Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle*

Victor Duarte
University of Illinois

Julia Fonseca
University of Illinois

Aaron Goodman
MIT

Jonathan A. Parker
MIT and NBER

Preliminary
July 13, 2021

Abstract

We develop a machine-learning solution algorithm to solve for optimal portfolio choice in a detailed and quantitatively-accurate lifecycle model that includes many features of reality modelled only separately in previous work. We use the quantitative model to evaluate the consumption-equivalent welfare losses of using simple rules for portfolio allocation across stocks, bonds, and liquid accounts instead of the optimal portfolio choices. We find that the consumption-equivalent losses from using an age-dependent rule as embedded in current target-date/lifecycle funds (TDFs) are substantial, around three percent of consumption, despite the fact that TDF rules mimic average optimal behavior by age closely (until retirement). The TDF portfolio does not improve on investing a constant 2/3 share in equity. Finally, optimal equity shares have substantial heterogeneity, particularly by wealth level and dividend-price ratio, implying substantial gains to further customization of advice or TDFs, particularly in these dimensions.

JEL codes:
Keywords: Portfolio allocation

*For useful discussions on model components we thank Matthias Doepke, Eric French, Fatih Guvenen, Rory McGee, Maarten Meeuwis, Mariacristina De Nardi, Lawrence Schmidt, Michele Tertilt, and Motohiro Yogo. 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: Department of Economics, 50 Memorial Drive, Cambridge MA 02142, agoodm@mit.edu; Parker: Sloan School of Management, MIT, 100 Main Street, Cambridge, MA 02142, JAParker@mit.edu.

How should people allocate their portfolios? The robust proscription from frictionless equilibrium models is that investors should invest in a weighted combination of the market portfolio and a risk-free asset, where the weight is determined by the investor's risk aversion relative the population. However, this advice is difficult to implement. The market portfolio is hard to determine and construct. Prices and returns may be altered by governments, institutions, and biased beliefs. Markets for insuring many individual rises are incomplete. As a result, most investment advice is based on partial-equilibrium, lifecycle models of saving and portfolio choice that take as given available investment opportunities, as pioneered in Samuelson (1969) and Merton (1969, 1971). This approach is also amenable to the incorporation of individual preference changes and idiosyncratic risks such as from non-traded labor income, home ownership and mortgages, health and mortality risks, pension income, family dynamics, and taxes, each of which can be quantitatively important for optimal investor behavior.

In this paper, we develop a machine learning method to solve for optimal consumption and portfolio allocation rules in a lifecycle model with a relatively complex economic environment that includes ***all*** of the above factors as well as other realistic features of the investment problem facing the typical US household. In addition to age, the model has 22 state variables (some of which only directly matter during retirement or during working life). We model a household that consists of a husband and wife who consume both housing and non-housing consumption and are each endowed with with a gender-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. Households can choose to rent or own with a mortgage or own outright, and pays a cost to sell or refinance. can in housing and borrow using mortgages, and can allocate some of their financial investments into retirement accounts with limited employer matching. The household faces a simple tax and benefit system and gets utility from bequests. During retirement, each individual receives a pension that is a function of lifetime labor income, and faces mortality risk and stochastic medical expenses.

We use the model to study optimal lifecycle portfolio choice and to measure how well simple portfolio rules, like those embedded in existing (low-fee, index) Target Date Funds (TDFs), do at approximating optimal-allocations and delivering similar levels of investor welfare. TDFs have become increasingly important for the typical retirement investor, rising from less than $8 billion dollars in 2000 to more than $2.3 trillion dollars in 2019.

Including balanced funds, which most TDFs turn into a few years after the target date, implies that these funds manage more than $4 trillion, or roughly %20 of the US mutual fund market.[1] We have three main substantive results.

First, our model suggests that on average, the share of financial wealth that the households we study should invest in equity is hump-shaped over the working life, peaking around age 40 and 80% and declining thereafter to just under 60% at retirement. However, for retirement wealth, TDFs do a good job of matching the average optimal behavior. While on average, households optimally hold almost all their retirement wealth in equity when they first start working, this share declines to 80% at age 45. A typical TDF maintains an 80-90% share in equity over this age range. After age 45, the average optimal share in equity declines linearly to about 50%, after which it is roughly constant. Equity shares in TDFs typically continue to decline after retirement to 30-40%, but other than this difference, TDFs are quite good at mimicking average optimal behavior in retirement accounts.

Second, the average optimal behavior masks substantial variation. The 90th percentile of optimal equity share is roughly 100% after age 45 for all financial wealth and is very close to 100% through the entire working life for wealth in retirement accounts. The 10th percentile of the distribution of optimal equity shares is much lower than that average optimal share; the tenth percentile declines across ages from almost 100% at age 25 to only 20% at retirement. As a result, the portfolios delivered by portfolio rules that condition only on age, like TDFs, are inappropriate for many investors, and increasingly so as investors age, so that the differences are substantial in the period of life when people have accumulated the most retirement wealth.

Our third set of results quantify the losses of imposing a common age-specific equity shares as in TDFs. Obviously, TDFs are not imposed on investors, but their rise was significantly driven by their adoption as default investments in retirement plans and investors in TDFs do not re-allocate into and out of TDFs following market returns (Parker, Schoar, Cole, and Simester, 2021a; Parker et al., 2021b; Mitchell and Utkus, 2021). Campbell (2016) (Section 5.1) discusses some of the strengths and weaknesses of TDFs as a solution to the life cycle portfolio problem, and we quantify the overall costs/benefits of these rules in terms consumption-equivalent welfare metrics in our (quite) holistic model.

In our model, a household that is invests its retirement wealth following the portfolio

[1] See Parker, Schoar, and Sun (2021b). The $2.3 trillion in TDFs includes $0.9 trillion in target date collective investment trusts (CITs) which invest like TDFs but have lower fees than the equivalent mutual funds and are primarily used by large employers. Dollar amounts are from Investment Company Institute (2020), figures 2.2 and 8.20, and Morningstar (2020).

of the typical (index, low-fee) TDF loses the equivalent of 2.6% of consumption on average in any given age (and state) when they can re-optimize all other behaviors. This loss rises to 3.4% of consumption if we do not allow the household to re-optimize their other behaviors.[2] While TDFs closely mimic optimal age-contingent allocations, these losses are similar to those of a simple rule that imposes a constant 2/3 equity share across all states and dates. We conclude that there is substantial room for improving investor behavior further by conditioning on more state variables, and from our analysis, dividend price ratios and wealth levels appear to be the state variables that could add the most value. This conclusion echoes the conclusion of Gomes, Michaelides, and Zhang (2021) which shows that there are large welfare gains to TDFs that take advantage of return predictability.

Our findings are made possible by our methodological contribution. Building on Duarte (2019), we develop a method that uses deep learning to solve dynamic stochastic models like this lifecycle model – models with many states and controls, and highly non-linear policy functions – 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 using fully connected feedforward neural networks as cells for two recurrent neural networks (one for working life and one for retirement). 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. In contrast, the traditional numerical dynamic programming (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. There are three advantages of our approach relative to traditional NDP.

First, and foremost, our method is much faster than traditional NDP in complex applications and so can accommodate a large number of state variables and shocks, highly non-linear policy functions, and both discrete and continuous actions. The gains in speed come from our use of new tools from the field of deep reinforcement learning, the field of machine learning that studies sequential decision making. For example, optimizing expected utility using simulated sample paths avoid computationally-slow numerical integration. Second, our method is far easier to use and program (and so less prone to

[2]Both of these losses are smaller – 0.74% and 0.81% respectively – from the perspective of the household at the start of life due to discounting.

error) than traditional numerical dynamic programming methods. For example, there is no need to specify the density and scale of grids over which policy functions can then be defined as matrices. Finally, this method arguably captures how investors, practitioners, or data scientists actually determine optimal behavior. Optimal behavior is determined by simulating and learning from the outcomes of a large number people living their lives according to different portfolio allocation and consumption rules, that is, from watching and learning from the lives of others. We describe our method in detail in Section 3.

While we apply our method to lifecycle consumption and portfolio choice, our method may readily be applied to a wide range of partial-equilibrium dynamic problems faced by managers, employers, small open economies, etc.

Our quantitative model of lifecycle portfolio choice is 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 quantitatively 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 important analyses of optimal portfolio behavior in lifecycle models that primarily focuses on inferring features and parameters of the model from observed behavior. Gakidis (1998), Campbell and Viceira (2002) (Chapter 6), Storesletten, Telmer, and Yaron (2004), Cocco, Gomes, and Maenhout (2005), Gomes and Michaelides (2005), and Davis, Kubler, and Willen (2006) all pioneered 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 accounts, Cocco (2005), Hu (2005), and Yao and Zhang (2005) incorporate housing, and Yogo (2016) focuses on the role of medical expenses during retirement.[3]

[3]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.

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 section 3.1.

1 The Lifecycle Model

Time is discrete and measures age in years. At each age, the household chooses how much to save and how to 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 saving. Each returns processes consists of an idiosyncratic shock and a loading on a common autoregressive process based on the dynamics of the dividend-price ratio of the aggregate stock market. Correlation among labor income and asset returns is captured by having both the dividend price ratio and the labor income process 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.

The most notable omissions 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 $t = T_R = 65$ and who both survive until age T_R with probability one. Households are also differentiated by the deterministic component of each spouses 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. Varying the choice of cohort and initial income quantile allows us to determine how optimal portfolio choice decisions differ for individuals with differently-shaped income profiles and at different places in the lifetime income distribution. Since this is a permanent characteristic, we generally omit q_0^1 and q_0^2 as state variables for notational simplicity.

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 stochastic processes. A stochastic process representing the business cycle generates correlation among labor income and returns on different assets. Second, a stochastic process representing aggregate effective risk aversion 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 aggregate dividend yield (divided 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 Guvenen, Ozkan, and Song document in earnings growth rates. Cyclicality (in particular, counter-cyclical 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

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

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.

1.1.3 Financial Assets

The household's financial wealth is composed of assets held in liquid accounts, a^L, or 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.[4] Households start 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. 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 at (the beginning of) time t, denoted by $R_{j,t}$, 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.

[4] We model traditional accounts only for now, and not the increasingly-common Roth-type accounts. Contributions to Roth accounts do not reduce current taxes, but withdrawals during retirement are tax free.

