

Before and After Target Date Investing: The General Equilibrium Implications of Retirement Saving Dynamics*

Adam Zhang
University of Minnesota

November 2023

Abstract

This paper quantifies the general equilibrium effects of financial innovation that increases access to equity markets. I study an overlapping generations model with both idiosyncratic and aggregate risk, solved with machine learning techniques. A benchmark economy with limited stock market participation and rebalancing frictions matches the current dynamics of macro aggregates, equity and bond returns, as well as wealth and portfolio concentration. A counterfactual experiment shows how widespread adoption of target date funds would improve risk sharing, reduce inequality, and generate substantial welfare gains for households in the bottom 90% of wealth distribution. The equity premium drops from 6.3% to 2.5%, while the standard deviation of equity returns stabilizes from 24.7% to 15.2%. Full adoption of target date funds would generate around 20% average welfare gains for people in the bottom 90% at the expense of the top 10% who lose by more than 50% through the redistribution of financial wealth. Asset pricing and welfare outcomes are very close between an economy with target date funds and one without any participation costs or rebalancing frictions.

*I am indebted to my advisors Monika Piazzesi, Martin Schneider, Michael Boskin and Chris Tonetti for their guidance throughout my graduate career. I also thank Adrien Auclert, Luigi Bocola, Sebastian Di Tella, Robert Hall, Chad Jones, Pete Klenow, Yucheng Yang, and other seminar participants at Stanford macro lunch, macroeconomic workshop, MM reading group, 2023 Minnesota Macro-Asset Pricing Conference, University of Minnesota Carlson School of Management, University of Texas at Austin McCombs School of Business, Rice University Jones Graduate School of Business, UC Santa Cruz Department of Economics, University of North Carolina Kenan-Flagler Business School, and Texas A&M Mays Business School. I gratefully acknowledge the financial support from the Fan Charitable Foundation, as well as the Gale and Steve Kohlhagen Fellowship in Economics and the Bradley Graduate Fellowship through grants to the Stanford Institute for Economic Policy Research. All errors are my own. Email: czadamj@umn.edu.

1 Introduction

Retirement saving dynamics have been changing. Prior to the recent introduction of target date funds, stock market non-participation and inertia were prominent features among retirement savers. Lately, the widespread adoption of target date funds has induced more stock holdings and more frequent rebalancing. To answer how different retirement saving dynamics matter for asset prices, welfare, and inequality requires modeling general equilibrium with aggregate risk and heterogeneous agents. Solving such models, particularly with overlapping generations (OLG), is computationally costly and, in some cases, impossible with conventional techniques.

This paper shows that limited stock market participation and infrequent rebalancing imply high equity premium and equity return volatility, consistent with the data. I show this result using a model with heterogeneous equity access that I discipline with household portfolio data by age and wealth. I solve the model with machine learning techniques to overcome the curse of dimensionality. In addition, the adoption of a simple financial product, target date funds, reduces equity premia and volatility, almost to the extent that these frictions are absent. The equity premium drops from 6.3% to 2.5%, while the standard deviation of equity returns stabilizes from 24.7% to 15.2%. Moreover, target date investing generates welfare gains of approximately 20% remaining lifetime consumption equivalents for households in the bottom 90% of wealth distribution through financial wealth redistribution. These outcomes are comparable to a world without any participation costs or rebalancing frictions.

I study a new rebalancing friction that fixes the asset allocation of flows into portfolios. In particular, agents allocate the flows of savings between equities and bonds using a fixed rule and do not rebalance portfolios. This setup captures three prominent features in household savings. The first feature is that the majority of U.S. households access financial products through retirement accounts, and they rarely change their contribution allocation rules or rebalance their portfolios (Ameriks and Zeldes, 2004, Choukhmane and de Silva, 2022). The second feature is that a substantial fraction of U.S. households do not participate in the stock market (Mankiw and Zeldes, 1991); hence, these households effectively have a flow allocation rule that is 0% in stocks. The third feature is that the very rich households have a stably high equity market share throughout booms and busts. For example, the equity market share of the top 10% richest households is around 80% between 1989 and 2019, which suggests that the richest households do not drastically change their flow allocations over time. The rebalancing friction in this paper differs from the existing literature which constrains portfolio weights and requires rebalancing to the targeted allocation.

Stock market non-participation and infrequent rebalancing generate inelastic asset demands and concentrate equity holdings, implying high stock return volatility, Sharpe ratio,

and equity premia. These frictions dampen the response of stock demand to aggregate shocks. As a result, the stock price responds dramatically to clear the market. Moreover, stockowners who have high savings are massively exposed to aggregate risk because stock returns tend to be higher than bond returns. Without rebalancing, the portfolio share in stocks grows as stockowners age, prompting these agents to demand high compensation for bearing volatile stock returns. The combination of high stock volatility and high price of risk leads to high risk premia in the economy.

After quantification of the benchmark model that features participation and rebalancing frictions, I show that target date funds reduce these frictions almost completely. The benchmark economy matches macro aggregates, equity and bond prices, and portfolio distributions by age and wealth. I then change portfolio choice constraints in the benchmark model to capture financial innovation that reduces these frictions. In the first counterfactual exercise, households by default invest in target date funds which have an age-dependent rebalancing strategy. The second counterfactual exercise removes both participation and rebalancing frictions, and households freely optimize portfolios. Asset prices, welfare, and inequality outcomes are similar under the two alternative asset market arrangements.

The OLG model in this paper connects lifecycle portfolios to asset prices in general equilibrium. The aggregate state of the economy switches between expansions and recessions. Households have time-separable CRRA preferences over consumption, and derive utility from bequests. While working, agents receive labor income, which features an age profile and idiosyncratic risk that is higher in recessions. After retirement, retirees receive social security payments. Households can save in stocks and riskfree bonds, subject to short-selling constraints. Competitive firms produce the consumption good with labor and capital. Firms finance their investments in capital with equities and bonds, choosing their capital structure and payout rules to maximize firm value and to smooth out payouts. The government balances its budget by collecting taxes, financing spending, and supplying government bonds.

The benchmark economy introduces participation and rebalancing frictions to capture stock market non-participation and inertia before target date funds. Specifically, households receive stock market participation shocks that are correlated with income. Before getting hit by a participation shock, households save in bonds only. When a participation shock arrives, a household sets up a contribution allocation rule that fixes the stock-bond ratio for future flows into the portfolio. Households do not actively rebalance portfolios afterwards.

Using parameter values that match household portfolio data, the benchmark model generates realistic macroeconomic, asset pricing dynamics; matches the lifecycle distributions of wealth and portfolio; and produces extreme concentration of equity holdings by wealth. The quantification process of the benchmark model involves two stages. In the first stage, I take

parameters either from the literature or from data. This set of parameters include firm parameters, government parameters, and most household parameters. Then, in the second stage, I use three household preference parameters to target three aggregate wealth moments: average wealth-to-income ratio, retiree wealth share, and top 10% wealth share.

Stock market non-participation and infrequent rebalancing imply that consumption processes observed in the data are compatible with high equity premia, high stock return volatility, and low riskfree rate, which typical consumption based asset pricing models fail to explain. In particular, participation and rebalancing frictions separate the pricing of risky and riskfree rates. The benchmark model in this paper deviates from a standard consumption based asset pricing model because rebalancing frictions prevent participating households from freely optimizing their portfolio weights. As a result, the usual Euler equations for optimal portfolio weights do not hold. Instead, households make consumption and savings decisions, taking portfolio weights as fixed. Therefore, their Euler conditions hold for returns on their portfolios, which are mixtures of the risky and riskfree rates. Nonparticipants, who tend to be low-wealth agents and who have strong precautionary savings motives, price the riskfree rate.

Participation and rebalancing frictions help the model produce realistic wealth and portfolio distributions. In the model, due to consumption smoothing incentives, agents save while they are working and dissave in retirement, which leads to a hump-shaped wealth age profile as seen in the data. Portfolio share in equity at any age, conditional on participation, is a consequence of the initial allocation rule and subsequent market outcomes. The model-implied equity market share by age closely track the data. Furthermore, the positive correlation between income and participation in the model implies that equity holders tend to be wealthy individuals who received lucky draws of income shocks and who have been accumulating assets at the equity return rate at a premium. Therefore, the model-implied equity holdings are even more concentrated than wealth.

I show the adoption of a simple financial product, target date funds, can mitigate or undo participation and rebalancing frictions, almost to the extent that the frictions are absent. To show this, I consider two alternative asset market arrangements. In the target date economy, households invest in target date funds by default but still face the same rebalancing frictions. As a result, all portfolios follow the target date strategy. The free access economy then further drops rebalancing frictions, allowing free choices of portfolio allocation at all times.

In the target date economy, where everyone by default invests in target date funds, the average annual equity premium is 2.5%, as opposed to 6.3% in the benchmark economy. The annualized standard deviation of equity returns drops from 24.7% in the benchmark economy to 15.2%. This stabilization in equity returns is one reason for the fall in the equity premium. The second reason is that the aggregate Sharpe ratio declines from 0.255 in the benchmark

economy to 0.164, suggesting that target date investing improves risk sharing relative to the benchmark economy.

Moving from the benchmark economy to the target date economy generates welfare gains for the bottom 90% of wealth distribution but inflicts welfare losses for the richest 10% of households at old ages. On average, agents in the bottom 90% gain around 20% remaining lifetime consumption equivalents (the majority of which is redistribution), while the top 10% agents lose up to 60% at old ages. The bottom 90% of households benefit both from increased equity market participation and from stabilized equity returns. In contrast, the dramatic reduction in the equity premium leads to much lower returns for households in the top 10% who are stock market participants in the benchmark economy. The richest 10% of households suffer welfare losses as a consequence.

The target date outcomes are very close to an economy where agents freely optimize portfolios, which I call the “free access economy.” In the free access economy, the equity premium declines further to 2.0%, equity returns become even less volatile with a 14.7% standard deviation, risk sharing improves, and the Sharpe ratio is 0.099. The similar asset pricing dynamics between the free access economy and the target date economy lead to similar welfare outcomes.

The target date and the free access economies are comparable in asset prices and in welfare for two reasons. The first reason is that the two economies both induce better risk sharing by redistributing equity share towards the young and towards the poor, who tend to be non-participants in the benchmark economy. Moreover, general equilibrium stabilizes stock returns, which mutes welfare differences among the two economies from portfolios deviations.

I apply machine learning tools to address the technical challenge of solving the high dimensional OLG model. The individual state variables are age, equity holdings, bond holdings, productivity, and contribution allocation rule. Just like any other heterogeneous agent model with aggregate risk, the challenge stems from the fact that agents need to keep track of the distribution of individual states as a state variable. This distribution function is an infinite dimensional object. Traditional techniques, such as approximating the distribution function with histograms or moments selected by the modeler, do not work well when OLG is present. I adapt a machine learning based algorithm, DeepHAM, and reduce the dimensionality of the problem (Han, Yang and E, 2021).

The algorithm has two components. In the first component, I replace the cause of the model’s high dimensionality, the distribution of individual states, with generalized moments. This method differs from the Krusell and Smith (1998) approach in two aspects. Firstly, the moments are more general than standard moments (e.g., mean, second moments, etc). Secondly, I instruct the computer to choose the moments rather than specifying these moments

ex-ante. The reason why approximating the distribution object with moments is sufficient is because agents do not interact directly with each other but rather interact through the market. Thus, instead of keeping track of how each individual matters for one another, I can focus on how each agent matters for aggregate dynamics. Given that the order of the agents does not matter, taking the moments suffices (Kahou, Fernández-Villaverde, Perla and Sood, 2021).

In the second component, reinforcement learning fits neural networks that parameterize these generalized moments and policy functions. Neural nets are functions that are flexible enough to approximate any continuous function, if sufficiently deep and wide (Cybenko, 1989, Hornik, Stinchcombe and White, 1989, Leshno, Lin, Pinkus and Schocken, 1993, Pinkus, 1999, Lu, Pu, Wang, Hu and Wang, 2017). In reinforcement learning, an artificial intelligence (AI) assumes the role of an agent and "lives" in the model environment, trying to maximize utility by adjusting neural nets that represent policy functions and generalized moments. After learning for a sufficiently long period of time, the AI produces the correct policy functions and the correct moments. In particular, I demonstrate that, after training, the computer has learned to distinguish wealthy from poor agents, young from old agents even for the same asset holdings.

Related Literature. This paper contributes to the existing literature on four fronts. Firstly, this paper connects general equilibrium with an extensive literature documenting inertia and stock market non-participation in household portfolio. In addition, this paper proposes and studies a new rebalancing friction, bridging the literature on access frictions in financial markets with empirical facts along two dimensions of heterogeneity: age and wealth. In doing so, this paper speaks to the implications of these frictions for asset prices, for inequality and welfare, and for lifecycle wealth and portfolio dynamics. Thirdly, this paper advances welfare analysis of target date funds from choice problem frameworks to general equilibrium. Last but not least, this paper adds to the literature studying stock prices in OLG economies, joining a series of recent papers that demonstrate success of machine learning-based algorithms in solving heterogeneous-agent models with aggregate risk.

Non-participation in the stock market and infrequent rebalancing are well-known patterns in micro data on U.S. household portfolios. Before the introduction of target-date funds, many households did not participate in the stock market (Blume, Crockett and Friend, 1974, Blume and Friend, 1978, King and Leape, 1985, Mankiw and Zeldes, 1991, Poterba and Samwick, 1995, Vissing-Jorgensen, 1998, 2002a,b, Agnew, Balduzzi and Sundén, 2003, Ameriks and Zeldes, 2004).¹ Moreover, many households select the portfolio allocation of their retirement plan contributions at enrollment and do not make any later changes to their contribution alloca-

¹Ameriks and Zeldes (2004) analyze Surveys of Consumer Finances between 1962 and 2001. They estimate the upper bound for stock market participation during this period to be 29.6% (1962), 43.7% (1983), 47.5% (1989), 49.6% (1992), 54.0% (1995), 57.0% (1998), and 59.7% (2001).

tion. More generally, households rarely rebalance their portfolios (Samuelson and Zeckhauser, 1988, Madrian and Shea, 2001, Choi, Laibson, Madrian and Metrick, 2002a,b, Agnew, Balduzzi and Sundén, 2003, Ameriks and Zeldes, 2004, Beshears, Choi, Laibson and Madrian, 2009, Brunnermeier and Nagel, 2008, Biliias, Georgarakos and Haliassos, 2009, Calvet, Campbell and Sodini, 2009, Mitchell, Mottola, Utkus and Yamaguchi, 2009, Bianchi, 2018).²

This paper also shows that non-participation in the stock market and infrequent portfolio rebalancing are quantitatively important for understanding risk premia and volatility in asset markets. These ideas go back to early work by Mankiw and Zeldes (1991) who document that data on consumption growth by stockholders is more volatile than consumption growth by non-participants. Early theoretical work assumes that non-stockholders save in bonds which are in zero net supply (Saito, 1995, Basak and Cuoco, 1998). For the bond market to clear, stockowners must therefore hold leveraged positions in stocks, which imply high risk exposures by few investors and therefore higher risk premia. Allen and Gale (1994) endogenize the participation decision with a fixed cost for participation. Vissing-Jorgensen (2002a) estimates these participation costs to be large. Heaton and Lucas (1996) study how transaction costs increase the equity premium in equilibrium. Guvenen (2009) adds heterogeneity in preferences as well as stochastic labor income of non-stockholders, which further concentrates risk exposures among stockholders. Also related is Gabaix and Koijen (2021) who demonstrate theoretically and empirically that inelastic asset demand can help understand high asset return volatility. This paper focuses on changes in participation and investment patterns over the lifecycle, and I study a new rebalancing friction that fixes the portfolio weights for flows.

Infrequent rebalancing goes back to Grossman and Laroque (1990) who introduce adjustment costs in consumption, which implies that assets are illiquid. As a consequence, households want compensation for holding illiquid assets in equilibrium, in addition to demanding the standard compensation for aggregate risk taking that is familiar from frictionless consumption-based models. Lynch (1996) analyzes the quantitative importance of these liquidity premia in discrete time, while Gabaix and Laibson (2002) derive analytical solutions in continuous time. Chien, Cole and Lustig (2012) demonstrate that infrequent rebalancing is also quantitatively important for understanding the high volatility in the Sharpe ratio of stock market and its countercyclicality. They study an economy with households that differ in their adjustment costs: some continuously rebalance, while others do so infrequently. This paper analyzes rebalancing frictions in a model with strong age heterogeneity. Moreover, I study whether target date funds can help address these frictions.

Finally, this paper adds to a literature that shows age is an important source of heterogene-

²For example, Ameriks and Zeldes (2004) study a 10-year panel dataset of retirement accounts in the U.S. They find that 73% of plan participants made no change to portfolio asset allocation during the ten years, and an additional 14% made only one change in ten years.

ity to understand equity valuations. [Abel \(2003\)](#) connects the baby boom with stock prices and shows that social security can potentially affect national saving and investment. [Geanakoplos, Magill and Quinzii \(2004\)](#) argue that population booms and busts cause bull and bear stock markets. [Gârleanu and Panageas \(2015\)](#) highlight the potential in preference heterogeneity across age cohorts to resolve some key asset pricing puzzles.

The recent introduction of target date funds has fundamentally changed the landscape of retirement investing. [Mitchell and Utkus \(2022\)](#) document that many households now invest in target date funds because their retirement plans enroll them into these plans as a default option. [Parker, Schoar, Cole and Simester \(2022\)](#) find that target date funds encourage stock market participation, especially among younger savers, and induce a decreasing age profile of stock holdings. These findings are in stark contrast with retirement portfolio patterns prior to the target-date era.

This paper studies the general equilibrium implications of target date investing in a world in which many households do not participate in the stock market and rarely rebalance their portfolios. The existing literature on the introduction of target-date funds studies consumption-portfolio choice problems with exogenous asset returns. Moreover, the literature compares the welfare of target date investing with optimal portfolio choice in the absence of any frictions in stock market participation and portfolio rebalancing. For example, [Gomes, Kotlikoff and Viceira \(2008\)](#) solve a lifecycle model with endogenous labor supply and conclude that the introduction of target date funds does not change welfare much relative to the optimal portfolio case. [An and Sachdeva \(2021\)](#) emphasize the costs associated with using the wrong vintage of target date funds, possibly due to incorrect assumptions about retirement age. [Duarte, Fonseca, Goodman and Parker \(2021\)](#) develop a machine-learning algorithm to compute a lifecycle model with inelastic labor supply and rich heterogeneity. They find that target date funds lower welfare. [Gomes, Michaelides and Zhang \(2022\)](#) find that target date funds should not just focus on selecting age-dependent portfolio shares but also exploit stock return predictability. In contrast, this paper shows that target date investing improves risk-sharing and reduces wealth inequality in an equilibrium with limited stock market participation and infrequent portfolio rebalancing. The equilibrium with target date funds has lower risk premia and asset price volatility, as well as higher welfare of households in the bottom 90% of the wealth distribution.

Quantitative papers that study equity valuation in OLG economies with aggregate risk have mostly used a version of the [Krusell and Smith \(1998\)](#) approach that finds a self-confirming equilibrium where agents form beliefs about a set of moments selected by the modelers. [Storesletten, Telmer and Yaron \(2007\)](#) study idiosyncratic risk and risk premia in an OLG economy with production and incomplete markets. [Favilukis \(2013\)](#) jointly considers the rise in wage inequality, decrease in stock market participation costs, and relaxation of borrowing

constraints. He finds that these observations have led to the sharp rise in wealth inequality, declines in interest rate and in equity premium. One exception in the quantitative strand of the literature is [Leombroni, Piazzesi, Schneider and Rogers \(2020\)](#) who solve for the temporary equilibrium of a model with exogenous expectations to study the entry of baby boomers into asset markets and inflation disagreement across age cohorts. The paper takes the joint distribution of income and initial endowments by age directly from the data, and feeds in survey forecasts to study equilibrium asset prices, wealth, and portfolios. In this paper, I use machine learning tools designed to approximate the rational expectations equilibrium.

