Belief Heterogeneity and the Income-Wealth Relationship*

Nigel McClung
Bank of Finland
Research Economist

Tristan Nighswander
Reed College
Assistant Professor

This Version: November 20, 2021

Abstract

Overlapping generations models with uninsurable idiosyncratic income risk struggle to match the right skew observed in the U.S. wealth distribution, and the weak correlation between income and wealth inequality we document in a cross-section of 28 countries. We argue that a realistic deviation from full information rational expectations may help standard models match these features of the data. A simple model in which agents have heterogeneous beliefs about personal employment probabilities, and learn from personal employment experience, can generate significant right skew in the wealth distribution, and a range of different wealth distributions for a given income distribution.

Keywords: Learning, Life-cycle Model, Heterogeneous Information, Wealth Inequality.

JEL codes: E10, E21, E71.

*The views expressed in this paper do not reflect the views of the Bank of Finland. We thank Erin Cottle-Hunt and Bruce McGough for helpful comments and discussion.

1 Introduction

It has been established empirically that the wealth distribution displays significantly more right skew than the distribution of income in the United States and around the world.[1] Further, it has been shown that there is a great deal of heterogeneity in the relationship between income and wealth inequality across countries.[2] Both features of the data are difficult to replicate using standard life-cycle models of the wealth distribution, which postulate uninsurable idiosyncratic income risk as a root cause of wealth heterogeneity. Specifically, standard life-cycle models fail to match key features of the US wealth distribution including the thick right tail (Huggett (1996)), and they further imply a strongly positive relationship between income and wealth inequality in contrast to the weak income-wealth relationship observed in the data (Benhabib et al. (2017)). These models typically assume full information rational expectations (FIRE), despite evidence of heterogeneous, non-rational beliefs.

This paper studies the implications of heterogeneous labor market expectations for aggregate savings and the wealth distribution in an overlapping generations model. We assume that in each period a fraction of newborn agents inherit correct information regarding the persistence of idiosyncratic employment shocks while the remaining newborn agents are born with incorrect beliefs and must forecast their perceived employment probabilities using personal labor market experiences and simple statistical learning tools.[3] Unlike most adaptive learning papers, which typically study economies with infinitely lived agents or finitely lived agents who share information perfectly across generations,[4] our paper features finitely lived agents who cannot learn to forecast rationally because they only learn from personal data.

[1]See Vermulen, 2016; Piketty 2014 for a discussion of the empirical income-wealth relationship and DeNardi and Fella, 2017, for a survey of the literature aimed at generating realistic degrees of wealth inequality in heterogeneous agent life-cycle economies

[2]See Benhabib et al., 2017; Jannti et al.; 2008

[3]Our approach to learning is consistent with the adaptive learning approach popularized by Evans and Honkapohja (2001), Marcet and Sargent (1989), among others.

[4]E.g. Branch et al., (2013), Hunt (2019) study environments where agents forecast aggregate variables using economy-wide adaptive learning rules. In these settings, agents can collectively learn to coordinate on a rational expectations equilibrium because their beliefs depend on the history of the economy and not on idiosyncratic information.

As a result, heterogeneous beliefs survive in the economy and initial beliefs have persisting effects on the overall capital stock and the distribution of resources.

This paper makes two contributions to the conversation on modeling the income-wealth relationship. First, we provide an updated estimate of the relationship between income and wealth inequality using a much larger cross-section of countries than previous work, such as Benhabib et al. (2017). Our analysis confirms earlier findings that the earnings Gini coefficient is **not** a statistically significant predictor of the wealth Gini coefficient. Second, we embed heterogeneous beliefs about about personal employment probabilities into an otherwise standard life-cycle model economy, and illustrate that simple deviations from full information rational expectations (FIRE) generate a far greater degree of wealth skew than traditional FIRE models, and can generate a wide range of wealth distributions for a given income distribution. Thus, belief heterogeneity may help rationalize right skew in the wealth distribution and the poor predictive power of income inequality for wealth inequality.

Section 2 discusses evidence of heterogeneous employment beliefs. Section 3 develops the model of belief heterogeneity we utilize throughout our analysis and outlines how beliefs evolve over the life-cycle in our model economy. Section 4 provides an updated account of the income-wealth relationship for 28 countries. Section 5 studies model-implied wealth distributions for a given income distribution.