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 $a^L_{j,t}$ and $a^I_{j,t}$ respectively so that:

$$a^L_t = \sum_{j=1}^{J} a^L_{j,t} \tag{10}$$

$$a^I_t = \sum_{j=1}^{J} a^I_{j,t} \tag{11}$$

Define $s^x_{j,t}$ 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), with $s^x_t = \sum_{j=1}^{J} s^x_{j,t}$. We can then write the state-evolution equation for financial assets as

$$a^L_{t+1} = \sum_{j=1}^{J} R_{j,t+1}(a^L_{j,t} + s^L_{j,t}), \tag{12}$$

$$a^I_{t+1} = \sum_{j=1}^{J} R_{j,t+1}(a^I_{j,t} + \Gamma(s^I_{j,t})), \tag{13}$$

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

$$\Gamma(x) = \begin{cases} x & \text{if } x \leq 0 \text{ or } t \geq T_R \\ (1 + \frac{k}{2})x & \text{if } 0 < x \leq lY_t \\ x + \frac{k}{2} lY_t & \text{if } x > lY_t. \end{cases} \tag{14}$$

We divide the match rate k by 2 on the assumption that only one spouse has access matching employer contributions at any given time. Households cannot short any asset (other than a mortgage, as described subsequently), so holdings in each asset class and account type must be nonnegative:

$$a^L_{j,t} \geq 0, \tag{15}$$

$$a^I_{j,t} \geq 0. \tag{16}$$

Finally, we cap annual contributions to tax-deferred retirement accounts, consistent with

US law:

$$s_t^I \leq s_{max,t}^I. \tag{17}$$

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.

Now define $B_{a,t}$ to be the net contribution to the household's time t budget constraint of all decisions regarding portfolio choice and saving in financial assets. We have:

$$B_{a,t} = -s_t^L - s_t^I - L'(Y_t) \sum_{j=1}^{J} \left(\frac{R_{j,t} - 1}{R_{j,t}} \right) a_{j,t}^L - \xi(s_t^I, t) + L'(Y_t) \max\{0, s_t^I\} \tag{18}$$

where $L'(Y_t)$ is the marginal tax rate, the first derivative of the tax function,[5]

$$L'(Y_t) = 1 - \lambda(1 - \kappa) Y_t^{-\kappa}. \tag{19}$$

and where $\xi \geq 0$ is a function describing the penalty associated with early withdrawal ($s^I < 0$) of retirement assets and the taxes paid pre-retirement:.

$$\xi(s^I, t) = \begin{cases} -(L'(Y_t) + 0.1)s^I & \text{if } s^I < 0 \text{ and } t < 60 \\ -L'(Y_t)s^I & \text{if } s^I < 0 \text{ and } 60 \leq t \\ 0 & \text{if } s^I \geq 0. \end{cases} \tag{20}$$

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 and any taxes or penalties incurred by withdrawing from retirement accounts. The cost of contributing to tax-deferred retirement accounts is partially offset by the reduction in current tax liability, which enters positively into the budget constraint. Note that ξ includes the 10% penalty

[5]Note that for simplicity, and to match the reality 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. Since withdrawals from the illiquid retirement account are likely to be a substantial share of income during retirement, we treat withdrawals during retirement equivalently to other income which accounts for second-order effects during retirement (see equation (41)).

associated with withdrawals from the illiquid account prior to age 60, but does not include the regular taxation of withdrawals post-retirement since (as described in Section 1.2.1) we treat these withdrawals as equivalent to labor income for tax purposes.

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 wealth held in short-term debt:

$$c_t \leq M a_{1,t}^L \text{ if } a_t^I > 0. \tag{21}$$

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. 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, and they also receive an additional utility benefit from any level h_t of housing.[6] 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

[6] This utility benefit is not in BGLV and is discussed subsequently as part of the objective function.

geometric random walk with drift:

$$p_{t+1} = \nu_t p_t, \tag{22}$$

$$\log(\nu_t) \sim N(\mu_\nu, \sigma_\nu). \tag{23}$$

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:

$$B_{h,t}(0) = -(1 - o_t)\left[\phi p_t h_t\right] - o_t\left[p_t\tilde{h}_t - d_{t+1} + \delta p_t\tilde{h}_t\right] \tag{24}$$

$$d_{t+1}\begin{cases} = 0 & \text{if } o_t = 0 \\ \leq (1-\iota)p_t h_t & \text{if } o_t = 1. \end{cases} \tag{25}$$

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).

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

$$\begin{aligned} B_{h,t}(1) = &-o_t[\delta p_t \tilde{h}_t + (1 - L'(Y_t))(R_{1,t} - 1 + \Delta r_m)d_t \\ &+ \mathbf{1}_{\{\tilde{h}_t = \tilde{h}_{t-1},\, d_{t+1} > d_t\}} f^d p_t \tilde{h}_t + \mathbf{1}_{\{\tilde{h}_t \neq \tilde{h}_{t-1}\}} p_t(\tilde{h}_t - (1 - f^h)\tilde{h}_{t-1})] \\ &+ (1 - o_t)[(1 - f^h)p_t \tilde{h}_{t-1} - \phi p_t h_t] - (d_t - d_{t+1}) \end{aligned} \tag{26}$$

$$d_{t+1} = \begin{cases} = 0 & \text{if } o_t = 0 \\ \leq (1 - \iota)p_t h_t & \text{if } o_t = 1 \text{ and } (\tilde{h}_t \neq \tilde{h}_{t-1} \text{ or } d_{t+1} > \chi d_t) \\ \leq \chi d_t & \text{otherwise.} \end{cases} \tag{27}$$

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 (27) 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

We can now write down a complete statement of the household's working life problem.

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. As described in Section 2, this profile is such that household size declines to 2 at retirement. Total consumption at each age is a Cobb-Douglas aggregate of non-housing consumption c_t and housing consumption h_t with share parameter α, where housing consumption is the service flow from housing, ϕh_t.[7] We assume constant relative risk aversion utility of total consumption in each year, with risk aversion coefficient γ.

[7] ϕ plays no role whatsoever so we set it to the rental rate.

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" Q_t as the liquid wealth available to the household at the beginning of time t, net of taxes and mortgage interest payments:

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

Policy functions are functions of the complete set of state variables at the beginning of each age t: $\Xi_t = \left\{ e_t, v_t, \{x_t^i\}_{i=1}^2, Y_t, \{\sum_{t'=T_0}^{t} y_{t'}^i\}_{i=1}^2, Q_t, a_{1,t}^L, a_t^I, p_t, o_{t-1}, h_{t-1}, d_t \right\}$. 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 (21) 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, $\left\{ c_t, o_t, h_t, \{a_{j,t+1}^L\}_{j=1}^J, \{a_{j,t+1}^I\}_{j=1}^J, d_{t+1} \right\}$ as a function of the state variables at each age, Ξ_t, to maximize the expected discounted sum of time-separable flow utility. Defining the consumption aggregate as

$$C_t = c_t^{\alpha} (\phi h_t)^{1-\alpha}, \quad (29)$$

the household maximizes:

$$\mathbb{E} \left[\sum_{t=T_0}^{T_R - 1} \beta^{t-T_0} \frac{w_t^{\gamma-1} C_t^{1-\gamma}}{1-\gamma} + \beta^{T_R - T_0} V_R(\Xi_{T_R}) \;\middle|\; \Xi_{T_0} \right] \quad (30)$$

subject to the comprehensive flow budget constraint

$$Y_t - L(Y_t) - c_t + B_{a,t} + B_{h,t}(o_{t-1}) = 0, \quad (31)$$

the constraints (15)-(17) and (21), and where the income process is defined in (3)-(6), the tax function is defined in (7), $B_{a,t}$ defined in (18)-(20), and $B_{h,t}(o_{t-1})$ is defined in (24)-(27).

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

$$Y_t^i = \nu_t^i Y(t; q_R^i), \tag{32}$$

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. We introduce new state variables, g_t^i, representing 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

$$\Pr(g_t^i = 1) = P_g(g_{t-1}^i, i, t; q_R^i). \tag{33}$$

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 is capped at $\bar{m}$. The process for medical expenses is given by:

$$m_t = \min[m_t^1 + m_t^2, \bar{m}] \tag{34}$$

$$\log(m_t^i) = \nu_t^i m(g_t^i, i, t; q_R^i) + \nu_t^i \sigma(g_t^i, i, t; q_R^i)\psi_t^i, \tag{35}$$

$$\psi_t^i = \zeta_t^i + \epsilon_{\psi,t}^i, \tag{36}$$

$$\zeta_t^i = \rho_\zeta \zeta_{t-1}^i + \epsilon_{\zeta,t}^i, \tag{37}$$

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

$$\epsilon_{\zeta,t}^i \sim N(0, \sigma_\zeta). \tag{39}$$

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.

To reflect means-tested government programs such as Medicaid that support the elderly, the household is guaranteed a minimum level of consumption when necessary expenditures like realized medical expenses and mortgage payments are sufficiently high. This government transfer rule is described by

$$n_t = \begin{cases} 0 \quad \text{if} \quad a_{t+1}^L + a_{t+1}^I > 0 \quad \text{or} \quad o_t \neq o_{t-1} \quad \text{or} \quad h_{t+1} \neq h_t \quad \text{or} \quad d_{t+1} \neq \chi d_t \\ \max\left\{0, \frac{\underline{c}}{w_t} + m_t + \tau_t - Y_t + L(Y_t) - B_{a,t} - B_{h,t}(o_t)\right\} \quad \text{otherwise,} \end{cases} \tag{40}$$

where n_t is the government transfer and $\underline{c}$ is the consumption floor (adjusted by the effective consumption shifter). The household is not allowed to receive government transfer payments if it carries any financial wealth to the next period ($a_{t+1}^L + a_{t+1}^I > 0$), adjusts its housing consumption ($h_{t+1} \neq h_t$), changes its home ownership status ($o_t \neq o_{t-1}$), or refinances or makes a larger-than-necessary payment on its mortgage ($d_{t+1} \neq \chi d_t$).

Otherwise, when the household's budget constraint is tight enough to push consumption below $\underline{c}$, the government transfer makes up the difference.