To solve the model numerically, I join a series of recent papers in solving heterogeneous-agent models with aggregate risk by using machine learning tools. [Kahou, Fernández-Villaverde, Perla and Sood \(2021\)](#) develop a deep learning algorithm that exploits symmetry in heterogeneous agent models and construct a concentration of measure in evaluating high-dimensional expectations. [Maliar, Maliar and Winant \(2021\)](#) solve dynamic economic models by reducing them into nonlinear regression equations fitted with neural networks. [Azinovic, Gaegauf and Scheidegger \(2022\)](#) design deep equilibrium neural nets that approximate functional rational expectations equilibria and demonstrate success in solving models with significant amount of heterogeneity, uncertainty, and occasionally binding constraints. Most closely related is the DeepHAM algorithm proposed in [Han, Yang and E \(2021\)](#). I use the method to solve a general equilibrium model in which the state space includes the distribution of individual states over a continuum of OLG households, and aggregate risk affects the distribution.

The remainder of the paper has the following layout. Section 2 sets up an OLG model of retirement savings in general equilibrium, under three asset market arrangements: benchmark economy, target date economy, and free access economy. Section 3 describes and evaluates a machine learning based algorithm that overcomes the curse of dimensionality. Section 4 quantifies the model and discusses why the benchmark model does not have puzzles that are common among consumption based asset pricing models. Section 5 compares outcomes across the three economies, in asset prices, inequality, and welfare. Section 6 concludes.

2 OLG Model with Idiosyncratic and Aggregate Risk

This section describes an overlapping generations model with idiosyncratic labor productivity shocks and aggregate risk. Firms and the government endogenously supply assets, which does not complete markets. Households consume and choose a portfolio of assets for their savings.

2.1 The Environment

To capture the booms and busts of the macroeconomy and asset returns, the model contains aggregate risk in continuous time, $t \in [0, \infty)$. The advantage of modeling in continuous time is that certain decision problems admit closed-form solutions, alleviating pressure from the task of model computation.

Aggregate State. The economy goes through expansions and recessions. The state of the economy $Z_t \in \{0, 1\}$ follows a two-state continuous time persistent Markov chain

$$Z_t = \sum_{i=1}^{N_t^Z} \xi_i^Z, \quad (2.1)$$

where N_t^Z is a counting process with intensity $\lambda^Z(Z_{t-})$, and $t-$ is the pre-jump time. Conditional on $Z(T_i^Z-)$ (using δ to denote the Dirac measure), the distribution of the jump size

$$\xi_i^Z \sim \begin{cases} \delta_1, & \text{if } Z(T_i^Z-) = 0 \\ \delta_{-1}, & \text{if } Z(T_i^Z-) = 1, \end{cases}$$

and T_i^Z is the stopping time of the i -th jump. Simply put, $\lambda^Z(Z_{t-})\Delta$ is approximately the probability that the economy switches from its current state Z_{t-} in the business cycle to the other state during the time Δ .

2.2 Household Sector

To model consumption-savings and portfolio decisions through the life cycle, this section introduces a continuum of OLG households that populate the economy. Benchmark and alternative asset market arrangements reflect frictions before and after financial innovations that increase access to equity markets.

Birth, Aging, and Death. A household starts working at age a^{entry} , retires at age a^{retire} , and lives at most until age a^{exit} . The household dies with an age-dependent probability $\eta(a_t)\Delta$ during the time Δ , where a_t is current age of the household at time t . New households enter to replace dying and exiting households, and the population distribution is stationary over time.

Preferences. A household has time-separable CRRA utility over consumption $u(\cdot)$, discounts the future at a constant rate ρ , and derives utility from bequests $u^B(\cdot)$. Therefore, for a con-

sumption process c , the discounted utility of a household at time t can be expressed as

$$E_t \left[\int_t^{t+a^{exit}-a_t} e^{-\rho(v-t) - \int_t^v \eta(a_s) ds} \left(u(c_v) + \eta(a_v) u^B(q_v) \right) dv \right], \quad (2.2)$$

where

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

$$u^B(q) = \frac{\bar{b}(\bar{b}+q)^{1-\gamma}}{1-\gamma}.$$

Income Dynamics. Before retiring, each household inelastically supplies labor and earns labor income. Household productivity $l(a_t, y_t)$ evolves according to a deterministic age profile and an idiosyncratic component y_t . Households are heterogeneous in idiosyncratic labor productivity, $y_t \in \{low, high, star\}$, where the “star” state captures top earners in the economy. Labor productivity y_t switches between the three states according to a Poisson jump process. The probability of switching depends on both the pre-switch productivity state y_{t-} and the aggregate state Z_{t-} . By allowing idiosyncratic and aggregate risk to be correlated, this setup accommodates cyclical movements in labor income risk (Constantinides and Duffie, 1996, Guvenen, Ozkan and Song, 2014). Following Huggett (1996), retired households receive constant social security payment \bar{s} . This assumption is a rough approximation of the progressive replacement rate schedule of the U.S. Social Security program. It has the advantage of dropping earnings history as a household state variable. Thus, at time t , a household receives income $m_t(a_t, y_t)$ which is either labor income $w_t l(a_t, y_t)$ or social security income \bar{s} . The wage rate w_t is the compensation for one efficient unit of labor supply.

A household consumes c_t and saves (or dissaves) s_t from income $m_t(a_t, y_t)$

$$c_t + s_t = m_t(a_t, y_t). \quad (2.3)$$

2.2.1 Asset Market Arrangements in Benchmark Economy

To capture non-participation in the stock market and inertia, the benchmark economy features frictions. Firstly, participation shocks determine whether a household can participate in the stock market. The arrival rate of participation shocks can depend on the productivity of the household. Secondly, households cannot rebalance their portfolios.

Assets. When households start working, they also start saving in a retirement account invested in bonds b_t and stocks e_t

$$n_t = e_t + b_t, \quad (2.4)$$

where n_t is the household net worth.

The riskfree rate is r_t^f . Equity payouts stay invested in equity (Duffie and Sun, 1990, Chien, Cole and Lustig, 2012). The cum-dividend return rate on equity follows

$$dr_t^e = \mu_t^e dt + \sigma_t^e dW_t,$$

where W is a standard Brownian motion, and the drift μ_t^e and the volatility σ_t^e are determined in equilibrium.

The contribution flow s_t into the account splits between bonds and stocks according to the allocation rule f_t , which indicates the fraction of the contribution that households invest in equity. A fraction of households immediately participate in the stock market at a^{entry} and choose a contribution allocation rule. The remaining households start with an allocation rule that has zero weight on equity, $f_t = 0$ at time t when $a_t = a^{entry}$. These households thus initially only save in bonds.

Households who do not participate in the stock market (with $f_t = 0$) may receive a participation shock, which is a counting process N_t^f . The intensity of this counting process may depend on the productivity of the household $\lambda^f(y_{t-})$. Once the household receives a participation shock, the contribution allocation switches to F_t , which the household chooses optimally. The contribution allocation rule of a household that receives a participation shock at time t thus switches from $f_t = 0$ to $f_s = F_t$ for $s > t$.

The jump intensity $\lambda^f(y_{t-})$ is dependent on pre-jump idiosyncratic productivity y_{t-} to capture that high-income individuals are more likely to participate in the stock market. Specifically, "star" earners can always participate and start to hold equity, if they have not already. Households with the *low* productivity state, however, do not receive participation shocks,

$$\lambda^f(y_{t-}) \begin{cases} +\infty & y_{t-} = star \\ \bar{\lambda}^f & y_{t-} = high \\ 0 & otherwise. \end{cases} \quad (2.5)$$

In addition, households do not actively rebalance existing assets.

Withdrawals are proportional to current portfolio shares (Chien, Cole and Lustig, 2012, Choukhmane and de Silva, 2022). Therefore, the effective flow allocation to equity \tilde{f}_t , depending on if the flow s_t is a contribution or a withdrawal, is

$$\tilde{f}_t = \begin{cases} f_t & s_t \geq 0 \\ \frac{e_t}{e_t + b_t} & s_t < 0. \end{cases} \quad (2.6)$$

Equity and bond holdings evolve according to

$$\begin{aligned} de_t &= (\mu_t^e e_t + \tilde{f}_t s_t) dt + \sigma_t^e e_t dW_t \\ db_t &= [r_t^f b_t + (1 - \tilde{f}_t) s_t] dt, \end{aligned} \tag{2.7}$$

and the household net worth n_t evolves by

$$dn_t = de_t + db_t = (\mu_t^e e_t + r_t^f b_t + s_t) dt + \sigma_t^e e_t dW_t.$$

In addition, households cannot short stocks or bonds

$$0 \leq e_t, b_t \text{ and } 0 \leq F_t \leq 1. \tag{2.8}$$

Appendix A formulates the household problem in recursive form.

Bequest Distribution. A small fraction of households receive bequests at age a^{entry} in the form of the average household portfolio. The probability of receiving bequests depends on idiosyncratic labor productivity y_t (Hendricks, 2007, De Nardi, 2004, Wolff and Gittleman, 2014). To capture the fact that a substantial fraction of estates does not pass on as inheritances but rather goes to expenses/charities that are not for production purposes, a certain amount of terminal wealth flows out of the economy (Joulfaian, 1994, Hurd and Smith, 1999, Hendricks, 2001).

2.2.2 Alternative Asset Market Arrangements

This section considers two alternative asset market arrangements. In the first alternative, the target date economy defaults any household savings into an appropriately chosen target date fund. This economy thus features recent financial innovations in how households can save for retirement. In the second alternative, the free access economy allows households to choose their portfolio optimally, without any participation and rebalancing restrictions.

Target Date Economy. Households, by default, invest in target date funds but still face the same rebalancing frictions as in the benchmark economy. Consequently, all household portfolios follow the target date strategy. In other words, a household still chooses consumption c_t and savings s_t to maximize (2.2), subject to budget constraint (2.3). Benchmark portfolio frictions (2.5)-(2.7) instead become

$$\frac{e_t}{e_t + b_t} = T(a_t) \quad \forall t,$$

where $T(a_t)$ is the exogenous target date glide path, the portfolio share invested in equity at age a_t .

The household net worth $n_t = e_t + b_t$ evolves accordingly,

$$dn_t = \{[r_t^f + T(a_t)(\mu_t^e - r_t^f)]n_t + s_t\}dt + \sigma_t^e T(a_t)n_t dW_t. \quad (2.9)$$

Short selling constraints (2.8) stay the same.

Free Access Economy. There are no participation and rebalancing frictions like in the benchmark economy. Households can choose stocks and rebalance their portfolio anytime. As a result, households choose consumption c_t , savings s_t , and portfolio equity share E_t to maximize (2.2), subject to budget constraint (2.3). Benchmark portfolio frictions (2.5)-(2.7) no longer exist. The household net worth $n_t = e_t + b_t$ evolves according to

$$dn_t = \{[r_t^f + E_t(\mu_t^e - r_t^f)]n_t + s_t\}dt + E_t n_t \sigma_t^e dW_t, \quad (2.10)$$

where the drift is the expected return on the portfolio invested in bonds and stocks plus any additional contributions (or minus any withdrawals). Any share E_t invested in stocks contributes to the volatility of net worth because of the volatility σ^N of stock returns. Short selling constraints (2.8) stay the same.

2.3 Production Sector

This section describes the supply of goods and assets in the economy. There is a continuum of identical production firms which decide about their capital structure and payouts. There are closed-form solutions for firms' optimal choices.

Technology. The firms produce consumption goods with capital and labor

$$Y_t = K_t^\alpha L_t^{1-\alpha}.$$

Firms own capital and hire labor at the competitive wage rate w_t .

Capital evolves according to

$$dK_t = [\iota_t - \Phi(\iota_t) - \delta(Z_t)]K_t dt + \sigma K_t dW_t,$$

where $\iota_t = I_t/K_t$ is the investment rate, and investment is subject to adjustment cost

$$\Phi(\iota_t) = \frac{1}{2}\phi[\iota_t - \delta(Z_t)]^2.$$

The depreciation rate $\delta(Z_t)$ is correlated with the aggregate state Z_t , while W_t is a standard Brownian motion that captures quality shocks to capital (Brunnermeier and Sannikov, 2014, Fernández-Villaverde, Hurtado and Nuño, 2019). Expected excess return on capital is the

marginal product of capital minus the riskfree rate, adjustment cost, and depreciation

$$ER_t = MPK_t - r_t^f - \Phi(\iota_t) - \delta(Z_t).$$

Payout and Capital Structure. Firms issue riskfree bonds B_t^f to finance investments in risky capital. Their balance sheets have assets

$$K_t = N_t + B_t^f,$$

where N_t is net worth of firms. Firms' leverage is the ratio of debt to their assets, B_t^f/K_t . Moreover, the ratio of capital to net worth is

$$\omega_t = \frac{K_t}{N_t} = \frac{1}{1 - \text{leverage}_t}.$$

The mapping from leverage B_t^f/K_t to the capital-to-net-worth ratio ω_t is one-to-one, and these two variables move in the same direction.

Firms can rent capital to each other through a competitive rental market and collect rents. Homogeneous firms all make the same decisions, and a representative firm exists. I describe below the problem of the representative firm.

Firms maximize their value and smooth their payouts (Brav, Graham, Harvey and Michaely, 2005, Farre-Mensa, Michaely and Schmalz, 2014). To capture this behavior, I model that the representative firm maximizes the expected present value of log payouts subject to its net worth. The log function captures the incentive to smooth out payouts intertemporally. Specifically, payouts D_t and the capital-to-net-worth ratio ω_t solve

$$\max_{D_t, \omega_t} E_t \left[\int_t^{+\infty} e^{-\bar{\rho}s} \log D_s ds \right], \quad (2.11)$$

subject to the evolution of net worth

$$dN_t = \left(\left[r_t^f + \omega_t ER_t \right] N_t - D_t \right) dt + \sigma \omega_t N_t dW_t. \quad (2.12)$$

The firm maximizes a log objective function which involves the intertemporal smoothing of payouts. To earn the excess return on capital, the firm would like to increase its leverage and thus its ratio of capital-to-net-worth ω_t . However, more leverage also involves more risk and the firm wants to smooth payouts. This trade-off leads to an interior solution for leverage and, hence, for capital-to-net-worth ratio

$$\omega_t \approx ER_t / \sigma^2,$$

which is increasing in the expected excess return on capital but decreasing in quality shock volatility σ . The optimal payout yield is

$$\rho_t = D_t/N_t. \quad (2.13)$$

The optimal payout yield equals to the firm's discount rate $\bar{\rho}$ on average but fluctuates over time due to the adjustment costs. The details of this derivation are in Appendix B. Finally, inflows/outflows from the household sector into the production sector open/close such identical firms. Section 2.5 explains the aggregation in mathematical terms.

2.4 Government

To model inter-generational risk sharing through government programs and to model bond supplies outside of the production sector, this section introduces the government which taxes, transfers, and supplies government bonds.

The government collects income taxes at a constant tax rate τ . In addition, the government borrows by issuing riskfree bonds B_t^g that make up a constant share g of the total bond market. Firms issue B_t^f . The total bond market is then $B_t = B_t^f + B_t^g$.

Fiscal spending has three components: social security payments, debt payments, and discretionary spending G_t . The government adjusts discretionary spending G_t to balance budget

$$\left((1 - \tau) \int_{\mathcal{I}_t} \bar{s} \mathbb{1}_{\{a_{it} > a^{ret}\}} di + r_t^f B_t^g + G_t - \tau w_t L_t \right) dt = dB_t^g,$$

where $i \in \mathcal{I}_t$ indexes households alive in the economy at time t .

2.5 Market Clearing and Aggregation

To prepare for the equilibrium definition, this section first describes the market clearing conditions and the evolution of aggregate variables.

Labor market clears by equating labor demand with labor supply

$$\int_{\mathcal{I}_t} l(a_{it}, y_{it}) di = L_t. \quad (2.14)$$

Bond market clears by equating household bond holdings with corporate and government bond supplies

$$B_t = \int_{\mathcal{I}_t} b_{it} di = B_t^f + B_t^g.$$

Equity market clears by equating household equity holdings with net worth of the firm

$$N_t = \int_{\mathcal{I}_t} e_{it} di = N_t. \quad (2.15)$$

The numeraire good market clears by Walras's Law.

In the benchmark economy, the aggregate inflow of equity from the household sector to the production sector is

$$F_t^e = \int_{\mathcal{I}_t} \tilde{f}_{it} s_{it} di + D_t, \quad (2.16)$$

which includes new purchases of stocks and reinvested payouts. The aggregate inflow of bonds is

$$F_t^b = \int_{\mathcal{I}_t} (1 - \tilde{f}_{it}) s_{it} di. \quad (2.17)$$

The market clearing conditions plus aggregated resource constraint

$$D_t + w_t L_t = C_t + F_t^e + F_t^b + G_t$$

imply that aggregate capital evolves according to

$$\begin{aligned} K_t &= \left[\iota_t - \Phi(\iota_t) - \delta(Z_t) - O_t \right] K_t dt + (F_t^e + F_t^b) dt + \sigma K_t dW_t \\ &= \left[\frac{Y_t - C_t - G_t}{K_t} - \Phi(\iota_t) - \delta(Z_t) - O_t \right] K_t dt + \sigma K_t dW_t, \end{aligned}$$

where O_t is the rate at which resources flow out of the economy because some estates do not pass on as bequests, as described in Section 2.2.

2.6 Equilibrium

This section describes the recursive competitive equilibrium of the economy. The equilibrium definition clarifies how prices and allocations operate in compatibility with supply and demand that arise from maximization problems laid out in previous sections.

In the benchmark economy, the household-specific individual state variables are age, positions in the bond and equity markets, contribution allocation rule, and idiosyncratic labor productivity. In the target date economy and in the free access economy, the household state variables are age, net worth, and idiosyncratic labor productivity. Denote individual state variables (and suppressing time subscript) of households as x , and the associated distribution is $\varphi(\cdot)$. Aggregate state variables consist of $X = (Z, W, \varphi)$. The entire collection of household state variables is then $X = (x, X)$.

The equilibrium consists of pricing functions (r, r^f, MPK, w) , household policy functions $(c, s, F$ in the benchmark economy; c and s in the target date economy; c, s, E in the free access economy), and firm policy functions (D, ω, K, L) , such that

- households maximize utility by solving (2.2)-(2.8);
- firms maximize discounted payouts (2.11)-(2.12);
- markets clear for labor, capital, bond, equity, and numeraire good;
- the law of motion for φ holds.

3 Computational Strategy

The high dimensionality of the model requires a computational strategy that is beyond conventional methods. To overcome the curse of dimensionality, this section uses machine learning tools to solve the model and evaluates the performance of the algorithm.