2 Evidence of Heterogeneous Employment Beliefs

Our view that labor market expectations are diverse and impacted by personal experience is supported by a growing literature. A 2019 College Pulse survey asked 7,000 students in the U.S. how much money they expect to make after graduation, and found their median salary expectation exceeded the actual median salary for graduates with 0-5 years experience by \$12,000.[5] Over-optimism is not unique to this survey: Jerrim (2015) shows that college-aged

[5] See Mike Brown's LendEDU report "Expectations vs. Reality: Early Career Salaries."

individuals overestimate life-time income by 40%; Alesina et al. (2018) provide evidence that Americans predict unreasonably high upward mobility; Mueller et al. (2018) document over-optimism among unemployed workers with respect to their re-employment beliefs. There is also compelling evidence of over-pessimism: Tortorice (2012) shows that Michigan Survey of Consumers respondents underestimate their re-employment probability after recessions. Rozyspal and Schlafmann (2019) note that individuals are more likely to be pessimistic than optimistic when forecasting their personal income trajectory.

There is also evidence that people's beliefs about employment prospects are always diffuse and highly dependent on individual employment experiences. Guvenen (2007) finds an individual's uncertainty about their personal income growth is slowly resolved over the life-cycle, in part because idiosyncratic income shocks are infrequent and not very persistent. Ellison and Macauley (2019) use Survey of Consumer Expectations (SCE) data to outline a great deal of dispersion in household expectations regarding the probability of re-employment following a theoretical job loss. This high degree of belief dispersion persists even after controlling for individual income, age, education, race, and a host of other demographic variables.

Other researchers have explored the implications of imperfect knowledge with respect to one's earnings in life-cycle model economies, however our paper marks the first attempt to tie earnings uncertainty to wealth inequality. The majority of research on this topic has focused on a specific type of earnings uncertainty associated with one particular earnings process which Guvenen (2007) refers to as heterogeneous income profiles (HIP). Under HIP, earnings are a function of both idiosyncratic noise and common age-specific factors[6] as well as ex-ante heterogeneity in both the slope and intercept of one's own personal income trajectory. Guvenen (2007) and Guvenen and Smith (2014) show how to match key features of consumption data by embedding HIP into life cycle models with Bayesian learning agents who seek to resolve uncertainty about their unobserved income processes. In their work, the

[6]This portion of individual earnings is closely related to the calibration utilized in Huggett (1996) and the vast majority of subsequent research focused on rationalizing wealth inequality in life-cycle model economies.

HIP process is necessary to slow down Bayesian learning and generate meaningful effects of imperfect information over the life cycle. Chang et al. (2018) utilize a HIP earning process to rationalize the high ownership of risky assets in model economies relative to US data. They find that the age-specific income uncertainty associated with learning a HIP income process leads to a much better fit of risky asset ownership than a standard model in which the income process is observed by agents. However, their mechanism leverages high uncertainty in earnings when young and thus leads to much higher savings rates amongst young households than empirical research supports.

Our approach strips away all differences in actual earnings potential across agents and instead focuses entirely on the role of differences in the *perception* of future earnings. As in Chang et al., Guvenen, and Guvenen and Smith young agents in our model have greater forecast errors than older agents (see Figures 1 and 2 in Section 3). However, this does not necessarily lead to higher savings rates for these households, as the diffusion of beliefs across optimism and pessimism leads to a model economy in which high precautionary savings is not the inevitable outcome of earnings uncertainty, it is a function of both one's initial beliefs and the distribution of beliefs of other agents! By avoiding a more complicated earnings process and eliminating the channel of underlying differences in earnings ability, we are able to show just how powerful the distribution of *initial* beliefs can be in re-shaping the asset distribution.

3 The Model and Belief Evolution

Here we outline a model in which agents hold heterogeneous beliefs about the evolution of personal employment. We then illustrate how an individual's beliefs evolve over the life-cycle as finite lifespans interact with agent learning rules, incorrect initial beliefs, and idiosyncratic employment experiences.

