points, or roughly 0.3 points of understated rule cost per point of baseline slack. Under-optimization therefore biases the reported losses toward zero, so the welfare costs in Table II are, if anything, understated.

Figure A.1: Measured Cost of Simple Rules and Optimization Slack

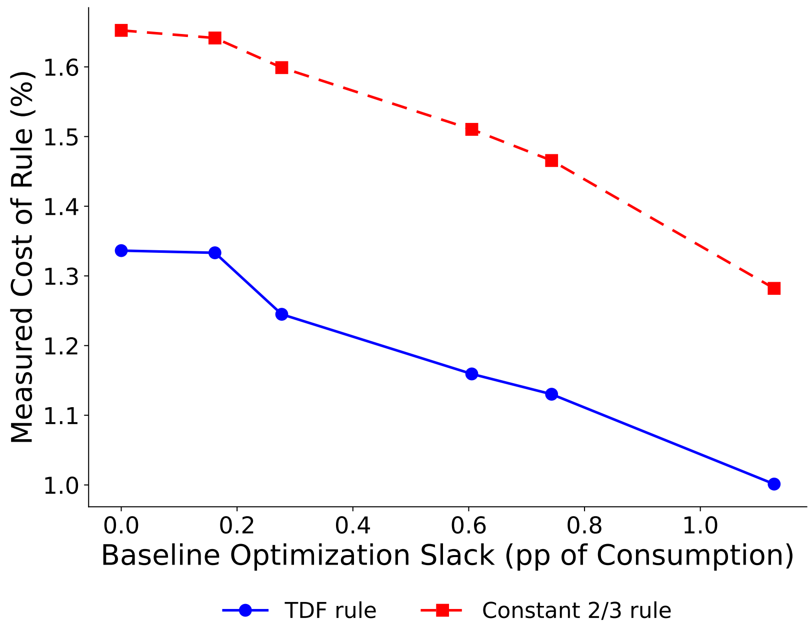

Notes: Each point is one solution of the median-income $\beta = 0.96$ calibration trained at a reduced optimization budget and given the identical downstream treatment. The horizontal axis is the run's baseline slack: the consumption-equivalent shortfall of its baseline relative to the best-achieved baseline. The vertical axis is the consumption-equivalent loss of the rule measured against that same run's baseline, from the perspective of age 25, on 100,000 simulated lives with common shock realizations. All points use the same seed; the tables in the paper use the best of the four independent solutions.

C.3 Euler-equation residuals

Figure A.2 examines the consumption Euler equation age by age. For the households that are in the interior set $\mathcal{I}_t$ at age t, we compute the residual

$$EE(t) \;=\; \beta\,\frac{\mathbb{E}[\,u_c(c_{t+1})\,\nu_{t+1}\,R_{1,t+1}\mid \mathcal{I}_t\,]}{\mathbb{E}[\,u_c(c_t)\mid \mathcal{I}_t\,]} \;-\; 1, \tag{60}$$

where u_c is the model's marginal utility of non-housing consumption, $R_{1,t+1}$ is the realized gross return on short-term bills, and ν_{t+1} equals one if the household still has a surviving member at $t+1$ and zero otherwise. The interior set excludes the households for which this first-order condition holds as an inequality rather than an equality: recipients of the government transfer, households at the consumption floor, households at the cash-in-advance kink, households with near-zero liquid wealth, and households saving essentially nothing into the liquid account. Constrained households are not solution errors, so they are reported as excluded shares, displayed below each panel of the figure, rather than included in the residual.

Interior residuals track the impatience wedge $\beta\,\mathbb{E}[R_1]-1$, with deviations of roughly one to two percent of one-period consumption, uniformly across ages, solutions, and cali-