Similar to other heterogeneous-agent models with aggregate risk, the distribution function φ over individual states is an aggregate state variable, an infinite-dimensional object which makes the computation of this class of models challenging. The OLG structure introduces a strong age dimension to the distribution, which makes it difficult to approximate. With the OLG structure in continuous time, there are infinitely many generations present at any point in time. In the benchmark model, there are five individual states, $x \in \mathbb{R}^5$. To solve their optimization problem, households thus need to keep track of the entire distribution φ of age, asset holdings, contribution allocations, and idiosyncratic labor productivity. Current technology is not capable of dealing with value function iteration in a setting with such high dimensionality within reasonable time. A feasible and sensible representation of φ is necessary.

The model solution has two components. In the first component, I reduce the dimensionality of the problem and approximate the distribution φ with its moments. The idea of replacing the distribution φ with some moments of the distribution is familiar from [Krusell and Smith \(1998\)](#). Their paper uses the first moment of the distribution, its mean, to compute an equilibrium model in which approximate aggregation holds. In their setting, the first moment alone is enough to well approximate the rational expectations equilibrium. However, the OLG structure prevents an approximate aggregation ([Krueger and Kubler, 2004](#)). In the setting of this paper, the first moment is not sufficient to approximate the rational expectations equilibrium.

To select the moments of the distribution φ , I use a machine learning algorithm. The algorithm instructs the computer to choose generalized moments

$$\tilde{\varphi} = E[\mathcal{G}(x)], \tag{3.1}$$

where \mathcal{G} is a basis function. In the case that \mathcal{G} is a polynomial function, $\tilde{\varphi}$ consists of standard moments (first, second, third moments, etc). \mathcal{G} can also be more general than polynomials, hence the name generalized moments (Han, Yang and E, 2021).

The intuition for why replacing φ by moments is sufficient to approximate the rational expectations equilibrium is that agents do not interact with each other but rather interact through the market. Thus, instead of focusing on how each individual matters for one another, keeping track of how each individual matters to the aggregate dynamics is sufficient. The interaction form in equation (3.1) is common in the mean-field literature. In the typical application in this literature, one generalized moment is sufficient. For the computation of the OLG model in this paper, the algorithm chooses two generalized moments.

In the second component of the model solution, reinforcement learning fits neural networks that parameterize the basis function \mathcal{G} and the policy functions (Han, Yang and E, 2021). In this component, the computer simulates the model environment with a cross section of agents. An artificial intelligence (AI) lives in the simulated environment and maximizes realized lifetime utility along simulated life paths. In attempts to maximize utility, the AI learns the utility-maximizing policy functions and the correct generalized moments.

Algorithm 1: DeepHAM (Han, Yang and E, 2021) Adapted to an OLG Economy

Input : 1) initialized neural nets \mathcal{C}^0 and \mathcal{G}^0 for policy functions and basis function; 2) duplicates of the two neural nets \mathcal{C}^{dup} and \mathcal{G}^{dup}

- 1 **for** $k = 1, 2, \dots, N^k$ **do**
- 2 simulate a panel of OLG agents (with replacement) for $T^B + T^E$ periods, using \mathcal{C}^{k-1} and \mathcal{G}^{k-1} (distributions of agents from $T^B + 1$ to $T^B + T^E$ represent the ergodic distribution of the economy, whereas the first T^B periods are burnouts)
- 3 **for** $m = 1, 2, \dots, N^m$ **do**
- 4 set $\mathcal{C}^{dup} = \mathcal{C}^{(k-1)N^m+m-1}$ and $\mathcal{G}^{dup} = \mathcal{G}^{(k-1)N^m+m-1}$
- 5 draw initial state variables of OLG agents $X_{\mathcal{I},0}$ from the ergodic distribution
- 6 initialize state variables of a single agent $X_{i,0}$ at age a^{entry}
- 7 **for** $t = a^{entry}, \dots, a^{exit}$ **do**
- 8 update state variables $X_{\mathcal{I},t+1}$ using \mathcal{C}^{dup} and \mathcal{G}^{dup}
- 9 update state variables $X_{i,t+1}$ using $\mathcal{C}^{(k-1)N^m+m-1}$ and $\mathcal{G}^{(k-1)N^m+m-1}$
- 10 collect realized utility for the single agent $u_{i,t}$
- 11 **end**
- 12 update neural nets to obtain $\mathcal{C}^{(k-1)N^m+m}$ and $\mathcal{G}^{(k-1)N^m+m}$, based on collected and discounted utility u_i for the single agent
- 13 **end**
- 14 **end**

Output: trained policy functions and basis function $\mathcal{C}^{N^k \times N^m}$ and $\mathcal{G}^{N^k \times N^m}$

Algorithm 1 shows the pseudo code of the computational strategy. There are a total of two

sets of two neural nets involved, with each set containing a policy neural net \mathcal{C} and a basis function neural net \mathcal{G} . Depths, widths, and activation functions of the two sets of nets are identical. The first set will go through reinforcement training, whereas the second set is for storage purposes.

The training process involves two loops. The outer loop prepares the ergodic distribution of the economy, by simulating for a long enough period of time which includes burnouts (Judd, Maliar and Maliar, 2011). In the inner loop, the AI and the cross section of OLG agents update their separate sets of neural nets iteratively, in the spirit of fictitious play (Brown, 1951, Han and Hu, 2020, Hu, 2021, Han, Yang and E, 2021). In each play, the OLG agents use neural nets from the previous iteration, whereas the AI tries to figure out the best response to the OLG agents. Specifically, the inner loop initializes by drawing from the ergodic set, copies parameters from the first set of neural nets to the second set, and adjusts the first set of neural nets based on realized utility via stochastic gradient descent. The loss function is the empirical counterpart of (2.2).

The purpose of the cross section of OLG agents is to provide the AI with the model environment from which the AI tries to learn. For this reason, the cross section of OLG agents and the individual AI use separate sets of neural nets to obtain a well defined loss function. Without doing so, general equilibrium prices would become manipulable to the AI who really should take prices as given instead. After training for N^m lifetimes, with each lifetime lasting from a^{entry} to a^{exit} , the algorithm falls back to the outer loop for a new ergodic set of individual states and repeats.

This adaptation deviates from the original DeepHAM algorithm by dropping the value function training. In technical terms, I use a plain vanilla version of policy gradient training as opposed to actor-critic. This decision is for theoretical and practical reasons. On the theoretical side, agents in the Han, Yang and E (2021) setup solve infinite-horizon problems, whereas the setup of this paper has a life cycle component. To obtain the value function at every age means the training must alternate between ages. Moreover, while training the value function recursively makes sense in an infinite-horizon setting, it is not obvious that their recursive definition can easily apply in the finite horizon case. Last but not least, analysis in this paper does not require the value function. On the practicality side, training the value function slows down the algorithm and takes up memory. The OLG plus heterogeneity structure in this paper is heavily demanding in memory. Given the hardware constraints, training additional neural nets for the value function would come at the cost of less accurate simulations by decreasing the size of the cross section. For these theoretical and practical reasons, the adaptation removes the value function training.

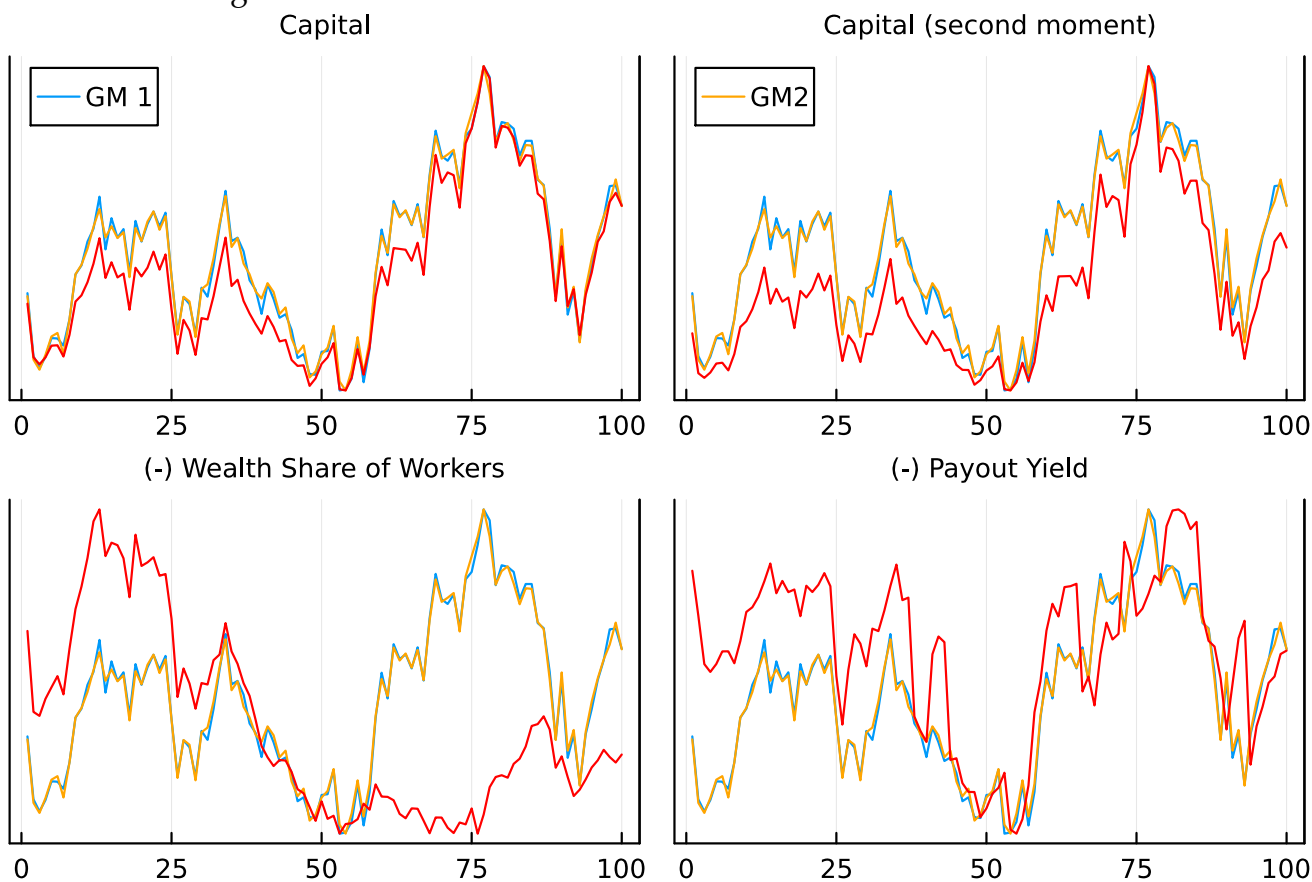
A second deviation of this adaptation is that the inner loop takes place multiple times

before re-simulating for a new ergodic set. In the original algorithm, [Han, Yang and E \(2021\)](#) obtain a new ergodic set after every neural net update. When the cross section is large, as in the setting of this paper which involves OLG and heterogeneity within each age, simulation becomes costly. This training scheme in the adaptation shortens the time spent on simulations.

3.1 Generalized Moment and Statistics of the Economy

This section investigates the relationship between the generalized moments and other statistics from the economy. Commonly used proxies for the distribution object in the literature show resemblance to the generalized moments picked by the computer, although not identical.

Figure 1: Time Series: Generalized Moment and Other Statistics



Notes. Time series plot the generalized moments and other statistics from the economy for 100 quarters. The blue and the orange lines represent the generalized moments, whereas the red line stands for alternative statistics. These statistics include: mean of capital holdings, second moment of capital holdings, fraction of constrained agents (in the bond market), wealth share of workers, and payout yield. Series are scaled.

Figure 1 plots scaled time series of the learned generalized moments and statistics of the economy. The statistics selected in the plot are common choices in the literature as proxies for the infinite dimensional distribution object. These proxies are (from the top-left to the

bottom-right): mean of capital holdings, second moment of capital holdings, wealth share of workers, and payout yield. Many of these statistics move closely with the generalized moments, although none of them completely coincide. Appendix F demonstrates that all statistics above affect the generalized moments, in addition to the capital stock in the economy. Appendix E investigates the learned basis function and finds that age is a strong dimension in the OLG economies, confirming the intuition in Krueger and Kubler (2004).

Among the four statistics shown in Figure 1, mean of capital holdings and payout yield show the closest relationship to the generalized moment. As argued in Storesletten, Telmer and Yaron (2007) who solve an OLG model with aggregate risk, the mean of capital holding has high predictive power for future capital stocks. Payout yield is known to predict stock excess returns (Campbell, 1991, Cochrane, 1992, Ang and Bekaert, 2007). The fact that the generalized moment shows similarities to these statistics indicates that the computer has learned to abstract important information from the distribution of individual states.

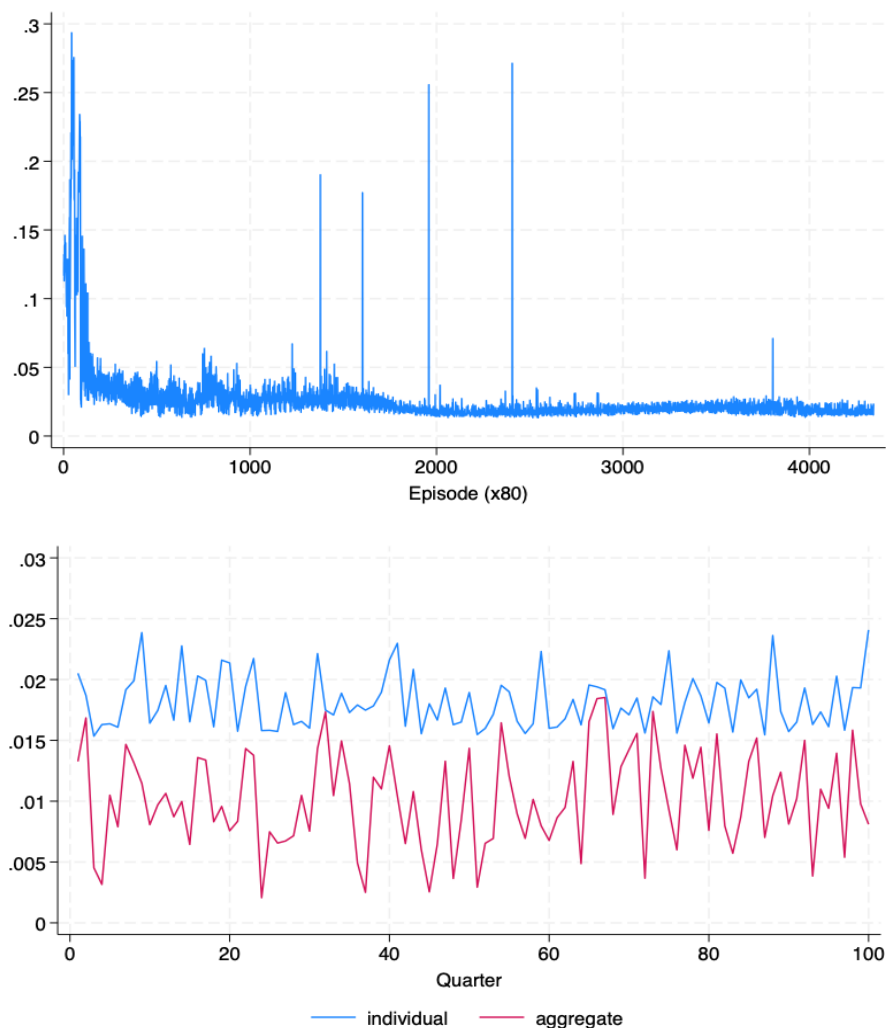
3.2 Implementation and Accuracy Check

This section discusses the implementation details of the training process, which includes the convergence speed and accuracy.

To solve the continuous time model numerically, I use the standard Euler-Maruyama discretization method. Following Krusell and Smith (1998), I use a large number of agents to replace the continuum. The difference is that given there is a strong age dimension in the model, I simulate a large number of agents per age cohort. Specifically, I use 200 agents for each of the 201 ages. The resulting total number of agents is 40200. Each episode takes about 40 seconds, and every 80 episodes make up an epoch. Model is trained for 4344 epochs to reach convergence. Total run time is around 24 days on an NVIDIA A100 graphic card with 40G memory.

After convergence, the neural nets start oscillating around the solution. To check for accuracy, I look at both the average relative consumption errors in the cross section and the average relative consumption errors across the 201 age cohorts. The advantage of using the relative Euler equation error is that this measure is invariant to the magnitude of consumption. The average relative consumption error across individuals is about 1.7% on the ergodic set of the economy, which means that the neural network determined consumption is on average 1.7% different from the Euler equation implied consumption. Since each agent carries trivial weight in a model with a continuum of agents, I also check the average relative aggregate consumption error which is 0.9%.

Figure 2: Average Relative Consumption Errors



Notes. Each epoch contains 80 episodes. The upper figure checks the relative consumption errors across all individual agents who are unconstrained. Unconstrained agents are those who consume less than 96% of all cash on hand (wealth and income combined). The lower figure checks the relative aggregate consumption error.

4 Quantification of the Model

This section discusses the quantification strategy. Macro aggregates and asset pricing dynamics, lifecycle savings/portfolio match empirical observations, untargeted. Fitting these untargeted moments: testifies to the validity of the underlying mechanisms in the model; provides the foundation for studying counterfactual asset market arrangements and the corresponding general equilibrium consequences.

I use a two stage procedure to quantify the model. In the first stage, I select parameter values from the existing literature and match parameters to their data counterparts. In the second stage, I estimate three household preference parameters to match three aggregate

wealth moments.

Table 1 shows the parameter values from the first stage of the quantification. For parameters in the household panel, I measure directly from the data before 2001, which was before the rise of target date funds. For the rest of the economy, I take parameters from the existing literature.

Table 1: Parameters from Literature and Data

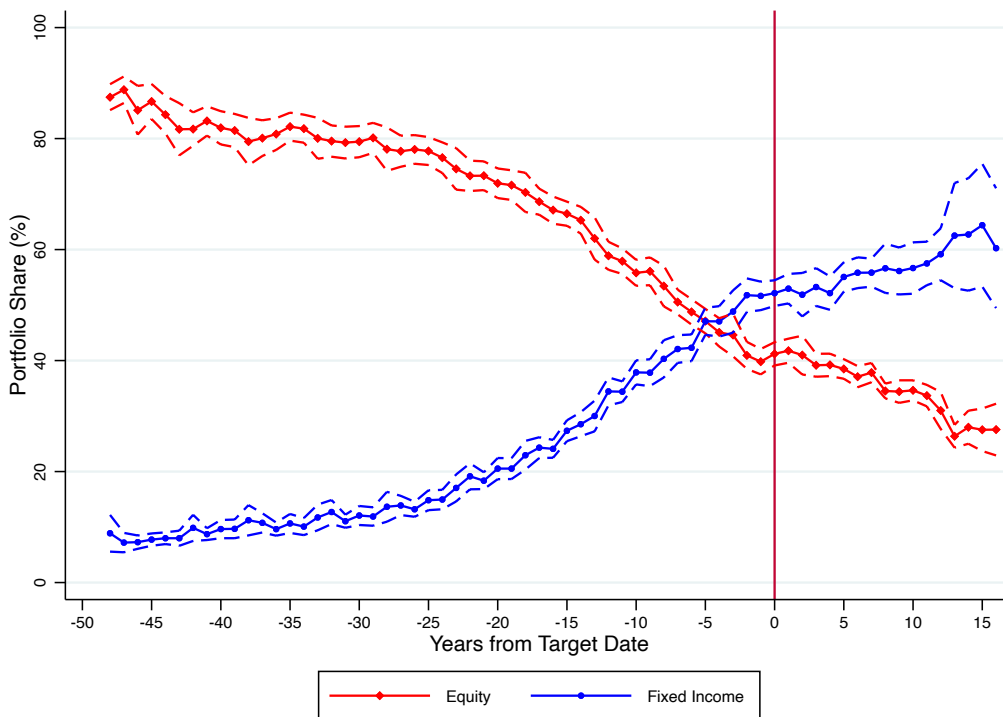
Parameter	Notation	Value	Source
Aggregate State			
switching intensity	$\lambda^Z(\cdot)$	0.125	Krusell and Smith (1998)
Households			
enter, retirement, death age	$a^{entry}, a^{retire}, a^{exit}$	30,65,80	
age distribution			1998 US Mortality Database
mortality risk	$\eta(a)$		1998 US Mortality Database
CRRA	γ	10	
Income age profile			Imrohoroglu, Imrohoroglu and Joines (1995)
labor productivity			Den Haan (2010) and Dávila, Hong, Krusell and Ríos-Rull (2012)
social security participation at a^{entry}	\bar{s}	0.3	SSA (See Appendix C for comparison)
participation intensity	$\bar{\lambda}^f$	0.5	participation rate, age 30
bequest arrival by type		0.002	participation rate, age 50
glide path	$T(\cdot)$	0, 0.05, 0.1	CRSP Mutual Fund Database, 2006-2021
Production Firms			
capital share	α	0.36	Kydland and Prescott (1982)
adjustment cost	ϕ	1	Hall (2002)
capital volatility	σ	0.1	Brunnermeier and Sannikov (2014)
depreciation	$\delta(Z)$	0.09, 0.11	Krusell and Smith (1998)
average payout yield	$\bar{\rho}$	0.049	Fernández-Villaverde, Hurtado and Nuño (2019)
Government			
income tax rate	τ	0.2	De Nardi and Yang (2014)
government bond	\bar{g}	1/3	SIFMA Research

Aggregate State. The switching intensity between expansions and recessions is from Krusell and Smith (1998). On average, expansions and recessions last 8 quarters. The switching probabilities are symmetric between the two states.