3.1 The Model

The modeling environment is a simplified version of the standard multi-period overlapping generations model first introduced in Huggett (1996). In each period, a continuum of households are born with no assets, a non-stochastic lifespan of length J, and no bequest motives. It follows that a generation of households also dies in each period after consuming all their resources at age J. There is no population growth. Labor is supplied inelastically for the first $j_R < J$ periods after which point households retire. Household, i, of age ($j \in \{1, ..., J\}$) chooses their savings allocation ($\{a^{j,i}_{t+j-1}\}_{j=1}^{J-1}$) by solving a standard intertemporal optimization problem:

$$\max_{\{a^{j,i}_{t+j-1}\}_{j=1}^{J-1}} \hat{E}_t \sum_{j=1}^{J} \beta^{j-1} u(c^{j,i}_{t+j-1}) \tag{1}$$

$$\text{s.t.} \qquad c^{j,i}_{t+j-1} + a^{j,i}_{t+j-1} \leq R_{t+j-1} a^{j-1,i}_{t+j-2} + \epsilon(s^{j,i}_{t+j-1}) h(j) w_{t+j-1} \tag{2}$$

where $\hat{E}_t$ denotes (potentially) non-rational expectations formed at t, R_{t+j-1} and w_{t+j-1} are the economy wide return on savings and labor, respectively, and $s^{j,i}_{t+j-1}$ is a two-state persistent exogenous Markov process governing the idiosyncratic employment risk faced by optimizing households. The transition out of state $s \in \{L, H\}$ s.t. $0 \leq \epsilon(L) < \epsilon(H) = 1$ is governed by the Markov transition probabilities P_L and P_H. The high employment state, $\epsilon(H)$, corresponds to full time employment and the low employment state, $\epsilon(L)$, corresponds to agents being "unemployed." We assume that agents observe their employment process and knows the value of $\epsilon(H)$ and $\epsilon(L)$.

All agents have the same CRRA utility function, given by $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$ if $\sigma \neq 1$ and $u(c) = \ln(c)$ otherwise. If we let $x = (a, \epsilon, j)$ then households' optimization problem can be written as a dynamic programming problem:

$$\begin{aligned} V(x) &= \max_{c,a'} \Big\{ u(c) + \beta \hat{E} V(x'|x) \Big\} \\ \text{s.t. } c + a' &\leq Ra + \epsilon w \end{aligned} \tag{3}$$

The policy functions for savings, $a(x)$, and consumption, $c(x)$, determine the allocations of savings that solve the household's optimization problem, which we recast as (3). From (3), it's immediate that agents' decisions depend on expectations of employment, which may be non-rational, as denoted by $\hat{E}$. In our analysis, we assume that agents form naive expectations of the real interest rate and wage (i.e $R^e_{t+j} = R_t$ and $w^e_{t+j} = w_t$) since this is consistent with rational expectations of prices in a stationary equilibrium of the model. We furthermore assume that agents understand that personal income is driven by a process of the form (2).

We depart from rational expectations by assuming that a proportion of agents lack knowledge of the true parameters P_H and P_L and instead solve their dynamic programming problems in each period by conditioning expectations on non-rational beliefs, $P^{e,j,i}_{H,t+j-1}$ and $P^{e,j,i}_{L,t+j-1}$. This implies $\hat{E}V(x'|x) = \sum_{\epsilon'} Pr^{e,j,i}_{t+j-1}(\epsilon'|\epsilon)V(a', \epsilon', j+1|x)$ depends on the household's subjective transition probabilities, $Pr^{e,j,i}_{t+j-1}(\epsilon'|\epsilon)$, which may vary across agents and over the course of each agent's lifetime. We chose to model a two-state earnings process with just two transition probabilities for agents to learn in order to display the power of belief diffusion in a simple modeling environment. To make things tractable, we assume that agents solve their optimization problem in each period under the (potentially false) belief that their current beliefs about transition probabilities will not change in future periods (i.e. Kreps' anticipated utility approach).

We introduce heterogeneous beliefs as follows. In each period, ϕ proportion of newborn agents are endowed with knowledge of P_H and P_L. These **fully informed rational expectations** (FIRE) agents form expectations using the true employment probabilities when solving (3). Of the remaining $1-\phi$ proportion of agents, λ proportion are **pessimists** (P) (i.e. born with initial employment transition probability beliefs (P_H^P,P_L^P) such that $P_H^P < P_H$ and $P_L^P > P_L$), and $1-\lambda$ proportion are **optimists** (O) (i.e. born with initial beliefs (P_H^O,P_L^O) such that $P_H^O > P_H$ and $P_L^O < P_L$).