Figure A.2: Consumption Euler Residuals on Interior States

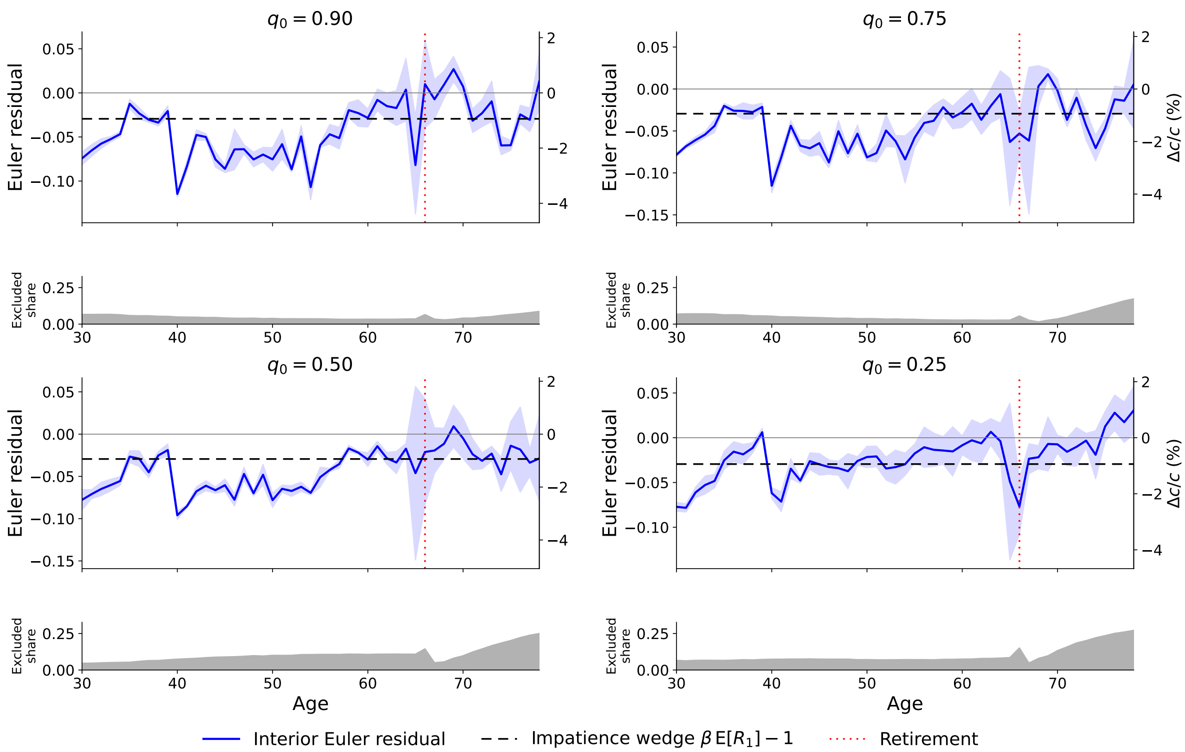

Notes: Each panel plots, for one permanent-income percentile and the $\beta = 0.96$ calibration, the age profile of the Euler residual in equation (60) on the interior set. Shaded bands span plus and minus one standard deviation across four independent solutions of the model. The gray strips below each panel show the share of surviving households excluded from the interior set: recipients of the government transfer, households at the consumption floor or the cash-in-advance kink, and households with near-zero liquid wealth or liquid saving. The dashed line marks the impatience wedge $\beta\,\mathbb{E}[R_1] - 1$; the dotted vertical line marks retirement. The right axis converts residuals into the equivalent one-period consumption misallocation using the curvature $1 + \alpha(\gamma - 1) \approx 3.13$. Ages 30–78.

brations, and with no age-localized artifacts beyond a small and economically interpretable transient at the retirement transition. The estimator in equation (60) already incorporates the covariance between marginal utility and returns as well as the survival transition; computing instead a simple ratio of mean marginal utilities with a constant mean return, or restricting the sample further to households with interior bond positions, leaves the profiles essentially unchanged. The $\beta = 0.93$ calibration displays the same pattern around its deeper impatience wedge of roughly -6%.

C.4 From residuals to welfare

A related question is how much lifetime welfare could remain unexploited when the Euler equation holds to within one to two percent of consumption on interior states. Because deviations from an optimum have no first-order effect on welfare, the answer is about 0.1 percent of lifetime consumption.

The residuals are stated in consumption units using the curvature of utility with respect to non-housing consumption. With period utility proportional to $(c^{\alpha}h^{1-\alpha})^{1-\gamma}/(1-\gamma)$ and housing services held fixed, the elasticity of marginal utility with respect to c is $1+\alpha(\gamma-1)\approx 3.13$ for our calibration, so a wedge of w in marginal-utility units corresponds to a consumption misallocation of $\varepsilon \approx w/3.13$: the measured interior wedges of up to 6% in marginal-utility units are deviations of up to roughly 2% in the level of consumption at that date.