Households. The start of the working life, retirement, and death age are 30, 65, and 80 years, respectively. The age distribution and mortality risk by age are from the 1998 U.S. Mortality Database. The risk aversion coefficient is 10, which is the upper limit considered by Mehra and Prescott (1985). The income age profile comes from estimates by Imrohoroglu, Imrohoroglu and Joines (1995). Three idiosyncratic labor productivity states and their transition matrix

come from [Den Haan \(2010\)](#) and [Dávila, Hong, Krusell and Ríos-Rull \(2012\)](#). The fraction of households in the low productivity state is roughly 3% in expansions and 10% in recessions. There are roughly 9% stars in expansions, and 6% in recessions. Social security is constant across agents and across time ([Huggett, 1996](#)). The replacement rate is higher for low income households to reflect progressivity of the Social Security System. As shown in [Appendix C](#), the replacement rates of social security payments in the model are very close to the ones estimated from the data across the majority of the income distribution ([Goss, Clingman, Wade and Glenn, 2014](#)).³ The arrival intensity of participation shocks for the high type matches the share of stock market participants aged 30 years and 50 years, which is when participation rate peaks (Survey of Consumer Finances, 1995-2001). Finally, 5% of high types and 10% of star types receive bequests upon entrance into the economy. [Appendix G](#) shows the (quarterly) transition matrices, states, stationary distributions, and normalization.

Figure 3: Target Date Funds Glide Path



Notes. Center for Research in Security Prices (CRSP) Mutual Fund Database. Annual fund summary file, 2006-2021. Target date funds are identified using Lipper class labels that lead with MAT.

For the target date glide path $T(a)$, I use data from the Center for Research in Security Prices (CRSP) Mutual Fund Database. [Figure 3](#) plots the average portfolio allocations of target

³Admittedly, the model does not capture the extremely high replacement at the very bottom. However, as later elaborated on in [Section 5.2](#), the model is not meant to capture the very bottom of the wealth distribution who do not hold much financial wealth at all. Even though this group of individuals are very important to study, they contribute trivially to the aggregate asset prices.

date funds along with 95% confidence intervals. The y-axis shows the portfolio share in percentage points, with the x-axis being the number of years from the targeted retirement date, with the target date normalized to 0. Target date funds mostly invest in stocks and fixed-income assets. Around 40 years out from the target date, target date funds invest about 80% of their portfolio in stocks. As time approaches the target date, the portfolio share in equity slides lower, reaching 40% at the retirement date, and continues to decline post-retirement.

Production Firms. The capital share is 0.36 as in [Kydlan and Prescott \(1982\)](#). The capital volatility is from [Brunnermeier and Sannikov \(2014\)](#) who also study quality shocks to capital. Previous literature has estimated adjustment cost to range from 0.5 to 2 using disaggregated data ([Shapiro, 1986](#), [Hall, 2002](#), [Cooper and Haltiwanger, 2006](#)). Adjustment cost 1 is approximately the average across industries as estimated in [Hall \(2002\)](#).⁴ The depreciation rate is equal to 10% on average, which is standard in the business cycle literature. Expansions (recessions) increase (lower) the depreciation rate by 1%, which amounts to the same size of the aggregate shock in [Krusell and Smith \(1998\)](#). The average payout yield is 4.9% as in [Fernández-Villaverde, Hurtado and Nuño \(2019\)](#), which is approximately the average payout rate of non-financial corporate businesses in the U.S. according to the Financial Accounts of the United States between 1970 Q1 and 2021 Q4.

Government. Income tax rate τ is 20%, which is in line with the literature ([De Nardi and Yang, 2014](#)). Government bonds make up around one third of the entire U.S. fixed-income asset market during the 1990s, according to the SIFMA Capital Markets Fact Book. So, the fraction of government bond supply as a fraction of the total bond market g equals to 1/3.

4.1 Targeted Moments

In the second stage of the quantification, I estimate household preferences to match moments of the wealth distribution. In particular, I estimate the household discount rate ρ , the bequest function intensity \underline{b} , and the bequest function intercept \bar{b} to match the average wealth-to-income ratio, the share of wealth owned by retirees, and the top 10% wealth share.

The calibration matches the wealth-to-income ratio, retiree wealth share, and top 10% wealth share using the household discount rate ρ , bequest intensity \underline{b} , and bequest intercept \bar{b} .⁵ The data moments exactly identify this GMM estimation. The household discount rate governs the patience of households, which affects their wealth accumulation and is thus closely related to the average wealth-to-income ratio. Bequests matter more for older households than for younger households. A higher bequest motive translates into lower withdrawals during

⁴In [Hall \(2002\)](#), he estimates the average adjustment cost, using quadratic specification, to be 0.91, which is in line with other estimates using micro data ([Shapiro, 1986](#), [Cooper and Haltiwanger, 2006](#)).

⁵Throughout this paper, household wealth is financial wealth which excludes housing and business wealth. The data counterpart is non-housing non-business net worth.

retirement. The bequest function intensity \underline{b} therefore targets the retiree wealth share. Finally, the bequest function intercept \bar{b} determines the bequest size of rich households. The role of the bequest function intensity is to break homotheticity in the household problem by changing the marginal utility of bequests relative to the marginal utility of consumption. For a positive bequest intensity parameter ($\bar{b} > 0$), bequests are a luxury good (Nardi, French and Jones, 2010). For a high value of the bequest intensity, rich agents save disproportionately more out of income compared to other households. The bequest intercept thus targets the top 10% wealth share.

Table 2: Targeted Moments
Targeted Moments

	Top 10% Wealth Share	Wealth- Income Ratio	Retiree Wealth Share
Data	0.695	4.033	0.268
Benchmark	0.610	4.337	0.232

Notes. Survey of Consumer Finances 1995, 1998, and 2001. Household wealth is calculated as non-housing, non-business net worth. Retiree wealth share is for households above age 65. The table excludes households with negative net worth.

Table 2 shows that the model-implied moments are close to their empirical counterparts. The model does not hit these targets exactly, because of long computational times. Another reason for these small mismatches could be the difference between actuarial survival probabilities and subjective mortality beliefs (Heimer, Myrseth and Schoenle, 2019, Grevenbrock, Groneck, Ludwig and Zimmer, 2021).

4.2 Untargeted Moments - Aggregates

The model is able to generate dynamics for macro aggregates and asset prices that compare well with the data. Moreover, the model does not feature standard asset pricing puzzles.

Table 3 shows that the benchmark model does well in matching the dynamics of macroeconomic aggregates and financial variables. The left panel displays the quarterly standard deviations of the growth rates in output, consumption, investments, and labor supply for the benchmark model and the data. The right panel shows the equity premium and its volatility, Sharpe ratio, and leverage. The quantitative fit of the model is reassuring. It indicates that the setup provides a useful tool to study the introduction of target date funds.

The model's asset pricing implications improve upon standard consumption-based asset pricing models. The equity premium is sizable, and the average riskfree rate is low. Moreover,

Table 3: Macroeconomic Aggregates and Financial Moments

	Quarterly SD (Growth Rate)				Annualized Asset Returns			
	Y	C	I	L	$E[r_t - r_t^f]$	$\sigma(r_t - r_t^f)$	SR	leverage
Benchmark	0.017	0.018	0.034	0.010	0.063	0.247	0.254	0.572
Data	0.012	0.012	0.041	0.014	0.066	0.178	0.371	0.560

Notes. The data sample contains 1970Q1 to 2022Q2 (210 quarters). Macroeconomic variables are from the Federal Reserve Bank of St. Louis. All data series are real and seasonally adjusted. Output is the gross domestic product. Consumption is the personal consumption expenditures. Investment is the gross private domestic investment. Labor supply is the hours worked for all employed persons (nonfarm business sector). For the model, simulation period is also 210 quarters. Asset prices come from CRSP value weighted index and the 1 month T-bill rate. Leverage is from estimate in (Graham, Leary and Roberts, 2015) for U.S. public firms in 2010.

the model implies stock return volatility and Sharpe ratio that are comparable with the data. The model slightly overstates the return volatility, implying a lower Sharpe ration than in the data. Overall, the properties of model-implied asset prices closely mirror their empirical counterparts.

The success in matching the equity premium does not come at the cost of unrealistic macroeconomic and financial aggregates. As the left panel shows, quarterly standard deviations for growth rates of output, consumption, investments, and labor supply are roughly consistent with the data. Furthermore, firm leverage in the model is very close to the empirical estimate for U.S. public firms in 2010 (Graham, Leary and Roberts, 2015).

4.2.1 Untargeted Moments - Aggregates: Discussion of Asset Pricing

To give intuition why the benchmark setup does not result in asset pricing puzzles, this section discusses how the benchmark model deviates from a standard consumption-based asset pricing model. Both the demand side of assets and the supply side can help understand the asset pricing dynamics in the model.

The benchmark model is able to simultaneously match asset prices and macro aggregates because participation and rebalancing frictions separate the pricing of the risky and the risk-free rates. This segmentation avoids puzzles seen in common consumption based asset pricing models which price both the risky and the riskfree rates with the same consumption process.

Households can adjust their savings at any time, implying that standard Euler equations hold for the return on savings, which are portfolios of stocks and bonds. For non-stockholders, the portfolio consists only of bonds, so that only the Euler equation for the riskfree rate holds. These households tend to be poor, both because lower productivity households have a lower arrival rate of participation shocks, and because these households earn a lower average return on savings than stock-owners. There are many states of the world in which these poor house-

holds may like to borrow but they face a borrowing constraint. To avoid these states of the world, poor households will save and thereby depress the riskfree rate in equilibrium relative to an economy without participation frictions.

The portfolio of stockholders contains both bonds and stocks, but their Euler equation only holds for the return on the entire portfolio, not for the return on each asset individually. The reason is that rebalancing frictions prevent stockholder-households from adjusting their portfolios. Instead, household choose how much to save and sell their portfolios, with fixed portfolio weights.

A two-agent economy with participation and rebalancing frictions can illustrate the intuition for these asset pricing dynamics at the extreme. Suppose one agent can only hold stocks, while the other agent only holds bonds. Stocks and bonds are in non-zero net supply. In equilibrium, the stockowner will price stocks, and the bond holder will price the riskfree rate. The stockowner has massive exposure to risk and demands a high compensation, pushing up the equity return rate. The intertemporal smoothing motives of the (poorer) non-stockowner will determine the riskfree rate. The segmentation of stock and bond pricing will avoid the well-known asset pricing puzzles seen in a standard consumption-based asset pricing model.

It remains to clarify how the stock owner ends up with the high exposure to risk, which the stylized two-agent model abstracts away from. Stock returns are on average much higher than the riskfree rate. As a consequence, mechanically, a stock market participant's portfolio share in equity increases with time, due to rebalancing frictions. This mechanism concentrates equity holdings even more, in addition to the concentration that participation frictions induce. Moreover, rebalancing frictions prevent stockowners from choosing portfolio shares, so these agents still do not price the riskfree rate. Thus, the segmentation intuition remains intact, even if stockowners in the full model hold both stocks and bonds.

On the supply side, firms invest in risky capital and issue bonds. They freely choose their optimal portfolio of capital and bonds, which implies that standard Euler equations for these assets hold with log preferences. However, firms' preferences are over payout streams which are highly volatile and almost perfectly correlated with returns on capital. As a result, the expected return on capital holdings is high. Since firms are also leveraged, expected returns on (levered) equity are even higher than expected returns on capital. Firms price the riskfree rate low due to high elasticity of intertemporal substitution associated with logarithmic utility.

A direct consequence of rebalancing frictions is that asset demand is relatively inelastic. Recent literature have shown both theoretically and empirically that inelastic asset demand can amplify asset return volatility ([Gabaix and Kojien, 2021](#)). The benchmark model generates inelastic demand because rebalancing frictions prevent households from adjusting asset positions to movements in asset prices. For example, when a bad shock hits the capital stock,

equity price drops. In the absence of rebalancing frictions, households would sell bonds and buy stocks because the expected equity premium is high. This rebalancing behavior pushes up the demand for equities, so the stock price does not have to fall all the way. In the benchmark model, however, asset demand is inelastic. The stock price falls deeper to clear the market. The logic for a positive shock is similar. Overall, the equity price is more volatile in the benchmark model than a typical consumption asset pricing model.

On the supply side, the firm asset demand elasticity is high, but participation frictions induce a high leverage. The result of limited equity holdings is that equity financing is expensive for firms. As a result, firms use a lot of debt financing for their investments, leading to volatile net worth processes of the firms.

4.3 Untargeted Moments - Life Cycle

Because of the OLG structure, the model has implications for lifecycle wealth and portfolio dynamics. This section compares wealth age profile and equity market share by age in the model and the data. Despite not explicitly targeting these lifecycle savings moments, the model matches both distributions very closely.

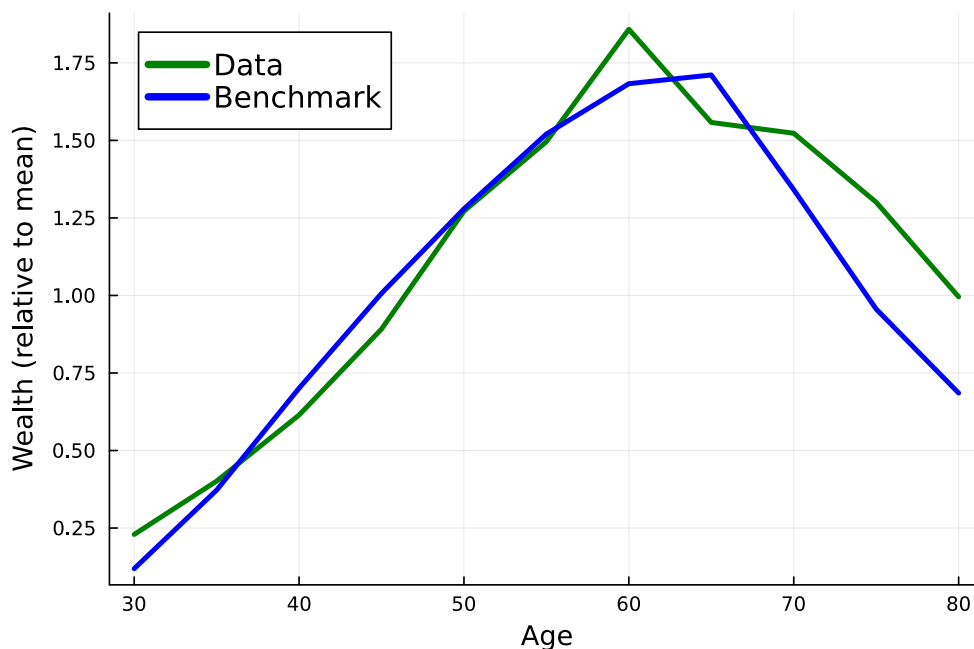
Figure 4 compares the wealth age profile in the model and in the data. Both the model and the data show a hump-shaped wealth age profile. Due to consumption-smoothing incentives, workers save while income is high and draw down savings in retirement. As a result, household wealth peaks around retirement age. During late retirement, the bequest motive becomes strong, and households do not consume all wealth.

Figure 5 studies equity market share by age, defined as the ratio of total equity holdings by an age cohort to the total equities outstanding in the economy. The dashed line shows the distribution in the benchmark model, while the solid line shows the distribution in the data. Mid-life households aged 50-60 years hold around 40% of equity shares in the economy, while younger and older agents hold relatively less. Overall, the distribution of equity holdings by age are similar in the benchmark economy and the data.

The main driver in the model that leads to such a realistic distribution is rebalancing frictions. Recall that the only portfolio-related targets used in quantification are stock market participation rates at age 30 and 50. Conditional on participation, rebalancing frictions in the model imply that portfolio equity share at any age depends only on the initial allocation and the subsequent market outcomes. In fact, Section 5.2 re-visits this plot when asset market arrangements change, the distribution of equity market share by age is drastically different. This finding confirms the importance of rebalancing frictions in understanding household lifecycle portfolio dynamics before the rise of target date funds.

It is important for the benchmark economy to match portfolio holdings by age. This good

Figure 4: Wealth Age Profile



Notes. Survey of Consumer Finances 1995, 1998, and 2001. Household wealth is calculated as non-housing, non-business net worth. Wealth, in both the model and data, has been normalized by the mean household wealth. The graph excludes households with negative net worth.

fit of the model provides a solid foundation for studying how alternative asset market arrangements modify inter-generational risk sharing and the accompanying general equilibrium effects.

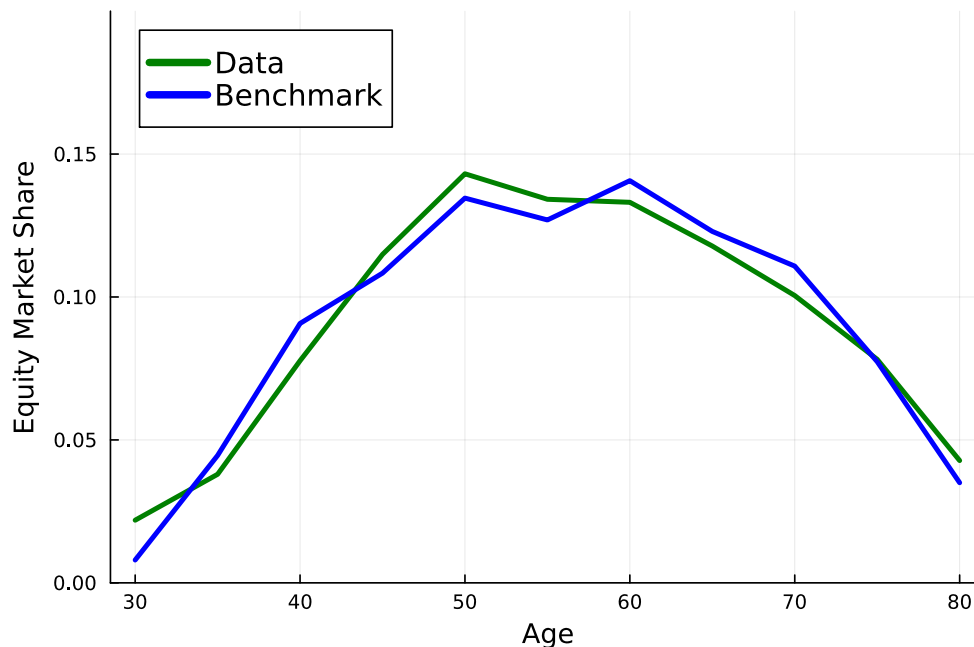
4.4 Untargeted Moments - Inequality

Besides an OLG structure, the model also features idiosyncratic risk and incomplete markets. This setup means that the model bears implications for inequality. This section demonstrates that the model replicates the extremely concentrated equity holdings by wealth in the data, without targeting any moments of the distribution of equity holdings.