Intuitively, pessimists underestimate the frequency and duration of the high employment state when forming expectations to solve (3), which leads pessimists to over-accumulate assets relative to other agents over the life-cycle. For analogous reasons, optimists under-accumulate wealth. Thus, non-rational employment expectations directly impact the aggregate wealth distribution. Finally, pessimists and optimists are assumed to update their beliefs about employment transition probabilities recursively using simple statistical tools and their own employment data, in the spirit of Evans and Honkapohja (2001).[7] Importantly, the learning mechanism enables agents to learn from personal employment experience but does not cause belief heterogeneity to vanish, since agents have finite lifespans and do not share information with other households within or across generations. We note that the household-side of the economy collapses to an entirely standard FIRE framework when $\phi = 1$.

It is important to note that one's initial status as an optimist or a pessimist need not persist over the course of an agent's life as individual's continually update their beliefs about idiosyncratic employment risk based on their personal employment experiences. Thus, a pessimist could become an optimist if they experience persistent employment spells. Similarly, an optimist could become a pessimist if they experience several consecutive periods of unemployment. The interaction of initial beliefs, labor market experiences, and newly formed employment expectations is discussed in detail in the following subsection. All other features of the economy are standard: output is determined by a Cobb-Douglas production

[7]Agents update beliefs about the transition probabilities using a simple recursive specification with a constant gain parameter and their own personal employment data given initial beliefs about the transition probabilities. See the Appendix for more on the learning specification and other modeling details. We follow Kreps' (1998) anticipated utility approach and assume that agents do not account for the fact that estimates are time-varying when solving (1).

function with labor and capital as inputs; factor prices are determined competitively, and labor, goods, and asset markets clear in each period. See the Appendix 1 for more details.

3.2 Evolution of Beliefs

Here we illustrate the interaction of finite lifespans, learning, incorrect initial beliefs, and savings decisions in an economy populated by learning agents. In panel (a)-(c) in Figure 1, we outline the evolution of beliefs for an optimistic agent (blue), a pessimistic agent (red), an agent with correct initial beliefs (green), and a FIRE agent (dashed black) all subject to the same employment shock history. Note that all non-FIRE agents experience evolving beliefs over the life-cycle, which suggests that non-FIRE households may over- or under-accumulate assets relative to FIRE households based on initial conditions (i.e. $P_{k,t}^{e,1,i} \neq P_k,\ k \in \{L, H\}$) or based on life experiences (i.e. $P_{k,t+j-1}^{e,j,i} \neq P_k,\ k \in \{L, H\}$).