Finally, we continue to assume that the household pays taxes during the year as during working life, but also makes an additional tax payment or receives a 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:

$$\tilde{L}(Y_t, s_t^I) = \begin{cases} L\left(\max\{0, -s_t^I\}\right) & \text{if } 0.5Y_t + \max\{0, -s_t^I\} < I_{0.5} \\ L\left(0.5Y_t + \max\{0, -s_t^I\}\right) & \text{if } I_{0.5} \leq 0.5Y_t + \max\{0, -s_t^I\} < I_{0.85} \\ L\left(0.85Y_t + \max\{0, -s_t^I\}\right) & \text{if } 0.5Y_t + \max\{0, -s_t^I\} \geq I_{0.85} \end{cases} \tag{41}$$

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).[8] The difference between this and what is paid during the year is the tax bill (or refund) that the household must pay in $t+1$:

$$\tau_{t+1} = \tilde{L}(Y_t, s_t^I) - L(Y_t) - L'(Y_t)s^I. \tag{42}$$

[8]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).

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 14 and 20). However, it must pay income taxes on withdrawals from the retirement account (as given in equation 20 and is subject to the IRS's rules regarding required minimum distributions from retirement accounts after age 70 (equation 17). Thus, $B_{a,t}$, as defined in Equation 18, 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

$$u_b(b_t) = \tilde{b}\frac{(b_t + \underline{b})^{1-\gamma_b}}{1-\gamma_b}, \tag{43}$$

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

$$b_t = a_t^L + a_t^I - L\left(\frac{a_t^I}{5}\right) - \tau_t + \max\{(1-\delta-f^h)p_t\tilde{h}_{t-1} - d_t, 0\}, \tag{44}$$

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,[9] τ_t gives the carry-forward tax bill determined in the previous period, and the final term gives the current value of the

[9]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 (44) negative, we set it to zero.

As in the pre-retirement period, consumption needs depend on 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, $w_t=\sqrt{v_t^1+v_t^2}=0$. We also allow consumption utility to depend on the health status of household members. In particular, to reflect the difficulty of maintaining a home and the increased likelihood of moving to assisted living environments when in poor health, the utility benefit of homeownership relative to renting decreases if the household's members are in the "sick" state. We implement this utility reduction by multiplying homeowners' housing consumption h_t by a scaling factor $\tilde{u}<1$, raised to the number of household members who are sick so that we redefine the consumption aggregate as

$$C_t=c_t^{\alpha}(\tilde{u}^{o_t(g_t^1+g_t^2)}\phi h_t)^{1-\alpha}. \tag{45}$$

This equation collapses to (29) during working because we assume everyone is healthy so that $g_t^1+g_t^2=0$.

The set of state variables during retirement is different from that during the 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 Q_t^R, this gives

$$Q_t^R=Q_t-m_t-\tau_t, \tag{46}$$

where Q_t is defined in (28). Finally, the health status (g_t^i) and survival (v_t^i) states of the household's members are now state variables.[10]

[10]Thus, during retirement the state at age t is given by Ξ_t =

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 at every age to maximize:

$$\mathbb{E}\left[\sum_{t=T_R}^{T_{max}} \beta^{t-T_0}\left(\mathbf{1}_{\{w_t>0\}}\frac{w_t^{\gamma}C_t^{1-\gamma}}{1-\gamma}+\mathbf{1}_{\{w_t=0,\,w_{t-1}>0\}}\tilde{b}\frac{(b_t+\underline{b})^{1-\gamma}}{1-\gamma}\right)\;\middle|\;\Xi_{T_R}\right] \tag{47}$$

subject to a modified version of the budget constraint during the working life (31) 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:

$$Y_t - L(Y_t) - c_t - m_t + n_t - \tau_t + B_{a,t} + B_{h,t}(o_{t-1}) = 0. \tag{48}$$

The household is subject to the contribution and withdrawal constraints (15)-(17) and (21). Income is given by (32), the tax function is defined in (41), the process for medical expenditures is defined in (34)-(39), the benefit n_t is defined by (40), b is given by (44), $B_{a,t}$ and $B_{h,t}(o_{t-1})$ are defined as during working life by (18)-(20) and (24)-(27), and finally consumption is given by (45).

The retirement value function referenced in the working-life objective (30), $V_R(\Xi_{T_R})$, is given by the maximized value of (47).

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.[11] We discuss our parameterization decisions here and provide a complete list of parameter values in Table I. 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

$\left\{e_t, v_t, Y_t, Q_t^R, a_{1,t}^L, a_t^I, \{g_t^i\}_{i=1}^2, \{q_R^i\}_{i=1}^2, \{\zeta_t^i\}_{i=1}^2, \{v_t^i\}_{i=1}^2, p_t, o_{t-1}, h_{t-1}, d_t\right\}$.

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

no financial assets, no housing wealth, no mortgage debt, and the persistent component of labor income variance $x^i_{T_0}$ set to 0.

2.1 Common Risks

We use official NBER recession-dating betweeen 1915 and 2015 to estimate the recession transition matrix P_e. To estimate the dividend yield process in (1) we use data on U.S. dividend yields during the 1915-2015 period from Jorda, Kroll, Kuvshinov, Schularick, and Taylor (2019).

2.2 Labor Income and Taxes

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

To specify the deterministic age profile, we use 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). After specifying a gender i, an initial income quantile q^i_0, and a cohort k, we simply read off the relevant 31-observation time series from GKSW and use it as our age profile, so that $\bar{y}^i_t = \bar{y}(t; i, q^i_0, k)$.[12]

Because the oldest cohort for which GKSW report income summary statistics turned 25 in 1957, and because GKSW only report summary statistics through age 55, we must perform some interpolation to populate our deterministic income profile for the 25-32 and 56-65 age ranges. 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.[13] 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.

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.

[12] 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.

[13] We obtain real wage growth rates from Peake and Vandenbroucke (2020).

Table I: Parameter Values

Parameter	Value	Description	Source
T_0	25	beginning of working life	calibrated
T_R	65	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 $j \in \{$long bond, short bond, equity$\}$	estimated
ρ_x	0.979	labor income persistence	GOS
$\mu_{\eta_1,0}$	0.119	first labor income innovation mean, expansion	GOS
$\mu_{\eta_1,1}$	-0.102	first labor income innovation mean, recession	GOS
$\mu_{\eta_2,0}$	-0.026	second labor income innovation mean, expansion	GOS
$\mu_{\eta_2,1}$	0.094	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	BGLV
σ_ν	0.039	house price innovation std. dev.	BGLV
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
$\bar{m}$	1,000,000	family medical expenditure cap	calibrated
$\underline{c}$	\$3,957	retirement consumption floor	DFJ
γ	3.84	risk aversion	DFJ
$\tilde{b}$	23.6	bequest utility intensity	DFJ (2016)
$\underline{b}$	369,000	bequest utility intercept	DFJ
$\tilde{u}$	0.90	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	0.888	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.

2.3 Financial Assets

We estimate the asset return processes in (8) with U.S. data over the 1915-2015 period. We take data on short-term government bills and equities from Jorda et al. (2019) and use the Dow Jones Total Corporate Bond index to estimate the long-term corporate bond process.

We set $k = 0.5$ and $l = 0.06$ in the employer-matching rule (14) to reflect the common employer policy of matching 50% of contributions up to 6% of income. As reflected in (14), 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 (17) 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.[14] We calibrate the cash-in-advance parameter M to 12 so that the household must carry the equivalent of 1 months' worth of consumption in safe, liquid wealth.

2.4 Housing

We take all 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.

2.5 Retirement

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 (41) and the required minimum distributions ($s^I_{max,t}$ for $t \geq T_R$ in equation (17)) 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)). We thus obtain a simulated distribution of working-life incomes. When an individual in the model reaches retirement, we sum their total labor income and let q^i_R be their rank within this simulated distribution of cumulative income. We then base $Y(t; q^i_R)$, and the other DFJ

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

processes described below, on q_R^i: specifically, we use the DFJ processes estimated at the permanent income quantile closest to q_R^i.[15]

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).[16] 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.

We set the homeownership scaling factor $\tilde{u}$ to 0.9, so that effective housing consumption is reduced by 10% if the household is a homeowner and one member is sick, and by 19% if the household is a homeowner and both members are sick.

3 Solution Method

To solve the investor's optimization problem, we build on the framework of Duarte (2019) which develops a method to use policy gradient algorithms to solve high-dimensional problems in economics and finance. In general, policy gradient methods work by parameterizing the agent's decisions with flexible functional forms and optimizing the expected value of a stochastic reward (see Sutton and Barto, 2018) . In our setting, the stochastic reward is the household's lifetime utility.

We parameterize the investors' policy functions as fully connected feedforward neural networks and update the networks' parameters with stochastic gradient descent. This machine learning approach greatly improves the efficiency and speed of the algorithm relative to a more traditional numeric dynamic programming approach of defining policy

[15] DFJ estimate their processes at each fifth percentile, so we choose the DFJ quantile in $\{0, .05, .1, ..., .95, 1\}$ that is closest to q_R^i.

[16] 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.

functions over grids of state variables and performing numerical integration to compute expectations. These efficiency gains make it feasible to solve models as rich and complex as the one described in Section 1.

There are however two key challenges that arise in applying machine learning tools to solve our complex dynamic stochastic problem (and also many other problems in economics and finance).

The first challenge is the presence of both discrete and continuous choices in a model with large state and action spaces. With only continuous choices, it is possible to compute the gradient of the objective function with respect to the parameters of the neural network representing policy functions using automatic differentiation, an algorithm used in most machine learning applications. With these gradients in hand, one could then adjust network parameters in the direction of the gradient, which is the direction that (locally) increases the objection function the fastest.[17] However, with discrete choices, the objective function is not differentiable with respect to network parameters and we cannot make use of automatic differentiation.

We overcome this first challenge by parameterizing the probability that a household makes a certain discrete choice, instead of parameterizing the discrete choice itself. This means that, while training our networks, discrete choices are stochastic, but the policy functions in our solution are deterministic and map states into the discrete choices that maximize the objective function. We describe this approach in detail in Section 3.1.