Figure 6 breaks down equity shares by wealth in the benchmark economy and in the data. An immediate pattern that stands out from the data is that the distribution of equity holdings is extremely concentrated. The top 10% richest households hold close to 80% of the equity shares in the economy. The model produces the same level of concentration in equity holdings as observed in the data due to frictions in participation and in rebalancing. Recall that the calibration procedure only targets the top 10% wealth share. For equity holdings to be even more concentrated than the wealth distribution, richer households must have higher portfolio shares in equity.

In the model, households who end up at the top of the wealth distribution are equity

Figure 5: Equity Market Share by Age



Notes. Survey of Consumer Finances 1995, 1998, and 2001. Household wealth is calculated as non-housing, non-business net worth. The graph excludes households with negative net worth.

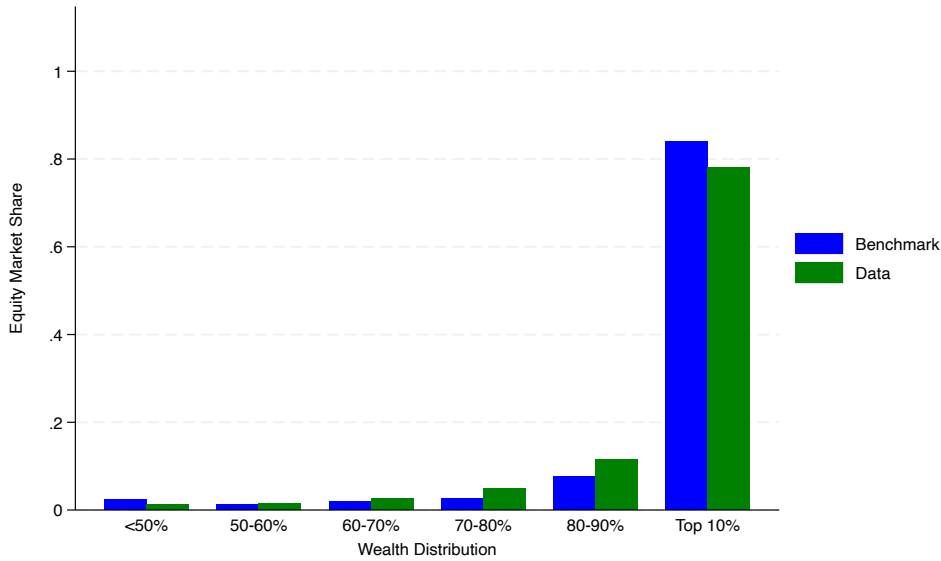
market participants who have been enjoying the equity premium for a long time. The positive correlation between equity participation and idiosyncratic productivity means that equity participants tend to have lucky histories of idiosyncratic productivity draws. In addition, rebalancing frictions imply that, conditional on participation, the portfolio share in equity results from the initial allocation and subsequent market outcomes. Given that equities tend to outperform bonds, the portfolio share in equities trends upwards as agents age. Agents around 50 to 60, just before they start drawing down savings, tend to be the richest households whose portfolios are also high in equities.⁶ Therefore, the distribution of equity holdings is more concentrated than the distribution of wealth.

The model does predict slightly higher equity holdings for the bottom 50%. This is because, in the data, there are impoverished households that do not hold financial products at all. Admittedly, this paper does not explicitly model this group of households, whose welfare is very important to study. From an asset pricing perspective, however, the contribution from households who are limited in investable wealth to general equilibrium asset prices is minimal. The slight mismatch towards the very bottom in Figure 6 does not cause a grave concern for general equilibrium analyses in this paper.

Fitting the cross-section equity holdings by wealth in Figure 6 validates again the portfolio dynamics in the model, which inertia and stock market non-participation govern, for house-

⁶Figure 7 further explores the second mechanism.

Figure 6: Equity Market Share by Wealth



Notes. Survey of Consumer Finances 1995, 1998, and 2001. Household wealth is calculated as non-housing, non-business net worth. The graph excludes households with negative net worth.

holds across the wealth distribution. It is important for the benchmark economy to produce a good fit of the inequality in asset holdings. The model can speak to how alternative asset market arrangements alter inequality measures and the accompanying general equilibrium effects.

5 Counterfactuals: Target Date and Free Access Economy

To assess the implications of inertia and stock market non-participation for asset prices, inequality, and welfare, this section conducts two counterfactual exercises that resemble recent and continued financial innovations that reduce these frictions.⁷ Widespread adoption of target date funds would improve risk sharing, reduce inequality, and generate substantial welfare gains for households in the bottom 90% of the wealth distribution. Outcomes are very close between an economy with target date funds and one without any participation costs and rebalancing frictions.

5.1 Counterfactuals: Asset Prices

This section studies how counterfactual asset market arrangements change equilibrium asset prices. Compared to the benchmark economy, target date investing lowers equity premium, stabilizes equity returns, and decreases the aggregate Sharpe ratio. Results are similar for the

⁷The main quantitative exercises considered here are comparisons between different stochastic steady states. Appendix I investigates outcomes along the transition path.

free access economy.

Table 4: Counterfactuals: Asset Prices

	Annualized Asset Returns						Sharpe Ratio and ω		
	$E[r_t]$	$\sigma(r_t)$	$E[r_t^f]$	$\sigma(r_t^f)$	$E[r_t - r_t^f]$	$\sigma(r_t - r_t^f)$	SR	ω_t	$\sigma(\omega_t)$
Benchmark	0.058	0.247	-0.004	0.007	0.063	0.247	0.255	2.339	0.728
Target Date	0.020	0.152	-0.005	0.008	0.025	0.152	0.164	1.527	0.011
Free Access	0.016	0.147	-0.005	0.008	0.020	0.147	0.143	1.471	0.004

Notes. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone.

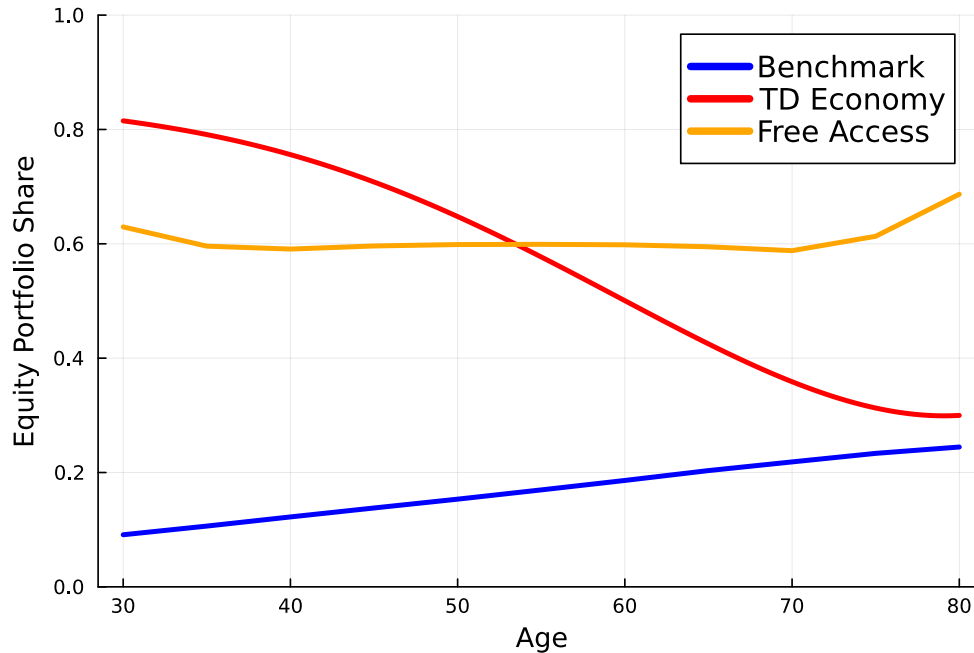
Table 4 compares moments on asset prices for the benchmark economy, the target date economy, and the free access economy. The left panel displays annualized average return rates and standard deviations of, from left to right, equity, bond, and risk premium. The right panel shows the aggregate Sharpe ratio, average capital-to-net worth ratio ω , and its quarterly standard deviation.

Compared to the benchmark economy, the two counterfactual worlds have drastically different asset pricing dynamics. In the target date economy, equity returns are lower and more stabilized: the average equity return rate is 2.5% with standard deviation 15.2%, compared to 6.3% and 24.7% respectively in the benchmark economy. The riskfree rate does not show noticeable differences between the benchmark and the target date economies. Consequently, lower and more stabilized equity returns translate into the smaller and less volatile equity premium in the target date economy. In addition, the aggregate Sharpe ratio diminishes by almost two thirds, diving from 0.255 in the benchmark economy to 0.164 in the target date economy. Accompanying all these changes in asset prices is a sharp decline in the firm leverage, with capital-to-net worth ratio cut to 1.527 from 2.339. Outcomes for the free access economy are very much comparable with those from the target date economy. The average equity return rate, its standard deviation, the equity premium and its volatility, the Sharpe ratio, and the capital-to-net-worth ratio ω drop further, but only to a limited extent.

To understand these movements in asset prices, Figure 7 examines average portfolio age profiles under the three asset market arrangements. In the benchmark economy, the average portfolio share in equity starts off around 20% at age 30 and goes up with age, reaching 35% at age 80. The equity premium mechanically drives most of this pattern. Given that equities tend to outperform bonds, the equity portfolio share goes up as stockowners age due to rebalancing frictions.

Compared to the benchmark economy, the target date economy shows substantially more

Figure 7: Counterfactuals: Portfolio Age Profile



Notes. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone. Y-axis is the average (unweighted) portfolio age profile.

equity holdings across almost all ages, especially for the young. The target date glide path slides from around 80% for 30-year-old agents to about 40% at retirement and continues declining to 25% at age 80. The higher average portfolio share in equities among the young reflects two margins of agents' portfolios: participation rate is lower in the benchmark economy (around 55% across all working ages); conditional on participation, the glide path sets the portfolio equity share higher than an agent would be at in the benchmark economy.

The average portfolio share in equities in the free access economy is similar to the glide path during working ages but differ substantially near retirement. Initially at 65%, the portfolio share in equities stabilizes around 60% across all ages, before bouncing up to around 70% at age 80. The initial decline in the equity portfolio share is a consequence of the decline in non-tradable, relatively safe human capital (Viceira, 2001). The reversal of its course is due to two reasons. Firstly, past a certain age, bequest motive starts to dominate. Secondly, post retirement, agents no longer face risk in social security payments. As retirees draw down risky financial savings, increasing risk exposure becomes optimal (Gomes, Kotlikoff and Viceira, 2008). Hence, post retirement, the equity portfolio share climbs back up.

The riskfree rate does not change as much in the two counterfactual economies due to two opposite forces. In the benchmark economy, stock market non-participants price the riskfree rate. In the target date economy, however, portfolios of all households follow the glide path.

As a result, everyone prices the return on their portfolio, which is a mixture of equities and bonds. The glide path is low in equities for retirees, suggesting mostly retirees price the riskfree rate in the target date economy. Compared to non-participants in the benchmark economy, retirees in the target date economy hold more equities, suggesting more volatile consumption processes. This first force tends to drive down the riskfree rate. In contrast, retirees do not have strong incentives to save. This second force tends to push up the riskfree rate. With the two forces counteracting each other, the riskfree rate stays roughly the same between the benchmark economy and the target date economy.

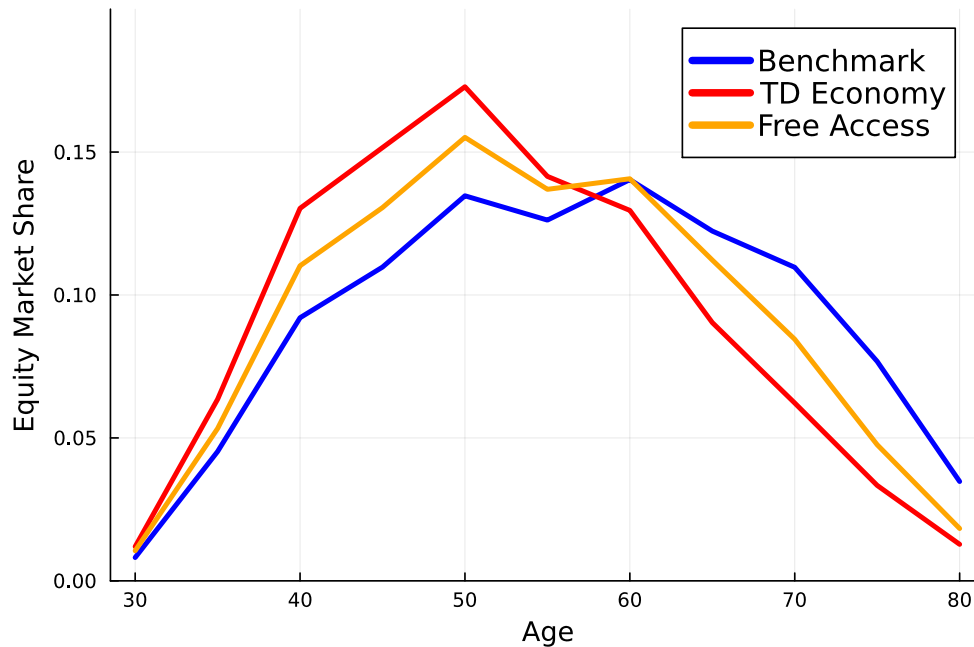
The reduction in the equity return volatility comes from changes in both the demand side and the supply side. Intuition from the demand side involves elastic/inelastic asset demand. As argued above, asset demand is inelastic in the benchmark economy, amplifying the equity return volatility. The asset demand elasticity is higher in the target date economy because households trade against market outcomes to stay on the glide path (Parker, Schoar and Sun, 2020). For example, when stocks outperform bonds, agents sell stocks and buy bonds to restore the equity-bond ratio that the target date glide path mandates. In the free access economy, demand elasticity further increases because agents can choose optimal portfolio weights. The intuition from the supply side is that the economies are less leveraged, which stabilizes equity returns, as equation (2.12) suggests. As demand rises for equities, equity financing becomes cheaper than before. Given the cost of debt financing, riskfree rate, is the same, firms respond by adjusting capital structure in favor of equities. Table 4 shows, the capital-to-net worth ratio ω descends to 1.5 (target date economy) and 1.296 (free access economy) from 2.121 (benchmark economy). Therefore, firms de-leverage, and equity returns become less volatile. Consequently, Table 4 shows the annualized standard deviation of equity dives to 14.6% (target date economy) and 13.1% (free access economy) from 21.9% (benchmark economy).

5.2 Counterfactuals: Risk Sharing

To study why the aggregate Sharpe ratios are lower in the counterfactual economies, this section investigates distributions of equity market share and Sharpe ratios by subgroup. Both the target date economy and the riskfree economy redistribute towards the young and towards households in the bottom 90% of the wealth distribution. Redistribution of equity shares towards the young affects the aggregate Sharpe ratio by altering individual Sharpe ratios, whereas redistribution of equity shares to the bottom 90% is mostly a compositional effect.

Figure 8 plots the equity market share by age for three different asset market arrangements. Compared to the benchmark economy, working agents hold more equity shares in the target date and the free access economies. Specifically, agents below age 60 hold about 65% of equities in the benchmark economy, where as this number jumps to 85% under the target date arrangement and 75% in the free access economy. In other words, both the target date and the

Figure 8: Counterfactuals: Equity Market Share by Age



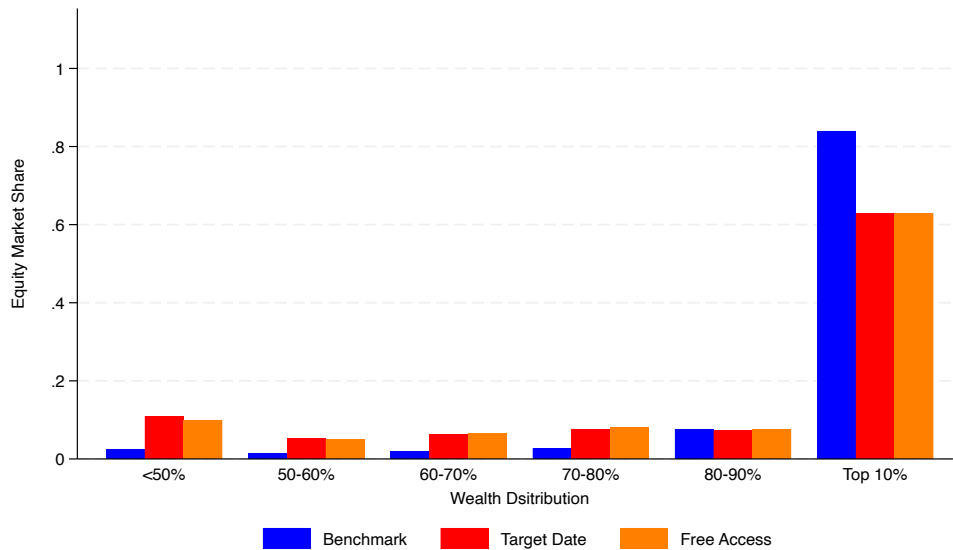
Notes. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone.

free access economies redistribute equity shares towards young workers, to a more aggressive extent under the target date arrangement.

Age patterns of equity shares depicted in Figure 8 are a consequence of increased equity holdings among the young. In fact, Figure 7 hints at these age patterns in the three economies. In the target date economy, both stock market participation and, conditional on participation, the portfolio share in equities are higher than those in the benchmark economy. This is especially true for workers for whom the glide path sets a high portfolio share in equities. As a result, equity shares redistribute towards the young in the target date economy. This redistribution is true to a lesser degree in the free access economy. Compared to the target date glide path, retirees in the free access economy hold higher portfolio shares in equities. Thus, redistribution of equity shares towards to the young is not as dramatic.

Figure 9 then breaks down equity market shares by the other dimension of heterogeneity: wealth. The striking pattern is that the two counterfactual asset market arrangements substantially reduce the concentration in equity holdings. The top 10% wealthiest households take up 84% of equities in the benchmark economy, while their market share falls to 63% in both the target date economy and the free access economy. In the mean time, households in the bottom 90% see consistent gains in equity market shares, particularly for the bottom 50%. On a whole, both alternative asset market arrangements redistribute equity shares to the bottom

Figure 9: Counterfactuals: Equity Market Share by Wealth



Notes. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone.

90%, leading to resembling distributions equity market shares by wealth.