Consider a household whose consumption at some date deviates from its optimal value by a fraction ε, with the resources absorbed by the adjacent saving margin. Because the optimum is a maximum, any budget-feasible perturbation has zero first-order effect on lifetime utility, and the loss is the second-order term

$$\Delta u \;\approx\; \tfrac{1}{2}\, u_{cc}\, (\varepsilon c)^2 \;=\; -\tfrac{1}{2}\, (1+\alpha(\gamma-1))\, u_c\, c\, \varepsilon^2, \tag{61}$$

or, in consumption-equivalent terms, $\frac{1}{2}(1+\alpha(\gamma-1))\,\varepsilon^2 \;\approx\; 0.06\%$ of that period's consumption at $\varepsilon = 2\%$. Persistence does not change the order of magnitude: a persistent intertemporal misallocation must respect the budget constraint, so it tilts the consumption path rather than destroying resources period after period, and the lifetime consumption-equivalent cost is the discount-weighted average of the per-period losses, of order 0.1% of lifetime consumption.

D Additional Figures

Figure A.3: Optimal Equity Shares of Financial Wealth

(a) Equity Share by State of the Cycle

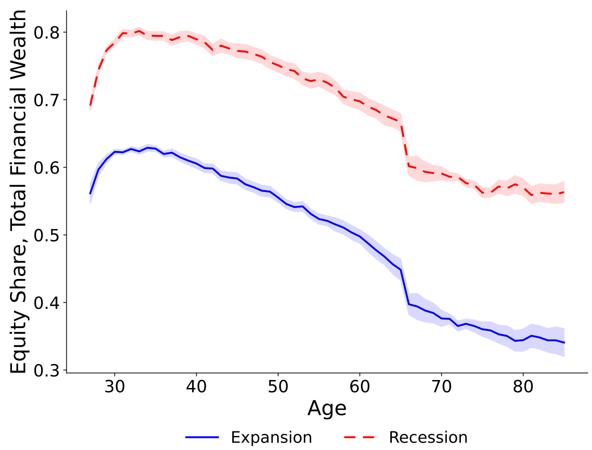

(b) Equity Share by Dividend-Price Ratio

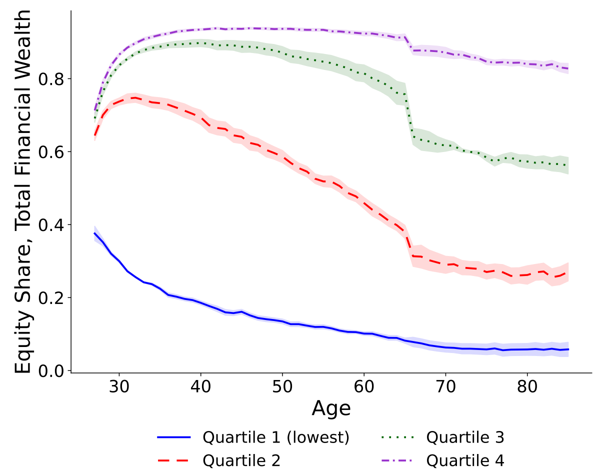

(c) Equity Share by Transitory Income

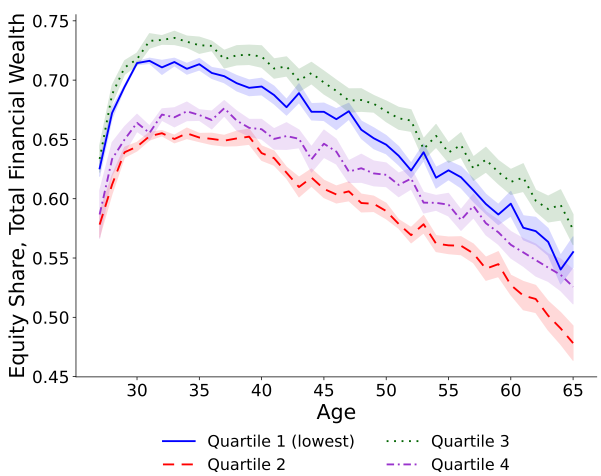

(d) Equity Share by Permanent Income

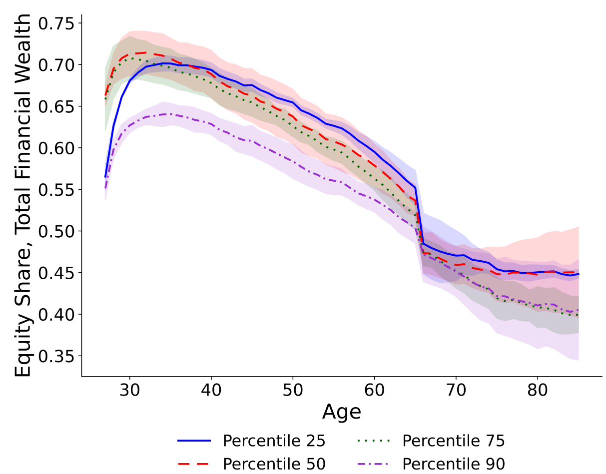

Figure A.4: Optimal Equity Shares of Financial Wealth

(a) Equity Share by Total Financial Wealth

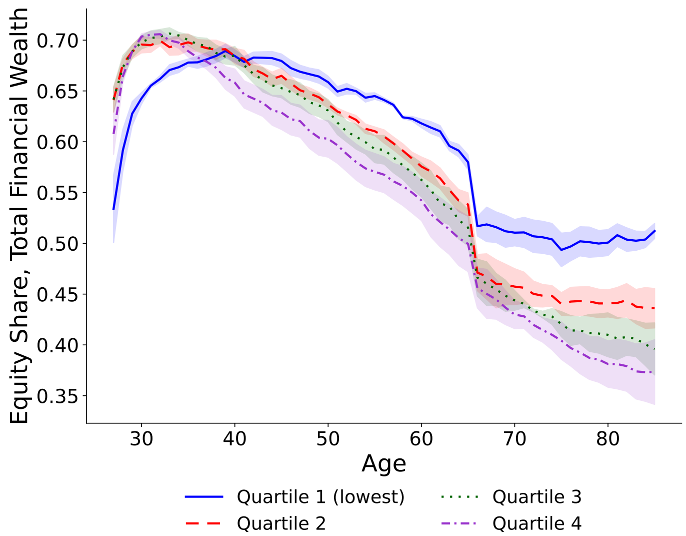

(b) Equity Share by Housing Leverage

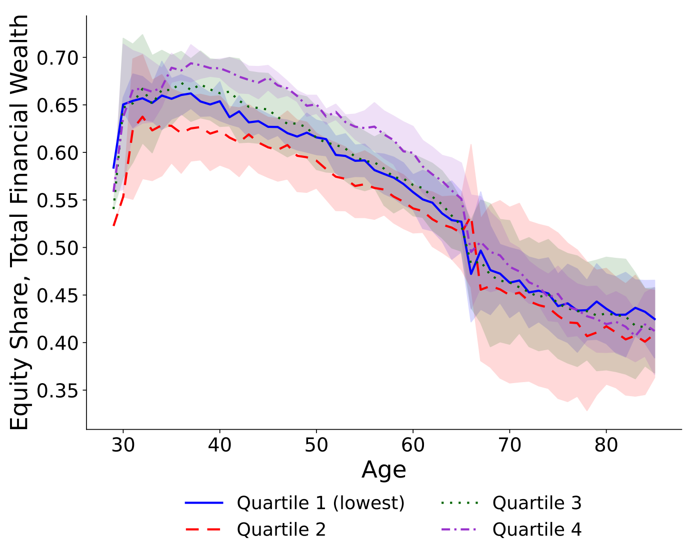

Figure A.5: Optimal Equity Shares of Wealth by Health, Mortality

(a) By Health and Mortality; Retirement Wealth

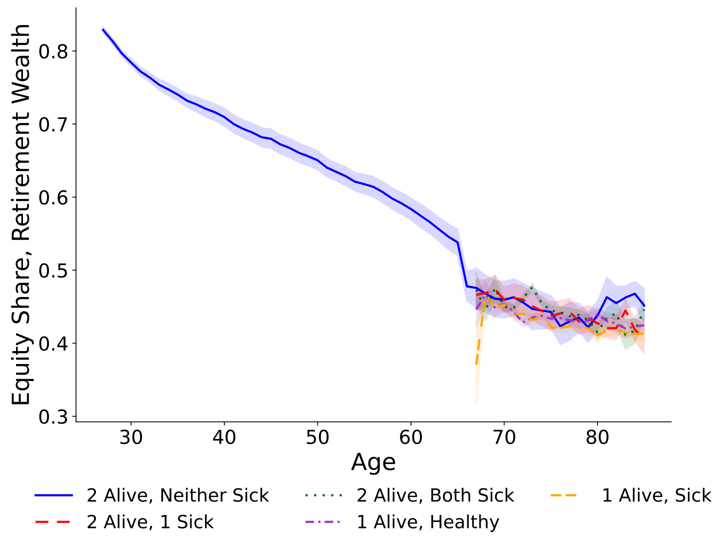

(b) By Health and Mortality; Financial Wealth