The second challenge is the combination of the long length of life and the variance and persistence of shocks (as well as the significant number of state variables). These features pose a problem because stochastic gradient descent uses simulated data (lifetimes of shocks) to compute approximations of the gradient of the objective function with respect to network parameters. And from these approximations, determine how to optimally adjust these network parameters. In an environment with a large finite horizon and high-variance shocks, the variance of these gradients is very large, making it hard to approximate them through simulations.

We address this second challenge by including time as an input of our network, instead of having a separate network for each time t, which is equivalent to using a recurrent neural network. With this structure, changing a parameter of the network affects actions across all time periods, and thus increases gradients without increasing their variance, making them easier to estimate. We discuss this further in Section 3.2.

[17] For a more detailed description of this approach, see Duarte (2019).

Further, because the set of control variables and the set of state variables both change at retirement, we use separate networks for the pre-retirement and post-retirement periods of life.

In the remainder of this section, we expand on this solution algorithm, describing how we simulate the model, choose network parameters to maximize lifetime utility over simulated sample paths, and deal with discrete choices. Finally, we present the architecture of our neural networks, which are chosen due to their computational efficiency, and describe how we implement this solution algorithm.

3.1 Algorithm

Let π be a policy function representing the choices of a household. The expected lifetime utility associated with policy π is

$$\mathbb{E}_{\varepsilon\sim\mathcal{D}}\left[\sum_{t=T_0}^{T_D}\beta^t u_t(\pi_t(\Xi_t),\Xi_t)+\beta^{TD}\mathcal{B}(T_D,\Xi_{T_D})\right], \tag{49}$$

where u_t is the flow utility at time t, $\mathcal{B}$ is the bequest function, T_D is the stochastic age at which the longest surviving member of the household dies, Ξ_t is the set of state variables at time t, and ε is the collection of all shocks. Shocks and initial states are sampled from the distribution $\mathcal{D}$.

Our goal is to find policy functions that maximize lifetime utility. To do so, we parametrize a given policy π using a neural network

$$\pi_t(\Xi_t)\equiv\pi(\Xi_t;\Theta^\pi). \tag{50}$$

For simplicity, we refer to the collection of parameters for all networks as Θ and to the sum of discounted utilities given a collection of parameters and shocks as $R(\Theta,\varepsilon)$. Given this parametrization, we define the loss function as

$$\mathcal{L}(\Theta)=-\mathbb{E}_{\varepsilon\sim\mathcal{D}}\left[R(\Theta,\varepsilon)\right]. \tag{51}$$

The set of parameters Θ is chosen to minimize the loss function. To find this minimum, we use stochastic gradient descent and proceed iteratively. For a set of initial parameters Θ, assuming that the loss function is differentiable with respect to Θ, one step of standard (non-

stochastic) gradient descent adjusts the parameters according to the following equation

$$\Delta\Theta = -\alpha \nabla_\Theta \mathcal{L}(\Theta), \tag{52}$$

where α denotes the learning rate and ∇_Θ denotes the gradient with respect to Θ. The idea behind gradient descent is to move Θ in the direction that (locally) reduces the loss function the fastest. However, since some of choices represented by π are discrete, the loss function given by Equation (51) is not differentiable with respect to Θ.

We circumvent this issue by relying on Parameter-Exploring Policy Gradient methods (Sehnke, Osendorfer, Rückstieß, Graves, Peters, and Schmidhuber, 2010). Instead of searching for a single vector of parameters Θ, these methods consist of searching for a distribution of parameters. Specifically, we assume

$$\Theta \sim N(\mu, \sigma^2 I), \tag{53}$$

where μ is a vector with the same dimensions as Θ, σ is a scalar, and I is the identity matrix. We chose μ so as to minimize the loss function, which we can now define as

$$\mathcal{L}(\mu) = -\mathbb{E}_{\Theta\sim N(\mu,\sigma),\varepsilon\sim\mathcal{D}}\left[R(\Theta,\varepsilon)\right]. \tag{54}$$

As in Salimans, Ho, Chen, Sidor, and Sutskever (2017), we treat σ as a fixed hyperparameter, initially setting σ to 0.001 and annealing it to zero during training. Importantly, as σ approaches zero, our algorithm settles on a deterministic vector of parameters Θ and policy functions map a set of state variables into deterministic actions. Sehnke et al. (2010) show that the loss function given by Equation (54) is differentiable with respect to μ even when the original loss function given by Equation (51) is not, and is given by

$$\nabla_\mu \mathcal{L}(\mu) = -\mathbb{E}_{\Theta\sim N(\mu,\sigma),\varepsilon\sim\mathcal{D}}\left[\nabla_\mu \log\varphi(\Theta) R(\Theta,\varepsilon)\right], \tag{55}$$

where φ is the probability density function of Θ.

It is infeasible to compute the expectation in the equation above for our model. Instead, the insight of stochastic gradient descent is to approximate this expectation with a small i.i.d. sample of simulated paths. Accordingly, we adjust μ at each iteration by

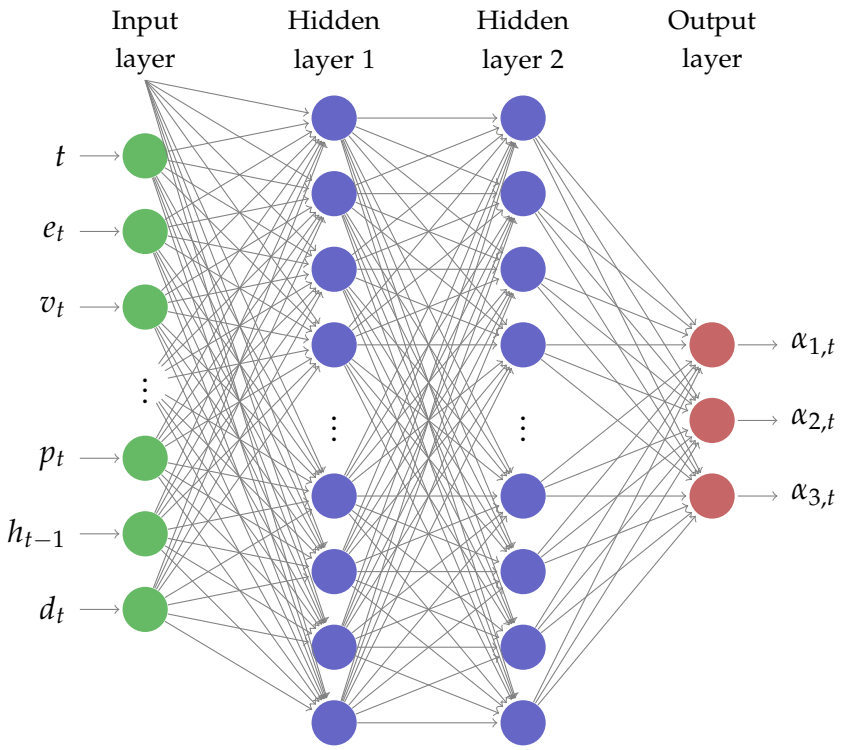
$$\Delta\mu \approx \alpha \frac{1}{N}\sum_{i=1}^{N} \nabla_\mu \log\varphi(\Theta_i) R(\Theta_i,\varepsilon_i) \tag{56}$$

and proceed by adjusting parameters according to this equation at each iteration until a suitable convergence criteria is met.

3.2 Network Architecture

We parameterize policy functions using two neural networks, one for the working age period and one for the retirement period. Every possible action the agent takes is modeled as a sequence of two hidden layers and one output layer. We illustrate this structure for the portfolio choice of an agent during the working age period in Figure I. As mentioned in Section 3.1, this network takes as inputs (represented as green nodes in Figure I) time (t) and the set of state variables. It produces as outputs (represented as red nodes) the portfolio shares of the three assets.

Figure I: Network Architecture



Note: This figure illustrates the architecture of a neural network representing the portfolio choice at time t of an agent during working life.

Note that this network has two hidden layers (represented by blue nodes in Figure I). Each node in the first hidden layer is a non-linear transformation of a linear combination

of the inputs. We normalize inputs using a slow-moving average of inputs over the first 10,000 iterations.[18] Each node in the second hidden layer is a non-linear transformation of a linear combination of the nodes of the first hidden layer. Finally, each output (or node) of the output layer is a non-linear transformation of a linear combination of the nodes of the second hidden layer. These non-linear transformations are commonly referred to as activation functions.

The three portfolio shares illustrated in Figure I sum to one due to our choice of softmax as an activation function for the output layer.[19] The softmax function is a function $\sigma : \mathbb{R}^n \to \mathbb{R}^n$ such that

$$\sigma(x)_i = \frac{e^{x_i}}{\sum_{j=1}^{n} e^{x_j}}, \text{ for } i \in \{1, \ldots, n\} \tag{57}$$

This function normalizes a vector x into a vector of non-negative fractions that sum to one. The usefulness of this transformation is immediate in our portfolio choice example, in which the agent must choose how much of his financial wealth to allocate across the three assets and the no-borrowing constraints of Equations 15 and 16 imply that portfolio shares must be non-negative.

We also use a softmax activation function when representing a household's decision to split his disposable income between consumption, financial investment, housing, and debt repayment. This implies that each of these variables will be a fraction of disposable income and, consequently, the budget constraint will always be satisfied. The activation function for the output layer associated with discrete choices is also a softmax, and the agent chooses the discrete case associated with the highest output.

Finally, as illustrated in Figure I, time is an input of our network and we do not have a separate network for each time t. The intuition for this structure is similar to what is known as parameter sharing (Caruana, 1993), recognizing that an output at time t will be very similar to an output at time $t+1$. This reduces the number of network parameters and increases efficiency (Tsitsiklis and Van Roy, 2001). In particular, this increases the magnitude of gradients without increasing their variance, since changes to a parameter will affect actions across all time periods. This allows us to better approximate gradients with a small i.i.d. sample of simulated paths.

[18] As in any application using neural networks, it is important for our algorithm that inputs have similar magnitudes. In our setting, some inputs are endogenous, so we do not know ex ante their respective means and variances. We thus use a slow-moving average of simulated states in this normalization.

[19] For the hidden layers, we use a hyperbolic tangent (tanh) activation function.