Increased access to equity markets drives the equity shares patterns along wealth portrayed in Figure 9. In the benchmark economy, participation in equity markets is restricted to rich agents who are either bequest receivers or who have enjoyed lucky draws of labor productivity. In contrast, everyone participates in the stock market by default in the target date economy. In fact, all portfolios are on the glide path. Therefore, the target date economy witnesses a drastic reduction the concentration of equity holdings.

To sum up the findings on risk sharing, the target date and the free access economies both redistribute equity market shares towards the young and towards households in the bottom 90% of the wealth distribution.

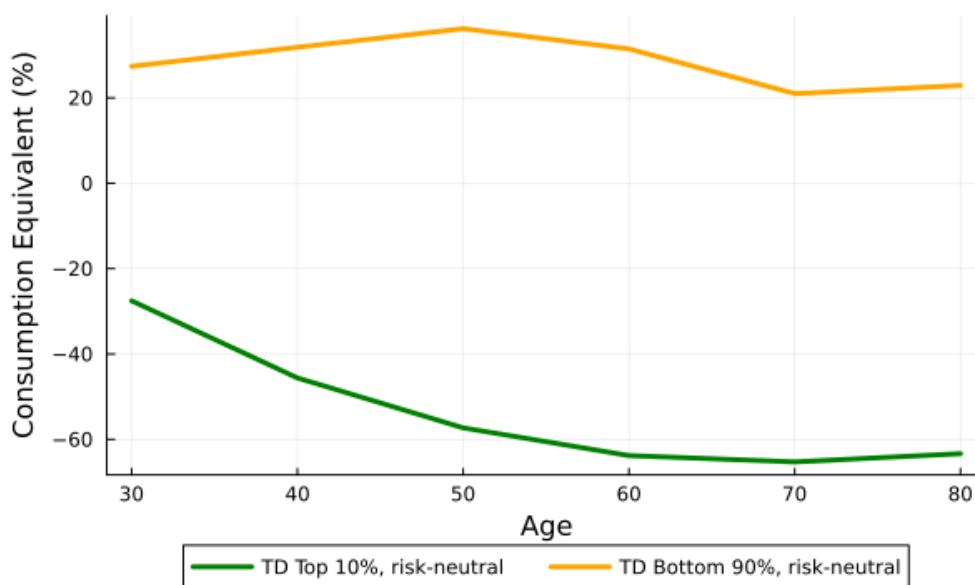
5.3 Counterfactuals: Welfare

This section studies how increased access to equity markets affect welfare through asset prices and risk sharing. On average, target date investing generates welfare gains in remaining life-time consumption equivalent for households in the bottom 90% at the expense of the top 10% richest households. Free access to asset markets results in very similar welfare outcomes compared to the target date economy.

To measure welfare implications, I calculate the remaining life-time consumption equivalents. The interpretation of this measure is the percentage boost in consumption for the rest of

the lifetime, so that a benchmark agent would be just as well-off as an agent from an alternative economy, both of whom are of the same age.

Figure 10: Counterfactuals: Consumption Equivalent - Target Date Economy



Notes. Consumption equivalents are calculated using risk neutral utility. Simulation keeps both aggregate and idiosyncratic shocks identical in the two economies. Cross sections of agents are taken from the ergodic set of the economy. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone. Consumption equivalent is defined for every age, in terms of remaining life-time consumption.

I start by examining consumption equivalents measured using risk-neutral utility. In other words, at any age, the welfare interpretation is the remaining lifetime wealth discounted to the present, using the intertemporal discount rate. Figure 10 plots the average risk-neutral consumption equivalents for all households and also by wealth distribution. Overall, Figure 10 indicates that there is substantial redistribution of wealth from the top 10% richest households to the bottom 90%, especially among the retirees.

The welfare outcome under target date investing implies heterogeneous gains and losses across the wealth distribution. For people in the bottom 90% of the wealth distribution, their remaining lifetime wealth is 20-25% higher than the benchmark case. The main driver is that everyone is participating in the stock market by default using target date funds. Therefore, households who save with riskfree rate only in the benchmark economy, due to participation friction, can accumulate assets at a much higher return rate. The slight dip towards the end of the lifecycle is a result of the target date glide path mandating a lower portfolio share in equities during retirement than the benchmark economy, as Figure 7 demonstrates. Notice that this number is substantially lower than the 50% wealth increase calculated in Mitchell and

Utkus (2022) because general equilibrium effects have substantially lowered stock returns, as suggested in Table 3. Still, the boost in remaining lifetime wealth, as measured by risk-neutral utility, for the bottom 90% of households is huge.

While the bottom 90% are able to accumulate much more wealth in retirement, the top 10% suffer tremendous wealth losses across all ages. In retirement, the top 10% wealth shrinks around 60%, compared to the benchmark economy. Even when retirement wealth is discounted to the beginning of the lifecycle, the top 10% richest households still lose close to 30% of remaining lifetime consumption equivalents. The culprit for these huge wealth and welfare losses is the dramatic reduction in the equity premium. As Table 3 suggests, the top 10% richest households, who are participants in both economies, accumulated wealth at a 6.3% annual rate in the benchmark economy but at only 2.5% in the target date economy.

Inspecting the age profile of wealth by wealth distribution confirms the drastic redistribution that target date investing leads to. Figure 11 plots the average wealth held at an age for the top 10% and the bottom 90% respectively in the two panels. Compared to the extreme unequal wealth distribution induced by participation and rebalancing frictions in the benchmark economy, the target date funds substantially increase wealth holdings by the bottom 90% at the cost of dramatically reducing the top 10% wealth stock, across all ages.

Figure 12 then compares the risk-neutral welfare outcomes with the case when the relative risk aversion is 10, which is the maximum risk aversion parameter that Mehra and Prescott (1985) consider to be acceptable. Examining the gap between the risk averse case against the risk neutral case, the bottom 90% of households enjoy smaller welfare improvements due to holding on to more risk in their portfolios. As people in the bottom 90% of wealth distribution tend to have small buffer stock, the welfare reduction due to higher exposure to risk is more pronounced (Carroll, 1997). Nonetheless, even when wealth gains in retirement are discounted to the beginning of the lifecycle, on average, these agents still gain 18% of remaining lifetime consumption equivalents. The top 10% richest households do not exhibit substantially different welfare outcomes than the risk-neutral case. Overall, using a relative risk aversion parameter as high as 10 leads to the similar conclusion on the welfare effects of target date investing from the general equilibrium.

Figure 13 repeats the welfare analyses in Figure 12 for agents in the free access economy. Opening up equity markets further by dropping rebalancing frictions leads to additional welfare gains during working ages but slightly lower gains at the beginning and the end of the lifecycle. Nevertheless, these differences are minute, less than 5 percentage points in consumption equivalents, compared to target date welfare gains. Much of these improvements comes from overall more equity holdings across all ages (Figure 7) and more stabilized equity returns (Table 4). Given that general equilibrium forces substantially stabilizes equity returns and re-

duces the equity premium in both the target date and the free access economies, differences in portfolio choices do not lead to drastically different welfare results. Welfare outcomes are very similar between target date and free access economies.

6 Conclusion

To conclude, this paper investigates the implications of increased access to equity markets for asset prices, inequality, and welfare. I set up an overlapping generations model with idiosyncratic and aggregate risk to study lifecycle portfolio choices in general equilibrium. I then solve the model by applying machine learning techniques to overcome the curse of dimensionality in solving the model. The benchmark economy features frictions in equity market participation and in rebalancing to replicate portfolio dynamics before the latest financial innovations. Two alternative asset market arrangements that resemble recent innovations then alter the two benchmark frictions one at a time.

Frictions in stock market participation and in rebalancing help explain puzzling asset pricing dynamics. After quantification using portfolio data between 1995 and 2001, the benchmark economy produces realistic dynamics for macroeconomic aggregates and for asset prices, matching wealth and portfolio concentration. The two benchmark frictions distinguish this model from standard consumption-based asset pricing models. Firstly, frictions in participation concentrate equity holdings among the wealthy who have high exposures to risk. Secondly, rebalancing frictions imply participants' Euler conditions hold for returns on portfolios but not for individual assets. For these reasons, the benchmark economy does not result in the classical equity premium/riskfree rate puzzle.

Target date investing improves risk sharing, reduces inequality, and generates welfare gains for the bottom 90%. The equity premium plunges from 6.3% to 2.5%, and the annualized standard deviation of equity returns falls from 24.7% to 15.2%. The stabilization comes as asset demand becomes more elastic, and firms adjust capital structure in response to changes in equilibrium asset prices. In addition, the aggregate Sharpe ratio plummets from 0.255 to 0.164. This result is due to redistribution of equity shares towards the young and towards households in the bottom 90% of the wealth distribution. The richest 10% of households suffer large welfare losses (up to 60% in remaining life-time consumption equivalent) as equity premium falls, while the rest of agents see 20% welfare gains.

Overall, outcomes are comparable between the target date economy and the free access economy. Free access economy removes frictions in participation and in rebalancing altogether, leading to further improvements in risk sharing, a bigger reduction in inequality, and more welfare gains across the economy. The equity premium dips to 2.0%, with annualized standard deviation of equity returns shrunk to 14.7%. The aggregate Sharpe ratio edges lower to 0.143.

All households enjoy welfare gains, in a similar fashion as the target date economy.

Findings in the paper suggest that increasing equity market access has large general equilibrium effects on asset prices, inequality, and welfare. Evaluations of retirement security policies that encourage the adoption of recent financial innovations, such as the 2007 Pension Protection Act and the 2022 Secure Act 2.0 (or known as the RISE & SHINE Act in the senate), should take into consideration these general equilibrium implications.

References

- Abel, Andrew B. (2003) "The Effects of a Baby Boom on Stock Prices and Capital Accumulation in the Presence of Social Security," *Econometrica*, 71 (2), 551–578.
- Agnew, Julie, Pierluigi Balduzzi, and Annika Sundén (2003) "Portfolio Choice and Trading in a Large 401(k) Plan," *American Economic Review*, 93 (1), 193–215.
- Allen, Franklin and Douglas Gale (1994) "Limited Market Participation and Volatility of Asset Prices," *American Economic Review*, 84 (4), 933–955.
- Ameriks, John and Stephen P Zeldes (2004) "How Do Household Portfolio Shares Vary with Age," Technical report, working paper, Columbia University.
- An, Byeong-Je and Kunal Sachdeva (2021) "Missing the Target? Retirement Expectations and Target Date Funds," *Retirement Expectations and Target Date Funds* (December 1, 2021).
- Ang, Andrew and Geert Bekaert (2007) "Stock Return Predictability: Is it There?," *Review of Financial Studies*, 20 (3), 651–707.
- Azinovic, Marlon, Luca Gaegauf, and Simon Scheidegger (2022) "Deep equilibrium nets," *International Economic Review*.
- Basak, Suleyman and Domenico Cuoco (1998) "An Equilibrium Model with Restricted Stock Market Participation," *Review of Financial Studies*, 11 (2), 309–341.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian (2009) "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States," in *Social Security Policy in a Changing Environment*, 167–195: National Bureau of Economic Research, Inc.
- Bianchi, Milo (2018) "Financial Literacy and Portfolio Dynamics," *Journal of Finance*, 73 (2), 831–859, [10.1111/jofi.12605](https://doi.org/10.1111/jofi.12605).
- Biliás, Yannis, Dimitris Georgarakos, and Michalis Haliassos (2009) "Portfolio Inertia and Stock Market Fluctuations," CEPR Discussion Papers 7239, C.E.P.R. Discussion Papers, <https://ideas.repec.org/p/cpr/ceprdp/7239.html>.
- Blume, Marshall E, Jean Crockett, and Irwin Friend (1974) "Stock Ownership in the United States: Characteristics and Trends," *Survey of Current business*, 54 (11), 16–40.
- Blume, Marshall E and Irwin Friend (1978) *The Changing Role of the Individual Investor: A Twentieth Century Fund Report*: Wiley.
- Brav, Alon, John R. Graham, Campbell R. Harvey, and Roni Michaely (2005) "Payout policy in the 21st century," *Journal of Financial Economics*, 77 (3), 483–527.
- Brown, George W (1951) "Iterative Solution of Games by Fictitious Play," *Activity Analysis of Production and Allocation*, 13 (1), 374.

- Brunnermeier, Markus K. and Stefan Nagel (2008) "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-evidence on Individuals," *American Economic Review*, 98 (3), 713–736.
- Brunnermeier, Markus K. and Yuliy Sannikov (2014) "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104 (2), 379–421.
- Calvet, Laurent E., John Y. Campbell, and Paolo Sodini (2009) "Measuring the Financial Sophistication of Households," *American Economic Review*, 99 (2), 393–398.
- Campbell, John Y (1991) "A Variance Decomposition for Stock Returns," *Economic Journal*, 101 (405), 157–179.
- Carroll, Christopher D (1997) "Buffer-stock saving and the life cycle/permanent income hypothesis," *The Quarterly journal of economics*, 112 (1), 1–55.
- Chien, YiLi, Harold Cole, and Hanno Lustig (2012) "Is the Volatility of the Market Price of Risk Due to Intermittent Portfolio Rebalancing?," *American Economic Review*, 102 (6), 2859–2896.
- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick (2002a) "Defined Contribution Pensions: Plan Rules, Participant Decisions, and the Path of Least Resistance," JCPR Working Papers 257, Northwestern University/University of Chicago Joint Center for Poverty Research, <https://ideas.repec.org/p/wop/jopovw/257.html>.
- (2002b) "For Better or For Worse: Default Effects and 401(k) Savings Behavior," JCPR Working Papers 256, Northwestern University/University of Chicago Joint Center for Poverty Research, <https://ideas.repec.org/p/wop/jopovw/256.html>.
- Choukhmane, Taha and Tim de Silva (2022) "What Drives Investors' Portfolio Choices? Separating Risk Preferences From Frictions," Technical report, Working paper.
- Cochrane, John H (1992) "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, 5 (2), 243–280.
- Constantinides, George M and Darrell Duffie (1996) "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104 (2), 219–240.
- Cooper, Russell W and John C Haltiwanger (2006) "On the nature of capital adjustment costs," *The Review of Economic Studies*, 73 (3), 611–633.
- Cybenko, George (1989) "Approximation by Superpositions of A Sigmoidal Function," *Mathematics of Control, Signals and Systems*, 2 (4), 303–314.
- De Nardi, Mariacristina (2004) "Wealth Inequality and Intergenerational Links," *The Review of Economic Studies*, 71 (3), 743–768.
- De Nardi, Mariacristina and Fang Yang (2014) "Bequests and heterogeneity in retirement wealth," *European Economic Review*, 72 (C), 182–196.
- Den Haan, Wouter J. (2010) "Comparison of solutions to the incomplete markets model with aggregate uncertainty," *Journal of Economic Dynamics and Control*, 34 (1), 4–27.

- Duarte, Victor, Julia Fonseca, Aaron S. Goodman, and Jonathan A. Parker (2021) "Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle," NBER Working Papers 29559, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/29559.html>.
- Duffie, Darrell and Tong-sheng Sun (1990) "Transactions costs and portfolio choice in a discrete-continuous-time setting," *Journal of Economic Dynamics and Control*, 14 (1), 35–51.
- Dávila, Julio, Jay H. Hong, Per Krusell, and José-Víctor Ríos-Rull (2012) "Constrained Efficiency in the Neoclassical Growth Model With Uninsurable Idiosyncratic Shocks," *Econometrica*, 80 (6), 2431–2467.
- Farre-Mensa, Joan, Roni Michaely, and Martin Schmalz (2014) "Payout Policy," *Annual Review of Financial Economics*, 6 (1), 75–134.
- Favilukis, Jack (2013) "Inequality, stock market participation, and the equity premium," *Journal of Financial Economics*, 107 (3), 740–759.
- Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuño (2019) "Financial Frictions and the Wealth Distribution," NBER Working Papers 26302, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/26302.html>.
- Gabaix, Xavier and Ralph Koijen (2021) "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis," cepr discussion papers, C.E.P.R. Discussion Papers.
- Gabaix, Xavier and David Laibson (2002) "The 6D Bias and the Equity-Premium Puzzle," in *NBER Macroeconomics Annual 2001, Volume 16*, 257–330: National Bureau of Economic Research, Inc.
- Geanakoplos, Author-Name: John, Michael Magill, and Martine Quinzii (2004) "Demography and the Long-Run Predictability of the Stock Market," *Brookings Papers on Economic Activity*, 35 (1), 241–326.
- Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira (2008) "Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds," *American Economic Review*, 98 (2), 297–303.
- Gomes, Francisco, Alexander Michaelides, and Yuxin Zhang (2022) "Tactical Target Date Funds," *Management Science*, 68 (4), 3047–3070.
- Goss, Stephen, Michael Clingman, Alice Wade, and Karen Glenn (2014) "Replacement rates for retirees: What makes sense for planning and evaluation," *Social Security Administration Actuarial Note Number 155*.
- Graham, John R., Mark T. Leary, and Michael R. Roberts (2015) "A century of capital structure: The leveraging of corporate America," *Journal of Financial Economics*, 118 (3), 658–683.
- Grevenbrock, Nils, Max Groneck, Alexander Ludwig, and Alexander Zimper (2021) "Cognition, Optimism, And The Formation Of Age-Dependent Survival Beliefs," *International Economic Review*, 62 (2), 887–918.

- Grossman, Sanford J and Guy Laroque (1990) "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58 (1), 25–51.
- Guvenen, Fatih (2009) "A Parsimonious Macroeconomic Model for Asset Pricing," *Econometrica*, 77 (6), 1711–1750.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song (2014) "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122 (3), 621–660.
- Gârleanu, Nicolae and Stavros Panageas (2015) "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing," *Journal of Political Economy*, 123 (3), 670–685.
- Hall, Robert E (2002) "Industry dynamics with adjustment costs."
- Han, Jiequn and Ruimeng Hu (2020) "Deep Fictitious Play for Finding Markovian Nash Equilibrium in Multi-Agent Games," in *Mathematical and Scientific Machine Learning*, 221–245, PMLR.
- Han, Jiequn, Yucheng Yang, and Weinan E (2021) "DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks," *arXiv preprint arXiv:2112.14377*.
- Heaton, John and Deborah J Lucas (1996) "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104 (3), 443–487.
- Heimer, Rawley Z., Kristian Ove R. Myrseth, and Raphael S. Schoenle (2019) "YOLO: Mortality Beliefs and Household Finance Puzzles," *Journal of Finance*, 74 (6), 2957–2996.
- Hendricks, Lutz (2001) "Bequests and Retirement Wealth in the United States," *Unpublished, Arizona State University* http://www.lhendricks.org/Research/bequedata_paper.pdf.
- (2007) "Retirement Wealth and Lifetime Earnings," *International Economic Review*, 48 (2), 421–456.
- Hornik, Kurt, Maxwell Stinchcombe, and Halbert White (1989) "Multilayer Feedforward Networks Are Universal Approximators," *Neural Networks*, 2 (5), 359–366.
- Hu, Ruimeng (2021) "Deep Fictitious Play for Stochastic Differential Games," *Communications in Mathematical Sciences*, 19 (2), 325–353.
- Huggett, Mark (1996) "Wealth Distribution in Life-cycle Economies," *Journal of Monetary Economics*, 38 (3), 469–494.
- Hurd, Michael D. and James P. Smith (1999) "Anticipated and Actual Bequests," nber working papers, National Bureau of Economic Research, Inc.
- Imrohoroglu, Ayse, Selahattin Imrohoroglu, and Douglas H Joines (1995) "A Life Cycle Analysis of Social Security," *Economic Theory*, 6 (1), 83–114.
- Joulfaian, David (1994) "The Distribution and Division of Bequests in the US: Evidence from the Collation Study," *OTA paper*, 71.

- Judd, Kenneth L, Lilia Maliar, and Serguei Maliar (2011) “Numerically Stable and Accurate Stochastic Simulation Approaches for Solving Dynamic Economic Models,” *Quantitative Economics*, 2 (2), 173–210.
- Kahou, Mahdi Ebrahimi, Jesús Fernández-Villaverde, Jesse Perla, and Arnav Sood (2021) “Exploiting Symmetry in High-Dimensional Dynamic Programming,” NBER Working Papers 28981, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/28981.html>.
- King, Mervyn and Jonathan Leape (1985) “Wealth and Portfolio Composition: Theory and Evidence,” CEPR Discussion Papers 43, C.E.P.R. Discussion Papers, <https://ideas.repec.org/p/cpr/ceprdp/43.html>.
- Krueger, Dirk and Felix Kubler (2004) “Computing equilibrium in OLG models with stochastic production,” *Journal of Economic Dynamics and Control*, 28 (7), 1411–1436.
- Krusell, Per and Anthony A. Smith (1998) “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106 (5), 867–896.
- Kydland, Finn E and Edward C Prescott (1982) “Time to Build and Aggregate Fluctuations,” *Econometrica*, 50 (6), 1345–1370.
- Leombroni, Matteo, Monika Piazzesi, Martin Schneider, and Ciaran Rogers (2020) “Inflation and the Price of Real Assets,” NBER Working Papers 26740, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/26740.html>.
- Leshno, Moshe, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken (1993) “Multilayer feed-forward networks with A Nonpolynomial Activation Function Can Approximate Any Function,” *Neural Networks*, 6 (6), 861–867.
- Lu, Zhou, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang (2017) “The Expressive Power of Neural Networks: A View from the Width,” *Advances in Neural Information Processing Systems*, 30.
- Lynch, Anthony W (1996) “Decision Frequency and Synchronization across Agents: Implications for Aggregate Consumption and Equity Return,” *Journal of Finance*, 51 (4), 1479–1497.
- Madrian, Brigitte C. and Dennis F. Shea (2001) “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior,” *The Quarterly Journal of Economics*, 116 (4), 1149–1187.
- Maliar, Lilia, Serguei Maliar, and Pablo Winant (2021) “Deep learning for solving dynamic economic models,” *Journal of Monetary Economics*, 122 (C), 76–101.
- Mankiw, N. Gregory and Stephen P. Zeldes (1991) “The consumption of stockholders and nonstockholders,” *Journal of Financial Economics*, 29 (1), 97–112.
- Mehra, Rajnish and Edward C. Prescott (1985) “The equity premium: A puzzle,” *Journal of Monetary Economics*, 15 (2), 145–161.

- Mitchell, Olivia S., Gary R. Mottola, Stephen P. Utkus, and Takeshi Yamaguchi (2009) "Default, Framing and Spillover Effects: The Case of Lifecycle Funds in 401(k) Plans," NBER Working Papers 15108, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/15108.html>.
- Mitchell, Olivia S. and Stephen P. Utkus (2022) "Target-date funds and portfolio choice in 401(k) plans," *Journal of Pension Economics and Finance*, 21 (4), 519–536.
- Nardi, Mariacristina De, Eric French, and John B. Jones (2010) "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118 (1), 39–75.
- Parker, Jonathan A., Antoinette Schoar, Allison T. Cole, and Duncan Simester (2022) "Household Portfolios and Retirement Saving over the Life Cycle," NBER Working Papers 29881, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/29881.html>.
- Parker, Jonathan A., Antoinette Schoar, and Yang Sun (2020) "Retail Financial Innovation and Stock Market Dynamics: The Case of Target Date Funds," nber working papers, National Bureau of Economic Research, Inc.
- Pinkus, Allan (1999) "Approximation Theory of the MLP Model in Neural Networks," *Acta Numerica*, 8, 143–195.
- Poterba, James M. and Andrew A. Samwick (1995) "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption," *Brookings Papers on Economic Activity*, 26 (2), 295–372.
- Saito, Makoto (1995) "Limited Market Participation and Asset Pricing," *manuscript, University of British Columbia, Department of Economics*.
- Samuelson, William and Richard Zeckhauser (1988) "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1 (1), 7–59.
- Shapiro, Matthew D (1986) "The dynamic demand for capital and labor," *The Quarterly Journal of Economics*, 101 (3), 513–542.
- Storesletten, Kjetil, Chris Telmer, and Amir Yaron (2007) "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," *Review of Economic Dynamics*, 10 (4), 519–548.
- Viceira, Luis M. (2001) "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income," *Journal of Finance*, 56 (2), 433–470.
- Vissing-Jorgensen, Annette (1998) *Limited Stock Market Participation* Ph.D. dissertation, Massachusetts Institute of Technology.
- (2002a) "Towards an Explanation of Household Portfolio Choice Heterogeneity: Non-financial Income and Participation Cost Structures," NBER Working Papers 8884, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/8884.html>.
- (2002b) "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 110 (4), 825–853.

Wolff, Edward N and Maury Gittleman (2014) "Inheritances and The Distribution of Wealth or Whatever Happened to The Great Inheritance Boom?" *The Journal of economic inequality*, 12 (4), 439–468.

Appendices

A Household Problem

$$\begin{aligned}
 V(X_t) &= \sup_{c, \bar{F}} E_t \left[\int_t^{t+a^{exit}-a_t} e^{-\rho(u-t) - \int_t^u \eta(a_s) ds} \left(u(c_u) + \eta(a_u) u^B(q_u) \right) du \right] \\
 c_t + s_t &= \begin{cases} w_t l(a_t, y_t) & a_t \leq a^{retire} \\ \bar{s} & a_t > a^{retire} \end{cases} \\
 f_t &= \begin{cases} 0 & t \leq T_1^f \\ F_t & t > T_1^f \end{cases} \\
 \tilde{f}_t &= \begin{cases} f_t & s_t \geq 0 \\ \frac{e_t}{e_t + b_t} & s_t < 0 \end{cases} \\
 de_t &= (\mu_t^e e_t + \tilde{f}_t s_t) dt + \sigma_t^e e_t dW_t \\
 db_t &= [r_t^f b_t + (1 - \tilde{f}_t) s_t] dt \\
 c_t, F_t, e_t, b_t &\geq 0,
 \end{aligned}$$

where T_1^f is the arrival time for the first jump in a Poisson counting process N_t^f with intensity

$$\lambda^f(y_{t-}) \begin{cases} +\infty & y_{t-} = star \\ \bar{\lambda}^f & y_{t-} = high \\ 0 & otherwise. \end{cases}$$

B Firm Problem

Each individual firm solves

$$\max_{D_t, I_t, \omega_t} E_t \left[\int_t^{t+s} e^{-\bar{\rho}s} \log(D_s) ds \right]$$

s.t.

$$dN_t = \left(\left[\frac{I_t}{K_t} - \Phi\left(\frac{I_t}{K_t}\right) - \delta(Z_t) \right] K_t - r_t^f B_t^f \right) dt + \sigma K_t dW_t \quad (\text{B.1})$$

$$D_t = MPK_t K_t - I_t \quad (\text{B.2})$$

$$\omega_t = \frac{K_t}{N_t} \quad (\text{B.3})$$

$$K_t = N_t + B_t^f \quad (\text{B.4})$$

$$\Phi\left(\frac{I_t}{K_t}\right) = \frac{1}{2} \phi\left(\frac{I_t}{K_t} - \delta(Z_t)\right)^2. \quad (\text{B.5})$$

Rewrite equation (B.1) with (B.2)-(B.5)

$$\begin{aligned} dN_t = & \left((MPK_t - \delta(Z_t)) \left[1 - \frac{1}{2} (MPK_t - \delta(Z_t)) \right] \omega_t N_t - r_t^f \omega_t N_t \right. \\ & \left. + [\phi(MPK_t - \delta(Z_t)) - 1] D_t - \frac{1}{2} \phi \frac{D_t^2}{\omega_t N_t} + r_t^f N_t \right) dt + \sigma \omega_t N_t dW_t, \end{aligned}$$

or, for simplicity,

$$dN_t = \mu_t^e dt + \sigma_t^e dW_t.$$

Let X_t^{agg} be the collection of aggregate state variables (distribution replaced by generalized moments) except jump Z_t , and

$$dX_t^{agg} = \mu_t^{agg} dt + \sigma_t^{agg} dW_t$$

Conjecture the firm value function as

$$V(X_t^{agg}, Z_t, N_t) = \chi_0(X_t^{agg}, Z_t) + \chi_1 \log(N_t) \quad (\text{B.6})$$

The firm HJB is

$$\begin{aligned} \sup_{\omega_t, D_t} \mu_t^{agg} V_X + \mu_t^e V_N + \frac{1}{2} \text{trace} \left[\begin{bmatrix} \sigma_t^{agg} \\ \sigma_t^e \end{bmatrix} \begin{bmatrix} \sigma_t^{agg} & \sigma_t^e \end{bmatrix} \text{Hess}_{X,N} V \right] \\ - \bar{\rho} V + \log(D_t) + \lambda_Z \left[V(X_t^{agg}, Z_t + \Delta Z, N_t) - V \right] \end{aligned}$$

Notice that conjecture (B.6) implies that the last row and the last column of $\text{Hess}_{X,N} V$ are populated by 0's except the bottom right corner element V_{NN} . Therefore, HJB can be re-

written as

$$\begin{aligned} \sup_{\omega_t, D_t} \mu_t^{agg} V_X + \mu_t^e V_N + \frac{1}{2} \text{trace} \left[\sigma_t^{agg} (\sigma_t^{agg})^\top \text{Hess}_X V \right] + \frac{1}{2} (\sigma_t^e)^2 \text{Hess}_N V \\ - \bar{\rho} V + \log(D_t) + \lambda_Z \left[V(X_t^{agg}, Z_t + \Delta Z, N_t) - V \right] = 0 \end{aligned} \quad (\text{B.7})$$

The implied first order conditions with respect to ω_t is

$$(MPK_t - \delta(Z_t)) \left[1 - \frac{1}{2} \phi(MPK_t - \delta(Z_t)) \right] - r_t^f + \frac{1}{2} \phi \frac{D_t^2}{\omega_t^2 N_t^2} - \sigma^2 \omega_t = 0$$

which means that payout is proportional to net worth:

$$D_t = \omega_t N_t \underbrace{\sqrt{\frac{2}{\phi} \left\{ r_t^f + \sigma^2 \omega_t - (MPK_t - \delta(Z_t)) \left[1 - \frac{1}{2} \phi(MPK_t - \delta(Z_t)) \right] \right\}}}_x \quad (\text{B.8})$$

The first order condition with respect to D_t is

$$\chi_1 [\phi(MPK_t - \delta(Z_t)) - 1] x - \chi_1 \phi x^2 + \frac{1}{\omega_t} = 0 \quad (\text{B.9})$$

The roots to (B.9) are

$$\frac{[\phi(MPK_t - \delta(Z_t)) - 1] \pm \sqrt{[\phi(MPK_t - \delta(Z_t)) - 1]^2 + \frac{4\phi}{\chi_1 \omega_t}}}{2\phi}$$

Given that $x > 0$,

$$x = \frac{[\phi(MPK_t - \delta(Z_t)) - 1] + \sqrt{[\phi(MPK_t - \delta(Z_t)) - 1]^2 + \frac{4\phi}{\chi_1 \omega_t}}}{2\phi}. \quad (\text{B.10})$$

Equating x implied by equations (B.8) and (B.10) yields

$$\begin{aligned} r_t^f = \frac{1}{8\phi} \left\{ [\phi(MPK_t - \delta(Z_t)) - 1] + \sqrt{[\phi(MPK_t - \delta(Z_t)) - 1]^2 + \frac{4\phi}{\chi_1 \omega_t}} \right\}^2 \\ + (MPK_t - \delta(Z_t)) \left[1 - \frac{1}{2} \phi(MPK_t - \delta(Z_t)) \right] - \sigma^2 \omega_t. \end{aligned} \quad (\text{B.11})$$

Equation (B.11) also implies that ω_t does not depend on N_t . This, combined with D_t being proportional to N_t , mean that

$$\chi_1 = \frac{1}{\bar{\rho}}$$

in order for HJB condition (B.7) to be true for all N_t . Therefore,

$$r_t^f = \frac{1}{8\phi} \left\{ [\phi(MPK_t - \delta(Z_t)) - 1] + \sqrt{[\phi(MPK_t - \delta(Z_t)) - 1]^2 + \frac{4\phi\bar{\rho}}{\omega_t}} \right\}^2 + (MPK_t - \delta(Z_t)) \left[1 - \frac{1}{2}\phi(MPK_t - \delta(Z_t)) \right] - \sigma^2\omega_t \quad (\text{B.12})$$

Substituting (B.12) into (B.8) gives

$$D_t = \frac{1}{2\phi} \omega_t N_t \left\{ [\phi(MPK_t - \delta(Z_t)) - 1] + \sqrt{[\phi(MPK_t - \delta(Z_t)) - 1]^2 + \frac{4\phi\bar{\rho}}{\omega_t}} \right\} \quad (\text{B.13})$$

Since a firm is a price taker, choice variables do not show up in μ_t^{agg} and the trace term in equation (B.7). Therefore, first order conditions for choice variables do not involve aggregate state variables (myopia). Plugging in $\chi_0(X_t^{agg}, Z_t)$ into HJB (B.7) along with optimal choices and condition (B.13) yields a system of two PDEs that do not involve N_t . Solving the system of PDEs pins down $\chi_0(X_t^{agg}, Z_t)$ but does not change the optimal decisions. One can verify sufficiency, following the same steps as in a Merton's problem with logarithmic utility.

C Social Security Replacement Rate

D Relative Consumption Error

E Basis Function

This section investigates how the trained basis function maps asset holdings into moments of the distribution. The computer learns to distinguish between heterogeneous agents from how they affect the aggregate dynamics. Figure E.16 shows the basis function after training. In this scatter plot, each dot represents an agent, and they differ in equity position (x -axis), bond position (y -axis), and age (color). Younger agents are in darker colors, while older agents are in brighter colors. The z -axis is the value of the basis function \mathcal{G} . 0 stands for the origin of the x - y plane, $(0,0)$.

The first pattern that stands out from Figure E.16 is that the computer has separated out low-wealth agents from their high-wealth counterparts. Dots representing agents with small equity and bond holdings cluster around a triangular plane in the top-left corner. From the perspective of the computer, these low-wealth agents do not affect the aggregate dynamics in the same way the wealthy agents do.

A second pattern is that the computer distinguishes households with different ages, even if they have identical equity and bond holdings. Dots of different colors stand apart in the z -direction for a given (x,y) coordinate. If the mean equity or mean bond position was sufficient for the model, the graph for the basis function would have shown lines in either the (x,z) plane or the (y,z) plane. If using both means was enough, the graph should have displayed a surface. The fact that we do not observe lines or a surface indicates that using first moments of equity and bond holdings does not suffice. Indeed, Figure F.17 and Figure 1 show that the generalized moments not only track the aggregate capital but also picks up information related to other statistics of the economy, such as the price-dividend ratio and the wealth share of workers, etc. The computer confirms that age is an important source of heterogeneity in this model.

F Generalized Moment and Aggregate Capital

Figure F.17 plots the generalized moments and the (negative) capital stock from simulated model periods. Each dot represents a quarter. The size of a dot shows the payout yield in that quarter, whereas the brightness of a dot signals the wealth share held by workers for the same quarter. The generalized moment is generally increasing in (negative) capital. Periods where workers hold more wealth (brighter dot) tend to register higher values in the generalized moment, although the exact values also depend on the price-dividend ratio.

Figure F.17 illustrates that the generalized moments are not a univariate function of capital. Despite a close relationship between the two variables, the generalized moment incorporates additional information. Figure 1 in Section 3.1 plots time series for the generalized moment and other statistics of the economy. Capital stock, price-dividend ratio, and workers' wealth share move closely with the generalized moment, although none of them completely coincide. This pattern implies the first moment of wealth is not a sufficient statistic from the computer's perspective.

G Idiosyncratic Labor Productivity

$Z = 0$

i, j	low	high	star	stationary distribution
low	0.6	0.4	0	0.104930
high	0.05	0.948625	0.001375	0.839444
star	0	0.02075	0.97925	0.055626
y	0.107914	0.719424	5.805755	0.938189

$Z = 1$

i, j	low	high	star	stationary distribution
low	0.3	0.7	0	0.031369
high	0.025	0.973625	0.001375	0.879255
star	0	0.02075	0.97925	0.089376
y	0.107914	0.719424	5.805755	1.005461

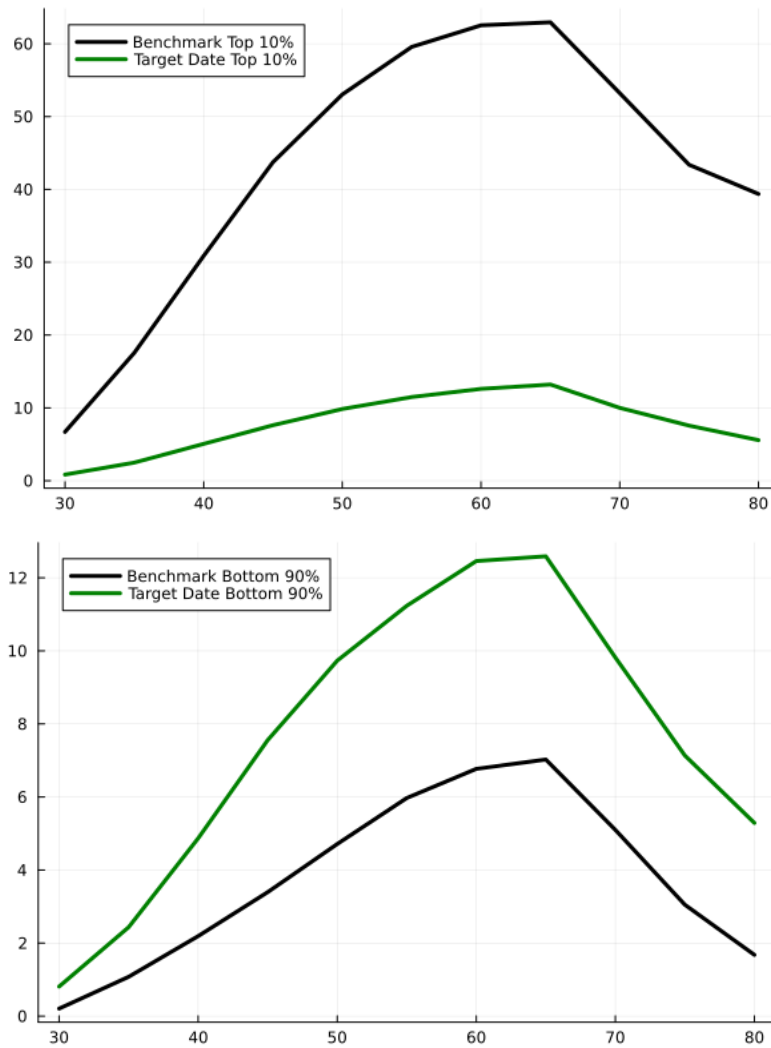
H Conditional Asset Prices

Table 5: Counterfactual: Asset Prices
Annualized Asset Return and Standard Deviation

		Equity Return		Riskfree Rate		Equity Premium	
		$E[r_t]$	$\sigma(r_t)$	$E[r_t^f]$	$\sigma(r_t^f)$	$E[r_t - r_t^f]$	$\sigma(r_t - r_t^f)$
Boom	Benchmark	0.060	0.243	0.006	0.005	0.054	0.242
	Target Date	0.026	0.152	0.006	0.006	0.020	0.152
	Free Access	0.022	0.148	0.007	0.006	0.014	0.148
Bust	Benchmark	0.055	0.251	-0.015	0.005	0.070	0.251
	Target Date	0.016	0.152	-0.015	0.006	0.031	0.152
	Free Access	0.015	0.146	-0.014	0.006	0.029	0.145

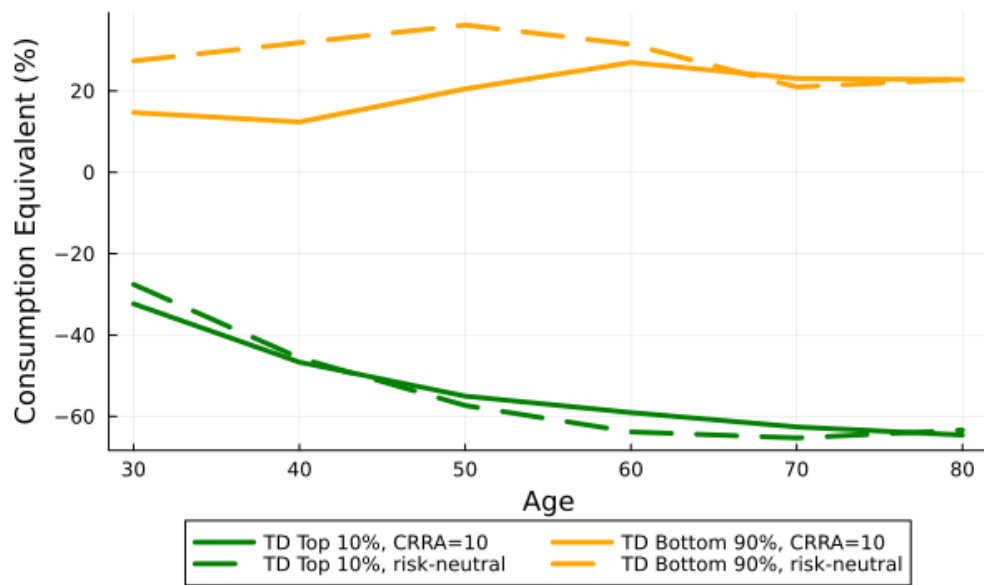
Notes. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone.