3.3 Solution

We run code to implement this solution in Google's TensorFlow Research Cloud, a cloud service intended for researchers.[20] At each iteration, we sample 2,048 vectors of parameters according to the distribution given by Equation 53.[21] As we describe above, a vector of parameters determines the policy functions of an investor and, for each vector of parameters, we simulate the lives of 256 investors who make decisions according to those policy functions. This amounts to simulating the lives of 524,288 investors (256×2,048) per iteration.

As we describe in Section 2.2, the deterministic age profile in our labor income process is conditional on an initial permanent income percentile $q \in \{25, 50, 75, 90\}$. We solve for optimal behavior separately for each permanent income percentile assuming that both members earn labor income according to the percentile q process.[22] The code that solves for optimal behavior given a permanent income percentile takes approximately 24 hours to run on Google's TensorFlow Research Cloud.

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 share of financial assets optimally invested in equity starts quite low but rises rapidly early in life, peaking around age 45 and declining thereafter. For retirement wealth, the average optimal share invested in equity starts high, and overall looks a lot like the glide path of a typical TDF. Second, however optimal portfolios differ significantly depending on other state variables, most importantly wealth level and the 'aggregate' variables that affect expected stock and bond returns.

As just 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

[20] The TensorFlow Research Cloud can be accessed at https://www.tensorflow.org/tfrc.

[21] We follow Google's recommendations for optimal performance and define the size of all arrays as multiples of 128. More information on the architecture of Google's cloud system is available at https://cloud.google.com/tpu/docs/system-architecture.

[22] This amounts to an assumption that assortative marriage matching causes all individuals to have the same permanent earnings capacity as their spouse.

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.

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).

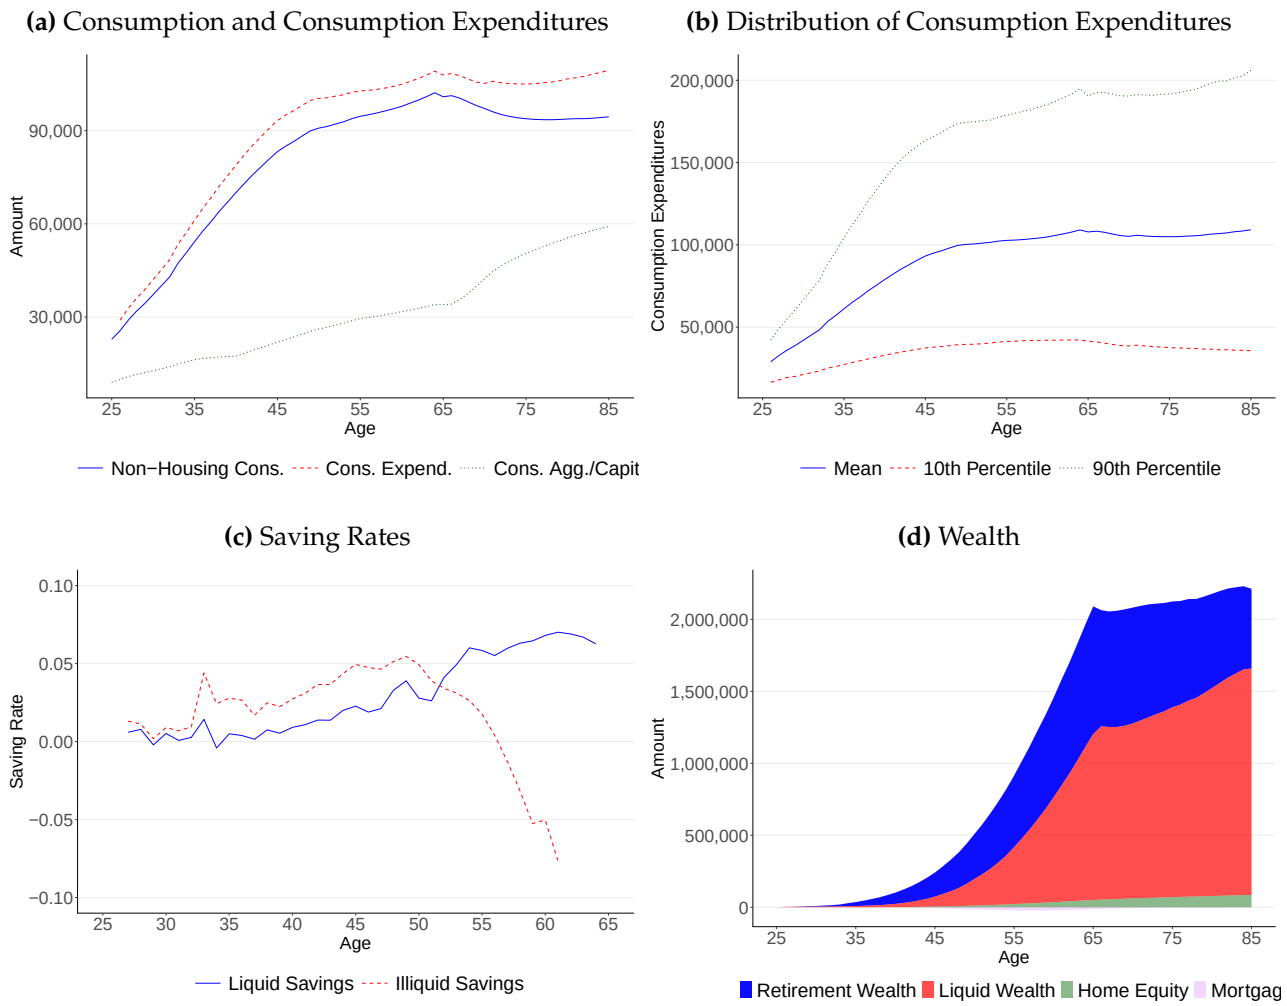
4.1 Optimal Behavior

Figure II.a plots the average level of non-housing consumption (c), total consumption expenditures (c plus mortgage and homebuying transaction costs, medical expenses, and rent or rental equivalent (ϕh)), and the consumption aggregate per effective family member, $c_t^{\alpha}(\tilde{u}^{o_t(g_t^1+g_t^2)}\phi h_t)/w_t$.[23] Because returns are high relative to households impatience, because of precautionary saving, and because of matching of retirement saving, the first two measures rise rapidly until the mid-40's and all three measures rise until retirement. Non-housing consumption, c, falls after retirement while average total consumption expenditures is roughly flat and the difference between the two series is occasional large medical expenses (which become more likely as people become older. Average per capita consumption rises steadily but more slowly throughout working life and then accelerates in old age largely in response to death of one spouse. Total consumption expenditures average about $\$30,000$ at the start of life, rise to just over $\$90,000$ during mid-life, and are around $\$110,000$ including medical expenses at age 85. Again these figure pertain to people retiring in 1989 and are in 2013 dollars.

Figure II.b shows the 10th, 50th and 90th percentiles in the distribution of consumption expenditures by each age. There is substantial increase in dispersion until around age 45, at which time consumption inequality remains roughly constant until the retirement period when it is steadily increasing due to medical expenses and death of one spouse. In mid life

[23] We exclude maintenance costs from consumption expenditures since they are part of the rental equivalence and when someone hits the consumption floor their consumption expenditures are $\underline{c}$

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

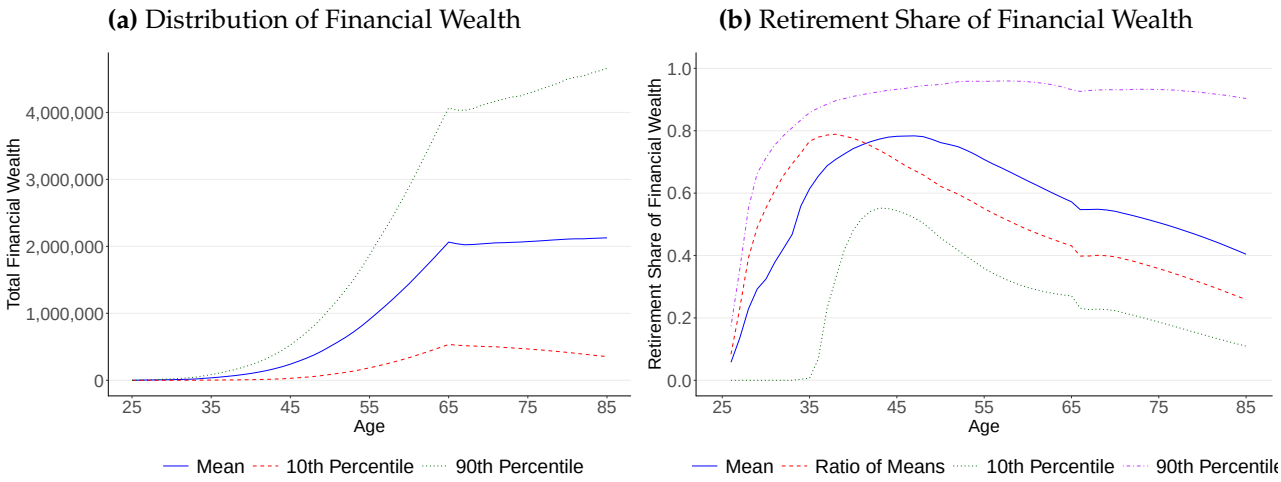


Notes: Consumption expenditures are defined as: non-housing consumption, health expenditures, house and mortgage transaction costs, and rent or rental equivalence (ϕh). The consumption aggregate per capita is the argument of the CRRA period utility function. Saving rates are averages of x^L/Y and x^I/Y dropping family labor incomes below \$10,000 and are plotted from 25 to 64.

the 10th percentile of expenditures is only about \$45,000 in a year while the 90th percentile is about \$180,000.

Turning to saving rates, Figure II.c shows that in their first few years of working life, young households build a little liquid wealth but do not save much on average in liquid wealth until age 40. Until late in working life, households primarily save in retirement accounts, with retirement saving rates rising from 2.5% at age 35 to peak at 5% at age 50. The average saving rate in retirement accounts turns negative as people approach retirement because those agents having saved the most in retirement accounts, experiencing

Figure III: Financial Wealth by Age



Notes: Financial wealth is liquid wealth plus retirement wealth.