Figure 11: Counterfactuals: Consumption Equivalent - Target Date vs. Free Access



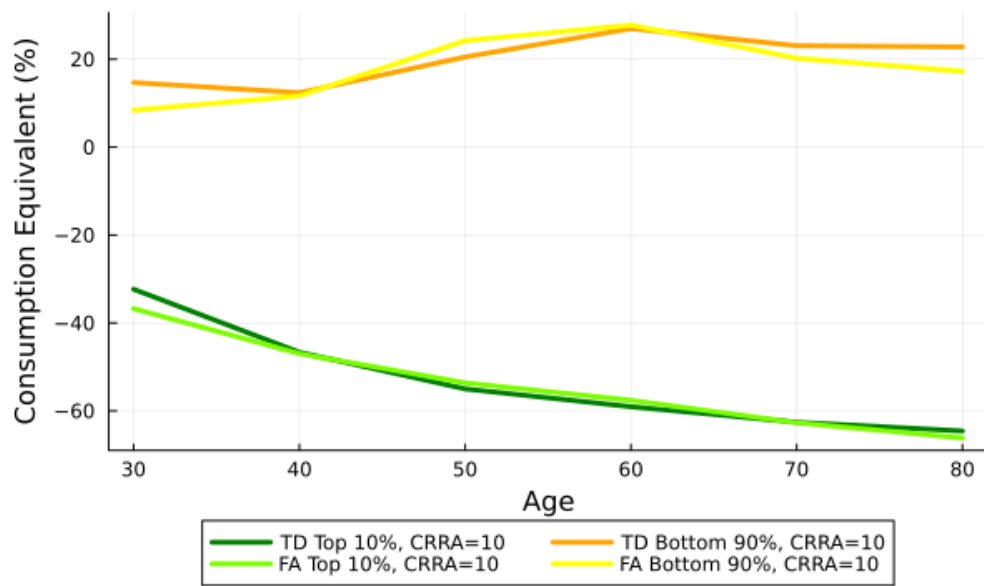
Notes. Average wealth by age and by wealth distribution is calculated from simulation cross sections in the ergodic set of the economy. Aggregate and idiosyncratic shocks are kept the same in both economies. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path.

Figure 12: Counterfactuals: Consumption Equivalent - Target Date vs. Free Access



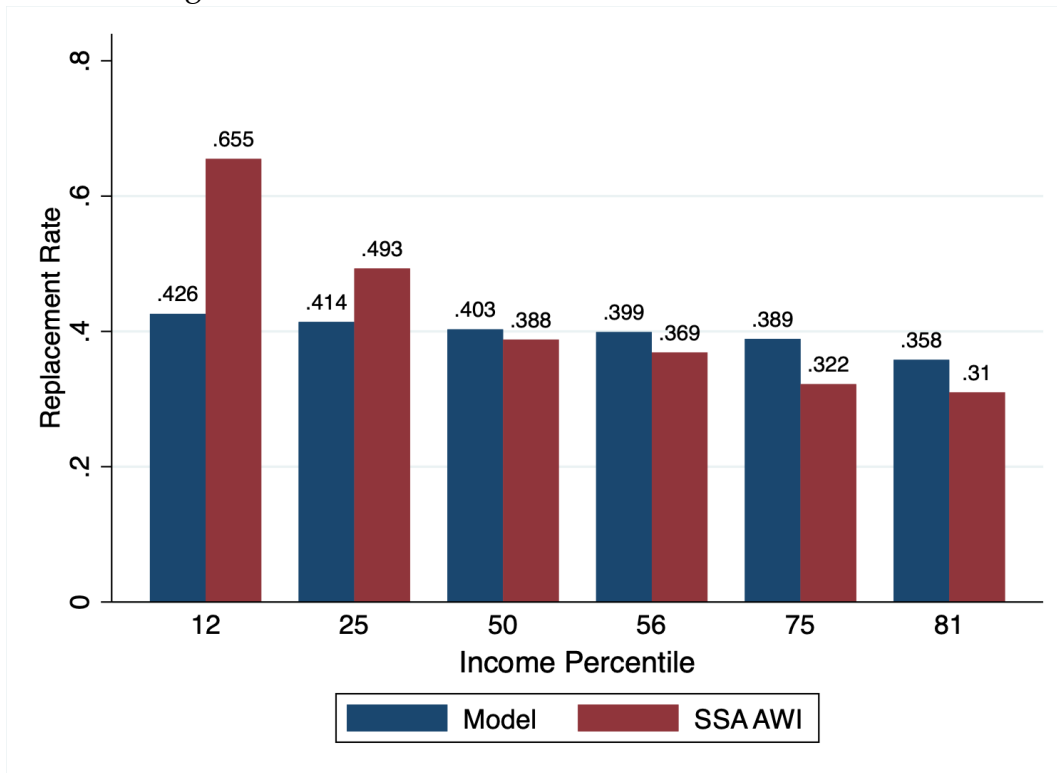
Notes. Consumption equivalents are calculated using risk neutral utility and CRRA utility with risk aversion 10. Agents are ranked by wealth according to their benchmark economy wealth holdings. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone. Consumption equivalent is defined for every age, in terms of remaining life-time consumption.

Figure 13: Counterfactuals: Consumption Equivalent - Target Date vs. Free Access



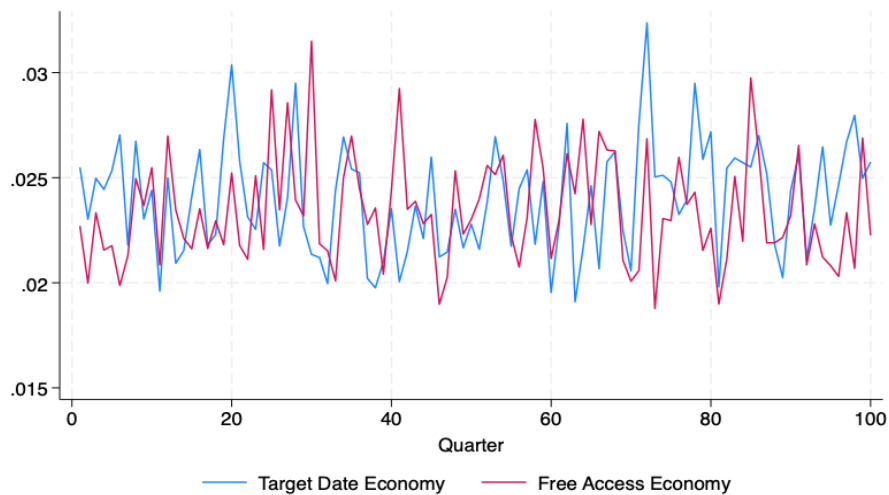
Notes. Agents are ranked by wealth according to their benchmark economy wealth holdings. Consumption equivalents are calculated through simulation of cross sections taken from the ergodic set of the economy. Benchmark economy features frictions in stock market participation and in rebalancing. Target date economy has all households following the target date glide path. Free access economy allows free participation and rebalancing for everyone. Consumption equivalent is defined for every age, in terms of remaining life-time consumption.

Figure C.14: Generalized Moment and Other Statistics



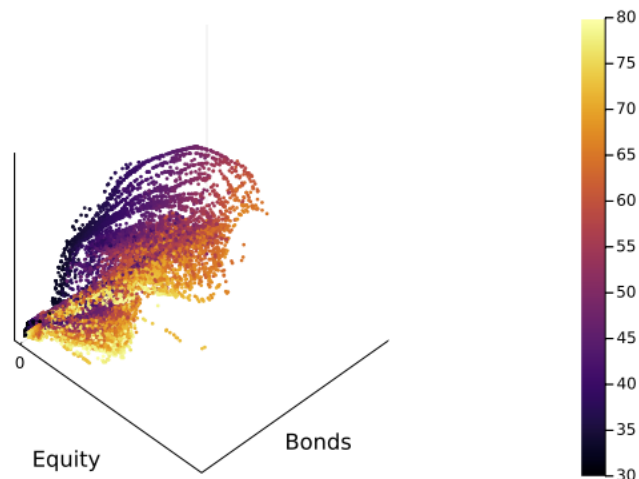
Notes. Data is from [Goss, Clingman, Wade and Glenn \(2014\)](#), calculated by the SSA Office of the Chief Actuary. Income is the high 35 years career-average earnings indexed by the national average wage index (adjusted for inflation).

Figure D.15: Average Relative Consumption Errors



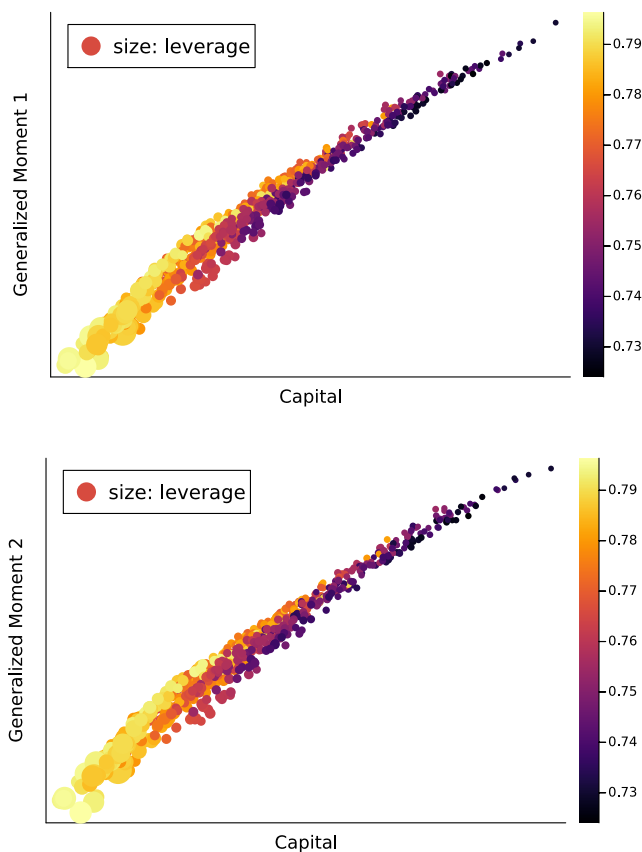
Notes. Each epoch contains 80 episodes. The upper figure checks the relative consumption errors across all individual agents who are unconstrained. Unconstrained agents are those who consume less than 96% of all cash on hand (wealth and income combined). The lower figure checks the relative aggregate consumption error.

Figure E.16: Trained Basis Function



Notes. The 3-D plot shows the trained basis function (z-value) by equity (x-value) and bond (y-value) positions. Each dot represents an agents in the economy. Young agents are in dark colors, while older agents are in brighter colors. 0 stands for the origin of the x-y plane.

Figure F.17: Generalized Moment and Other Statistics



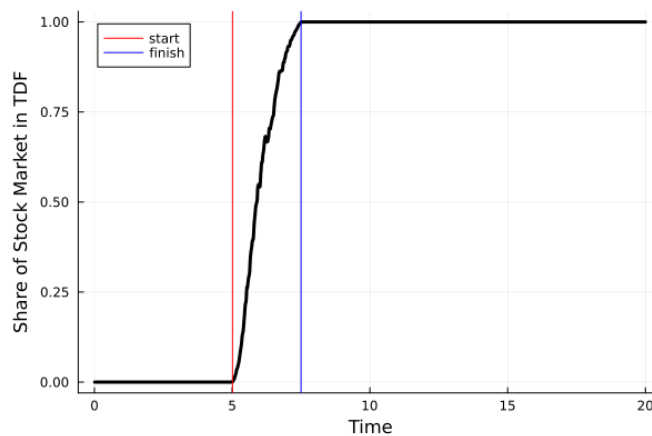
Notes. Scatter plot displays the generalized moment and the (negative) capital stock for 800 quarters. Each dot represents a quarter. The size of a dot describes the payout yield, where as the brightness of a dot shows the wealth share held by workers.

I Transition Dynamics

So far the analysis has been focused on comparing stochastic steady states under the benchmark, target date, and free access cases. In this section, I consider the transition dynamics as target date funds are adopted gradually, which reflects the potential path we are on after the passage of the 2006 Pension Protection Act and the 2022 Secure Act 2.0.⁸

As an experiment, I increase the target date users gradually in the economy and inspect the transition of asset prices and inequality. Specifically, starting from the stochastic steady state of the benchmark economy, each new cohort that enters the model is invested in target date funds by default. The full transition into the target date economy takes exactly one generation.

Figure I.18: Counterfactuals: Consumption Equivalent - Target Date vs. Free Access



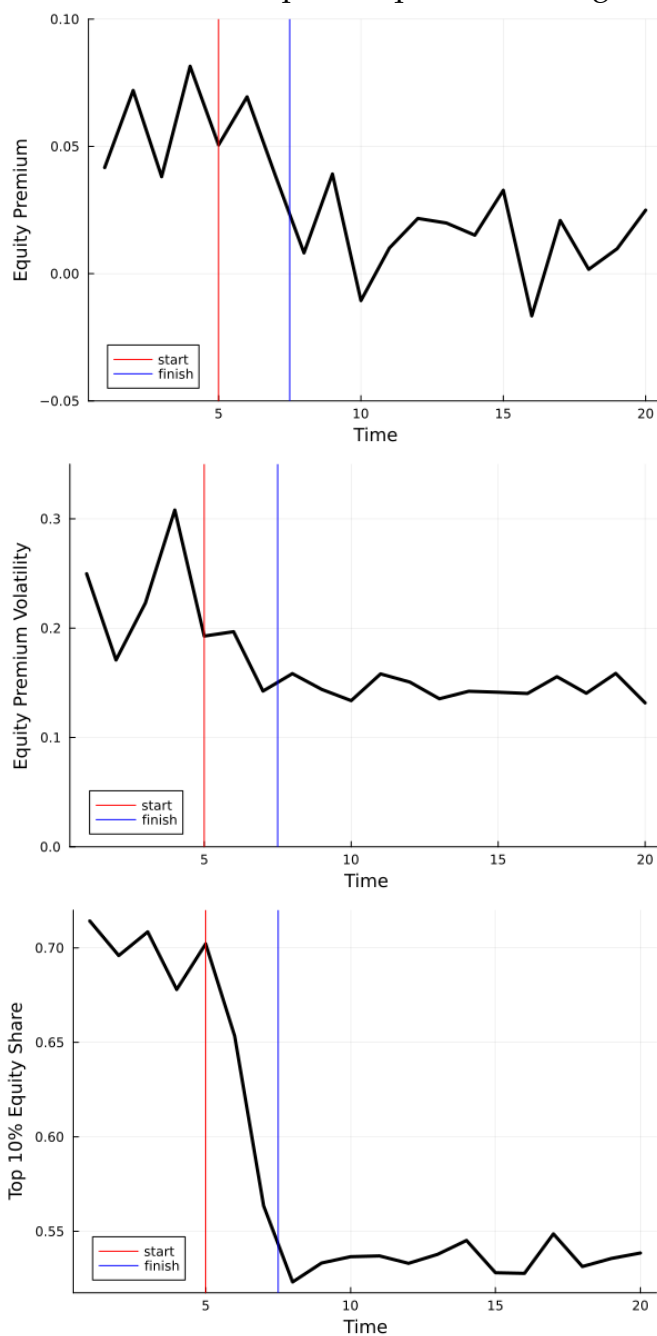
Notes. The first 5 periods (100 years) drawn from the stochastic steady state of the economy. Beginning in period 5, new cohorts enter the economy and hold target date funds by default. The transition finishes at period 7.5. From then on, the economy is in the target date stochastic steady state.

Figure I.18 depicts the stock market share in target date funds as the economy transitions from the benchmark frictional scenario to the full target date economy. Each period is two decades in the graph. The equity market share of target date funds rises dramatically for two reasons. Firstly, each entering cohort replaces the old exiting cohort, so the number of target date users increases as time passes. Secondly, as Figure 7 shows, target date cohorts enter the economy with much higher portfolio shares in equities. The consequence is that the distribution of equity holdings shift dramatically from benchmark agents to target date users, which explains the rapid transition into the target date stochastic steady state.

These new cohorts participate in risk sharing and provide stock demand elasticity, driving down the equity premium and the stock return volatility. Asset prices and inequality moves in the direction of the target date outcomes almost linearly. Figure I.19 demonstrates the transition by plotting outcomes of interests for every twenty years. The equity premium slides smoothly from 6% in the benchmark economy to 2% in the target date stochastic steady state, while the equity premium volatility drops from 25% to around 15%. At the same time, the top 10% equity market share plummets from above 70% to below 55%.

⁸The 2006 Pension Protection Act designated target date funds as a qualified default investment alternative (QDIA) to money market funds. The 2022 Secure Act 2.0 mandated automatic enrollment into retirement accounts for any new retirement plans in the U.S.

Figure I.19: Counterfactuals: Consumption Equivalent - Target Date vs. Free Access



Notes. The first 5 periods (100 years) drawn from the stochastic steady state of the economy. Beginning in period 5, new cohorts enter the economy and hold target date funds by default. The transition finishes at period 7.5. From then on, the economy is in the target date stochastic steady state.