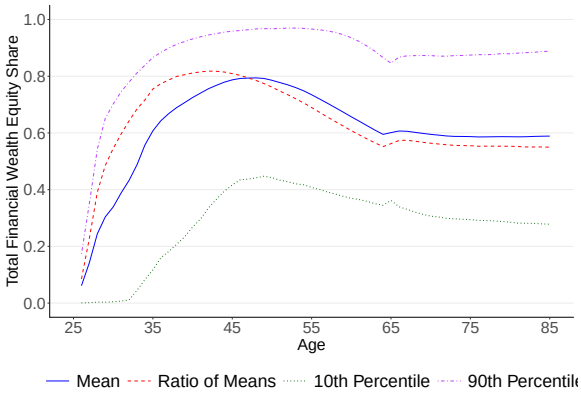
the highest returns, and getting low labor income start withdrawing substantial amounts which lowers the average despite most households still contributing to their retirement account prior to retirement.

Figure II.d shows the average wealth accumulated at each age. On average, households build liquid and illiquid wealth steadily and rapidly during their working lives, and at a slower ate during retirement. The share of households that are homeowners remains very low until around age 35, at which age the home ownership rate rises over the next 20 years so that by (and during) retirement just over 60% of households own homes. Mortgage debt peaks about age 50, but overall remains small relative to average wealth (partly because of the low homeownership rate, partly because wealthy households hold little mortgage debt). At retirement, average wealth is just above $2 million, with less than half of that in retirement accounts. Because II.d averages each type of wealth, Figure III does not capture the experience of the typical family (it is a wealth-weighted average share).

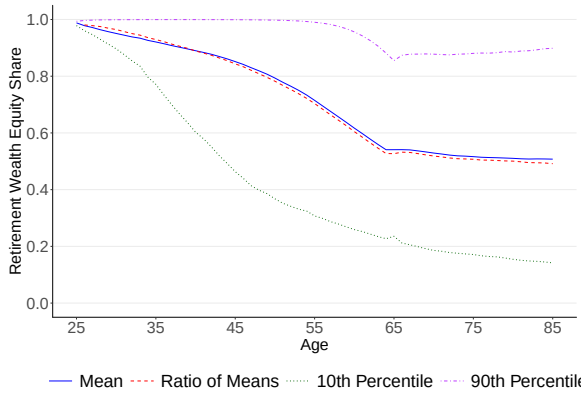
More importantly, underlying these averages are substantial differences in financial wealth across households at every age. Figure III.a shows that, during retirement, the 90th percentile of the wealth distribution is double the average, above $4 million, and the 10th percentile is half the average, only around $500,000. Further, wealth inequality rises during retirement as the wealthy earn high returns and save for bequests while the medical costs matter more for relatively low-wealth households who deccumulate wealth as they age. While the average household holds most of its financial wealth in retirement accounts from

Figure IV: Distribution of Equity Shares of Wealth

(a) Equity Share of Financial Wealth



(b) Equity Share of Retirement Wealth



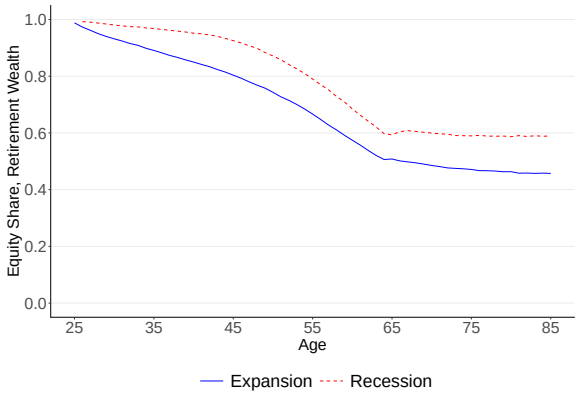
age 30 to 60, III.b shows that both early in working life and during retirement households hold most of their wealth in liquid assets. As agents accumulate financial wealth, they accumulate more in retirement accounts than in liquid wealth so that the average (and median) household maintains more than 70% of its financial wealth in retirement accounts between ages 40 and 55. 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 at each age only have about half their wealth in retirement accounts in their early 40's and have much less at most other ages.

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 share of equity in retirement accounts is quite similar to that provided by TDFs. Figure IV.a shows that the share of financial wealth at the beginning of life is below 30%, a low level driven entirely by non-retirement wealth and due primarily to the need for liquidity. But the average equity share rises rapidly with age to over $80% by the early 40's, then declines to roughly 60% at retirement and remains steady thereafter. Figure IV.b shows that the optimal equity share in retirement accounts looks quite different, and is strikingly similar to that proscribed by target date funds. For young households, the average equity share in retirement accounts is 100%. That optimal share declines to reach 80% at age 45, and thereafter declines more rapidly to just over 50% at retirement with only a slight decline therafter.

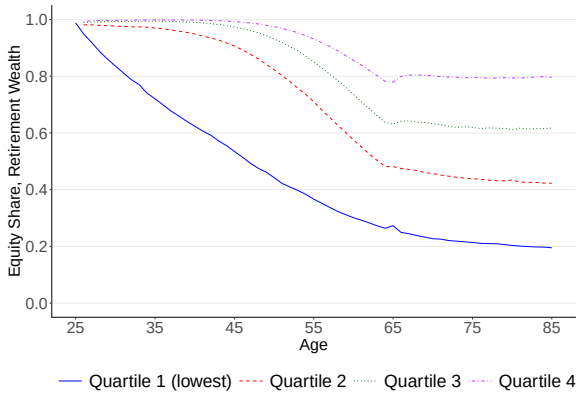
Importantly however, there is a lot of variation in the optimal equity share across people at each age, variation that is far from the homogeneous allocations provided by current

Figure V: 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



TDFs. Figure IV.a shows that the 90th percentile of the optimal equity share in financial wealth rises rapidly early in life to reach 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 90%. In contrast, the 10th percentiles of optimal equity share of retirement wealth in equity falls from 100% to below 40% before age 50, and reaches 20% at retirement, declining only slightly more thereafter.

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

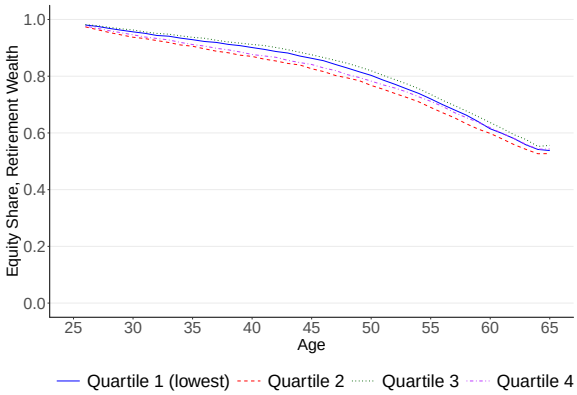
4.2 Optimal Equity Shares of Retirement Wealth by State

This subsection shows how the substantial variation in optimal equity share at each age shown in Figure IV relates to differences in economic circumstances. We focus on equity shares in retirement accounts, but Appendix Figure A.1 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 IV.a instead of Figure IV.b.

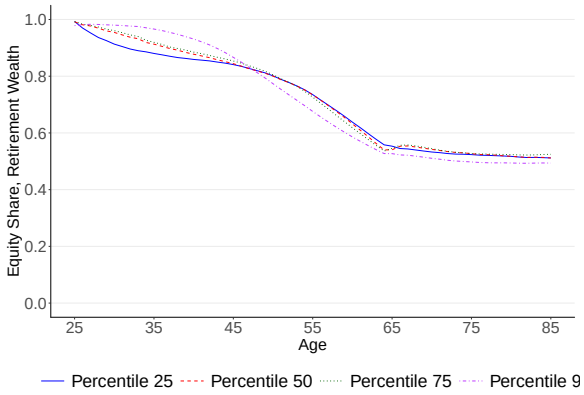
The largest cause of differences in optimal portfolio at each age are differences in the aggregate state variables that determine expected returns. Figure V.a shows that equity shares are about 10% higher in a recession than an expansion. Differences in response to variation in the dividend price ratio are even larger. Figure V.b shows that when the

Figure VI: Optimal Equity Shares of Retirement Wealth

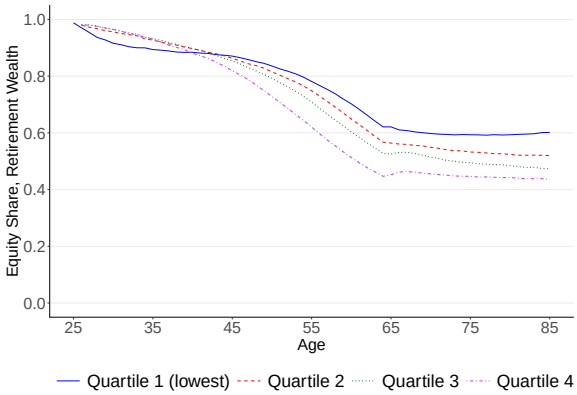
(a) Equity Share by Transitory Income Shock



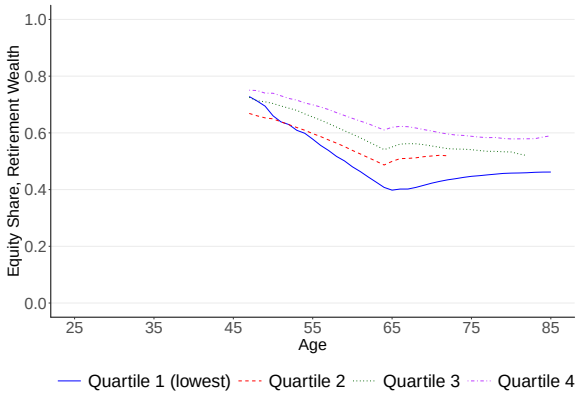
(b) Equity Share by Expected Income Profile



(c) Equity Share by Total Financial Wealth



(d) Equity Share by Housing Leverage



dividend price ratio is in the bottom quarter of its distribution – so expected stock returns are high – the share of retirement wealth invested in equity should be below 40% on average for households over 45 and as low as 20% for households over 70. When the dividend-price ratio is in the top quarter of its distribution, the equity share should be 100% for these households.

In contrast, there are only small variations in the optimal portfolio of equity when households receive good or bad transitory income shocks (Figure VI.a). There are also only small differences across percentiles of the income distribution which are associated with different lifecycle profiles of average income (Figure VI.b).

Figure VI.c shows more significant differences in portfolio allocations associated with different wealth levels, with wealthier households maintaining higher equity shares until around age 40 and lower equity shares thereafter, relative to high wealth households. There

are two reasons that high wealth households hold less of their retirement wealth in equity. First, the risk aversion in the bequest function is higher (because of the intercept) than in the flow utility function, so high wealth households are more risk aversion once likely to leave bequests. Second, high wealth households have less bond-like pension income relative to their accumulated assets, so they hold more of their accumulated in bonds.

Figure VI.d shows that at retirement, homeowners in the top quartile of housing leverage hold about 10% more of their retirement portfolios in equity than households in the bottom quartile. Finally, 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.2).[24]

5 The Welfare Costs and Benefits of Simple Portfolio Rules

Suppose 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 simply portfolio rules? Such simple rules are embedded in Target Date Funds (TDFs) which now serve as default investment options in most employer-sponsored defined-contribution retirement plans.

We begin by considering how much worse the household would do if their retirement saving were allocated to stocks and bonds following the mix proscribed by current TDFs. Specifically, following closely the design of the Vanguard Target Retirement Fund series, we impose that up to age 40 retirement wealth is invested 90% in stocks and 10% in bonds. 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 49.5% at age 65, and then declines by 3% per year to reach 32% at age 70. The equity share remains at 30% from age 71 onward.[25]

Panel A of 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. Columns 1 – 5 show these figures from the perspective of the household at age 25, both for the average household

[24] Also, between age 35 and 55, as households transition from renting to home ownership, renters should consistently hold very close to 100% of their retirement wealth in equity while the optimal equity share for homeowners declines linearly from 93% to 78% over these ages on average.

[25] TDFs do not all follow the same glide path, and do not all deliver the same returns, so that 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).

Table II: Consumption certainty equivalent loss of imposing simple 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
					Panel A: Impose TDF on retirement portfolio					
None	0.74	0.60	0.74	1.02	0.74	2.61	2.33	2.70	2.97	2.65
Liq. portfolio	0.81	0.63	0.80	1.12	0.83	3.19	2.85	3.23	3.84	3.11
All	0.81	0.63	0.80	1.14	0.84	3.43	3.10	3.55	4.64	3.24
					Panel B: Impose constant 2/3 equity share on retirement portfolio					
None	0.79	0.59	0.81	1.10	0.81	2.64	2.14	2.71	3.38	2.78
Liq. portfolio	0.82	0.60	0.83	1.16	0.87	2.73	2.27	2.94	3.18	2.92
All	0.82	0.60	0.84	1.18	0.87	3.06	2.40	3.36	4.03	2.91

Notes: Average family is a weighted average of the given percentiles as described in the text (the 25th percentile represents 32.5% of the population, . . . the 90th percentile represents 17.5%.) 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 $\beta = 1$.

and 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.[26]

The utility loss from following a TDF's proscriptions rather than the optimal state-dependent portfolio rule would cost the average household 0.8 percent of all their consumption in present discounted value or from the perspective of age 25, and between 2.6 and 3.4 percent of consumption in any year chosen at random. 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 Panel A shows that the average household loses the equivalent of only 0.74 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,

[26] In this calculation, we simply use the objective function and include an equal percent loss in bequests.

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 2.61 percent of consumption at all ages and in all state of the world.

However, this first experiment is in some ways inconsistent. We are implicitly 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. The TDF might be inferred to do little because the household in the model to some extent 'unwinds' the portfolio imposed by the TDF with their portfolio choice in their non-retirement, liquid account. Thus we run a second experiment where we impose the TDF investment choices also on liquid wealth. As shown in row 2 of Panel A in Table II, in practice this increases the welfare loss from the TDF, by 7 basis points of consumption from the perspective of age 25, and by just under 0.6 percent from the equally-weighted lifetime perspective. The reduction is significant, but not as large as imposing the investment behavior only in the retirement account because most households accumulate the majority of their financial wealth in their retirement accounts and because the cash in advance constraint keeps low wealth households from taking too much 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 its the old 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 own choices of retirement portfolio – the same (conditional on state variables). As shown in the third row of Panel A of Table II, this leads to a loss equivalent to roughly 0.81% of consumption from the perspective of household at age 25 and to a larger 3.43% of consumption for the average household simply averaged across all years of life.

Panel A of Table II shows that all of these losses tend to be largest for households 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 modelled), this is a segment of the population that has leaned most heavily on TDFs. Finally, Panel B of Table II shows that the glide path of the TDF does not actually improve outcomes for the average investor relative to the pre-existing advice to maintain an age-independent share of wealth in equity, advice

still embedded in balanced funds. Panel B of Table II shows that the consumption-based welfare losses of a constant equity share are sometimes slightly larger than the equity shares in the TDF, but the constant equity share for some households is better than the equity shares embedded in the TDF.

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 from more individualized financial advice or from more customized TDFs.

6 Concluding Remarks

We see four important areas for future work, all of which we view as omissions that are potentially quantitatively significant. First, it would be interesting to also consider the choice between traditional and Roth IRA/401k's. We currently assume that Roth-type investments are not available to the investor, but instead that all retirement saving is in 'traditional' accounts where contributions are made pre-tax and withdrawals are taxed as income.

Second, we do not consider different exposures of labor income to aggregate returns, nor long-run correlation between dividends and labor income. The labor income of different occupations, industries, and skill sets appear to be differentially exposed to fluctuations in both the dividend price ratio and dividends. Our current analysis contains this differential exposure (through variances and serial correlations) but assumes that investors are unaware of their specific covariances (and do not learn about them). With respect to the long-term relationship between dividends and labor income, our current analysis completely omits it. We hypothesized that this relationship would have little effect on portfolio choice because of the small correlation between individual income risk and aggregate returns.

Third, we consider a canonical time-separable flow utility function with constant relative risk aversion. It would be interesting to see how robust our results are to variations in this assumption. In particular, Meeuwis (2019) shows that risk aversion appears to decline with age and that this tends to increase equity shares as wealth and consumption on average increase over the lifecycle.[27]

The final issue for our analysis is that we have assumed that consumption and saving

[27] Meeuwis (2019) also goes further and allows the agent to suffer flow disutility from anxiety about future uncertainty by adopting Epstein-Zin preferences.

behavior are optimal, and the deviations from this assumption that we analyze are minimal. Presumably if investors make significantly sub-optimal choices in saving, then these choices would affect portfolio choices. While our optimal proscriptive rules can condition on current asset levels, they are optimal under the assumption that future saving behavior is optimal. An important question is then whether – either through learning from advice or the experiences of others – people will arrive at close-to-optimal rates of saving and consumption, or whether our model should also be used proscriptively in providing savings advice to investors.

References

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.

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

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," .

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.

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.

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.

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.

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.

DUARTE, V. (2019): "Machine Learning for Continuous-Time Finance," Working paper.

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

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. AND A. MICHAELIDES (2005): "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *Journal of Finance*, 60, 869–904.

GOMES, F. J., A. MICHAELIDES, AND Y. ZHANG (2021): "Tactical Target Date Funds," *CEPR Discussion Papers 13019*, 869–904.

GUVENEN, F., G. KAPLAN, J. SONG, AND J. WEIDNER (2017): "Lifetime Incomes in the United States Over Six Decades," *NBER Working Paper No. w23371*.

GUVENEN, F., S. OZKAN, AND J. SONG (2014): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122, 621–660.

HEATHCOTE, J., K. STORESLETTEN, AND G. VIOLANTE (2017): "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 132, 1693–1754.

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.

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.

MITCHELL, O. S. AND S. UTKUS (2021): "Target date funds and portfolio choice in 401(k) plans," *Journal of Pension Economics and Finance*, 1–18.

MORNINGSTAR (2020): "2020 Target-date strategy landscape," .

PARKER, J. A., A. SCHOAR, A. COLE, AND D. SIMESTER (2021a): "Household Portfolios and Retirement Saving over the Life Cycle," .

PARKER, J. A., A. SCHOAR, AND Y. SUN (2021b): "Retail Financial Innovation and Stock Market Dynamics: The Case of Target Date Funds," .

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.

SEHNKE, F., C. OSENDORFER, T. RÜCKSTIESS, A. GRAVES, J. PETERS, AND J. SCHMIDHUBER (2010): "Parameter-exploring policy gradients," *Neural Networks*, 23, 551–559, the 18th International Conference on Artificial Neural Networks, ICANN 2008.

SHOVEN, J. B. AND D. B. WALTON (2020): "An Analysis of the Performance of Target Date Funds," *NBER Workingpaper, 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. AND A. G. BARTO (2018): *Reinforcement Learning: An Introduction*, The MIT Press, second ed.

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.

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

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.

Appendix: Proofs and derivations

Detailed notation for retirement objective:

$$\begin{aligned}
&\max\ u(c_t, h_t; 1) \\
&+\mathbb{E}\Bigg[\sum_{f=1}^{T_{max}-T_R} \beta^f \Bigg(\prod_{k=-1}^{f-2} \pi(g^i_{t+k}, i, t+k; q^i_R)\Bigg) \\
&\qquad * \Big(\pi(g^i_{t+f}, i, t+f; q^i_R) u(c_{t+f}, h_{t+f}; 1) + \big(1 - \pi(g^i_{t+f}, i, t+f; q^i_R)\big) z(b_{t+f})\Big)\Bigg],
\end{aligned} \tag{58}$$

where T_{max} is the maximum attainable age (i.e., $\pi(\cdot, \cdot, T_{max} - 1; \cdot) = 0$). In the more complicated case where both individuals are alive at time t, the objective is:

$$\begin{aligned}
&\max\ u(c_t, h_t; 2) \\
&+\mathbb{E}\Bigg[\sum_{f=1}^{T_{max}-T_R} \beta^f \Bigg(\Big(\Pr(\text{one individual alive at time } t+f-1)\Big) \\
&\qquad * \Big(\pi(g^i_{t+f}, i, t+f; q^i_R) u(c_{t+f}, h_{t+f}; 1) + \big(1 - \pi(g^i_{t+f}, i, t+f; q^i_R)\big) z(b_{t+f})\Big) \\
&\qquad + \Big(\Pr(\text{both individuals alive at time } t+f-1)\Big) \\
&\qquad * \Big(\pi(g^1_{t+f}, 1, t+f; q^i_R)\pi(g^2_{t+f}, 2, t+f; q^i_R) u(c_{t+f}, h_{t+f}; 2) \\
&\qquad\quad + \big(1 - \pi(g^1_{t+f}, 1, t+f; q^i_R)\big)\pi(g^2_{t+f}, 2, t+f; q^i_R) u(c_{t+f}, h_{t+f}; 1) \\
&\qquad\quad + \pi(g^1_{t+f}, 1, t+f; q^i_R)\big(1 - \pi(g^2_{t+f}, 2, t+f; q^i_R)\big) u(c_{t+f}, h_{t+f}; 1) \\
&\qquad\quad + \big(1 - \pi(g^1_{t+f}, 1, t+f; q^i_R)\big)\big(1 - \pi(g^2_{t+f}, 2, t+f; q^i_R)\big) z(b_{t+f})\Big)\Bigg),
\end{aligned} \tag{59}$$

where

$$\begin{aligned}
&\Pr(\text{one individual alive at time } t+f-1) = \\
&\prod_{k=-1}^{f-2} \pi(g^1_{t+k}, 1, t+k; q^i_R)\Big(1 - \prod_{k=-1}^{f-2} \pi(g^2_{t+k}, 2, t+k; q^i_R)\Big) + \\
&\Big(1 - \prod_{k=-1}^{f-2} \pi(g^1_{t+k}, 1, t+k)\Big) \prod_{k=-1}^{f-2} \pi(g^2_{t+k}, 2, t+k)
\end{aligned} \tag{60}$$

and

$$\Pr(\text{both individuals alive at time } t+f-1) = \prod_{k=-1}^{f-2} \pi(g^1_{t+k}, 1, t+k; q^i_R) \prod_{k=-1}^{f-2} \pi(g^2_{t+k}, 2, t+k; q^i_R). \tag{61}$$

Unpublished Appendix for

Optimal Portfolio Choice Over the Lifecycle

by

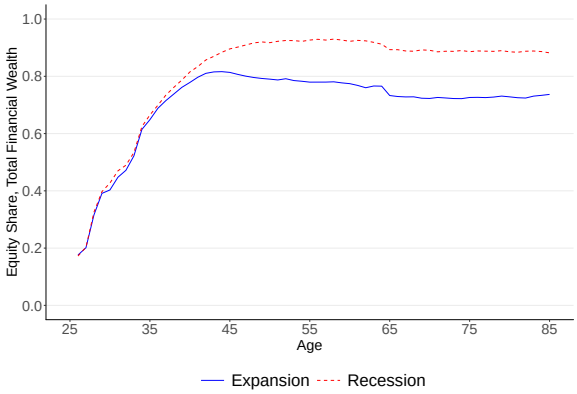
Victor Duarte Julia Fonseca Aaron Goodman Jonathan A. Parker

July 13, 2021

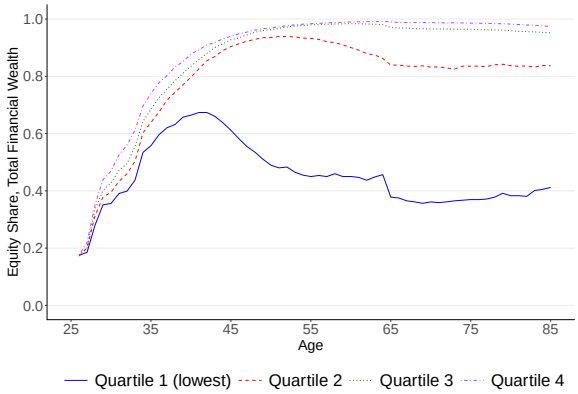
A Additional Figures

Figure A.1: Optimal Equity Shares of Financial Wealth

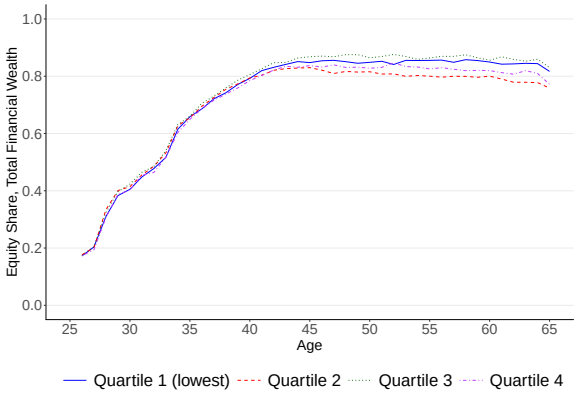
(a) Equity Share by State of the Cycle



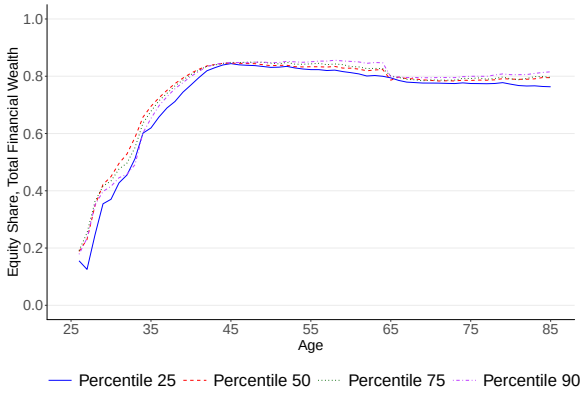
(b) Equity Share by Dividend-Price Ratio



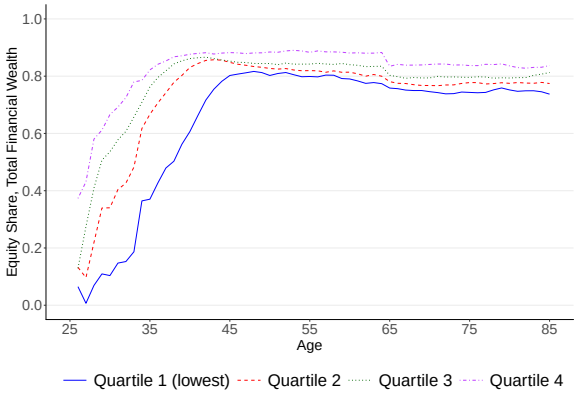
(c) Equity Share by Transitory Income



(d) Equity Share by Permanent Income



(e) Equity Share by Total Financial Wealth



(f) Equity Share by Housing Leverage

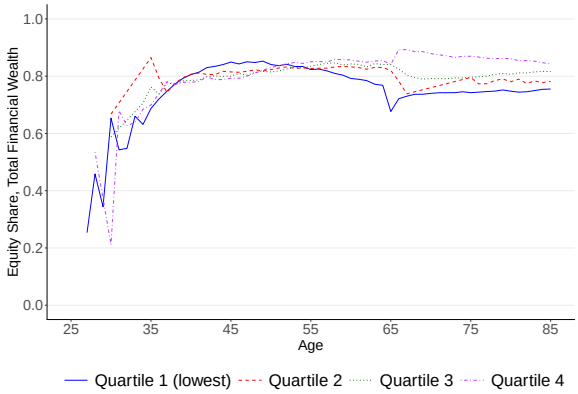
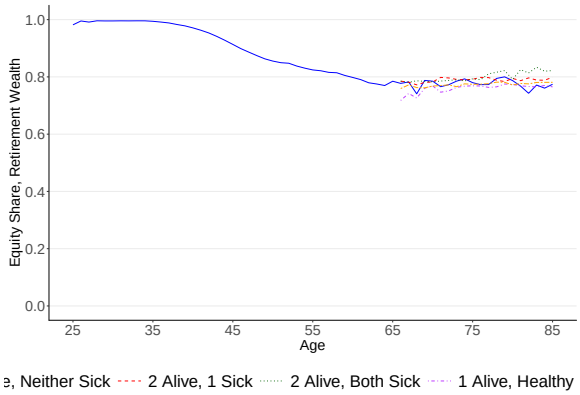


Figure A.2: Optimal Equity Shares of Wealth by Health, Mortality

(a) By Health and Mortality; Retirement Wealth



(b) By Health and Mortality; Financial Wealth

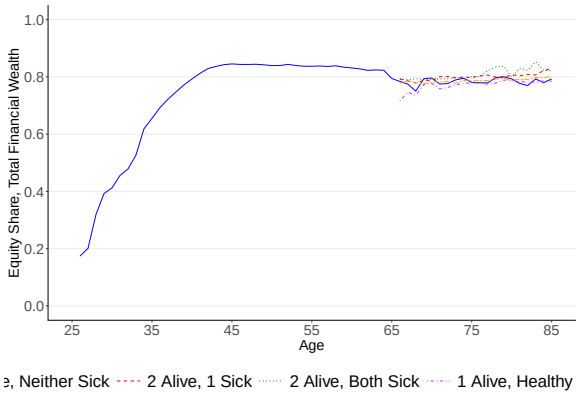


Figure A.3: CAPTION

This figure illustrates the architecture of a neural network representing the portfolio choice at time t of an agent during working life.

```
@tf.function
def train():
  with tf.GradientTape() as tape:
    V, other_outputs = simulate(training=True)
    negative_mean_V = -tf.reduce_mean(V)

  grads = tape.gradient(negative_mean_V, parameters)
  optimizer.apply_gradients(zip(grads, parameters))
  return -negative_mean_V, other_outputs
```

https://www.overleaf.com/project/5e3dbd63f0631c0001bff755

Figure A.4: CAPTION

This figure illustrates the architecture of a neural network representing the portfolio choice at time t of an agent during working life.

```
def brain(n_state_vars):
  state = Input(n_state_vars)
  t = Input(1)
  shared_features = Sequential(
      [BatchNormalization(), Dense(256), Activation('relu'),
       Dense(128), BatchNormalization(), Activation('relu'),
       Dense64), BatchNormalization(), Activation('relu'),
       ])(state)

  inputs = tf.concat([t, shared_features], 1)

  # output layer dimensions
  classes = [4, 2, 4, 4, 5, 3, 4, 6, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3]
  n_classes = len(classes)
  activations = [gumbel_softmax] * 2 + [tf.nn.softmax] * (n_classes - 2)
  latent = tf.split(Dense(sum(classes))(inputs), classes, axis=1)
  actions = [activations[i](latent[i]) for i in range(n_classes)]
  return Model([t, state], actions)
```