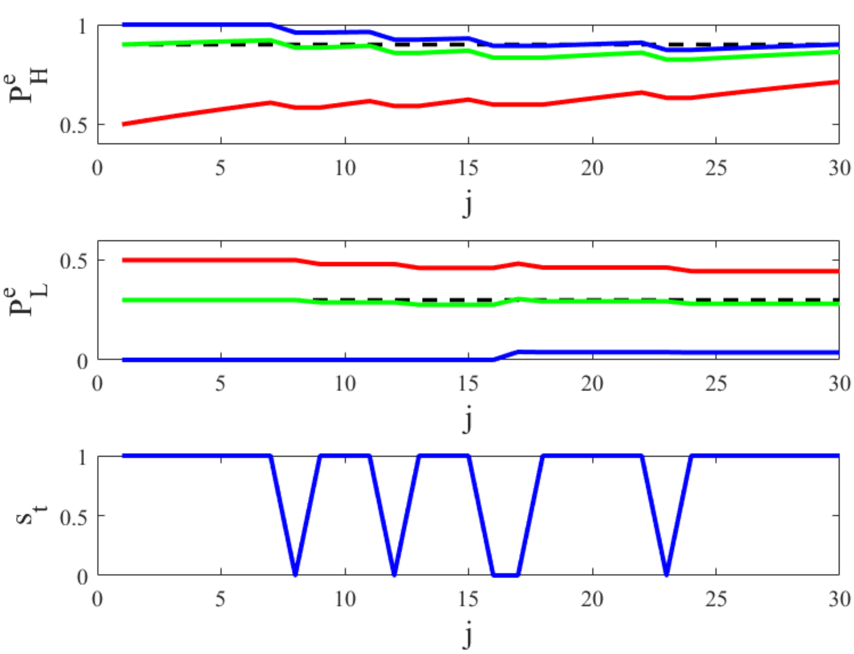

Figure 1: Belief Evolution- 30 year economic life

Panel (a) shows the evolution of beliefs regarding P_H^e, panel (b) shows the evolution of beliefs for P_L^e, and panel (c) shows the employment shock history. In panel (c), 1 corresponds to drawing the high employment state, $\epsilon(H)$, and 0 corresponds to drawing the low employment state, $\epsilon(L)$. For the shock history experienced by these agents, each type continually updates their beliefs about P_H^e in the first 8 periods, with beliefs about the persistence of the high employment state consistently increasing for agents with correct initial beliefs (green) and pessimistic (red) agents until period 8 when the low employment state is experienced for the first time. Following a one period experience of unemployment, agents revise their beliefs about the persistence of the low employment state upwards and their beliefs about the persistence of the high employment state downwards. This process continues over the life-cycle.

Figure 1 shows us that when we interact learning, finite lifespans and incorrect beliefs, we end up with a population of agents who may live and die with incorrect beliefs, and hold beliefs that are anchored to their initial beliefs. Finally, since each expiring generation is replaced by a new generation of uninformed agents, the heterogeneity in beliefs never disappears from the economy. Figure 1 also shows beliefs about P_H converging more quickly across agent types than beliefs about P_L. This occurs as agents only learn about states of the world they visit, so the less frequent low employment state provides agents fewer opportunities to learn about the dynamics into and out of said state. How large must J be for the heterogeneity-preserving effects of finite lifespans to vanish from the economy? Figure 2 repeats the exercise of Figure 1 for $J = 100$.

The greater convergence in beliefs displayed in Figure 2 is a helpful reminder of how our modeling environment is distinct from previous papers in the learning literature. Typically, the initial beliefs of agents are unimportant because agents are either infinitely lived or share information perfectly across generations (like our FIRE households). However, when a finite life with imperfect initial information is imposed, individuals may not be able to overcome their poor initial beliefs as they simply don't have enough time to learn about the true transition probabilities in their idiosyncratic earnings process.

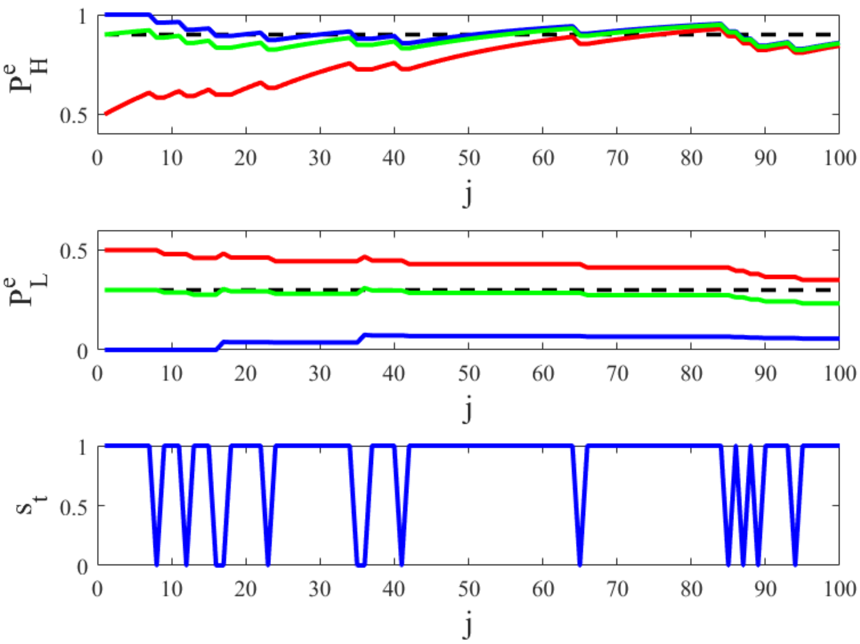

Figure 2: Belief Evolution- 100 year economic life

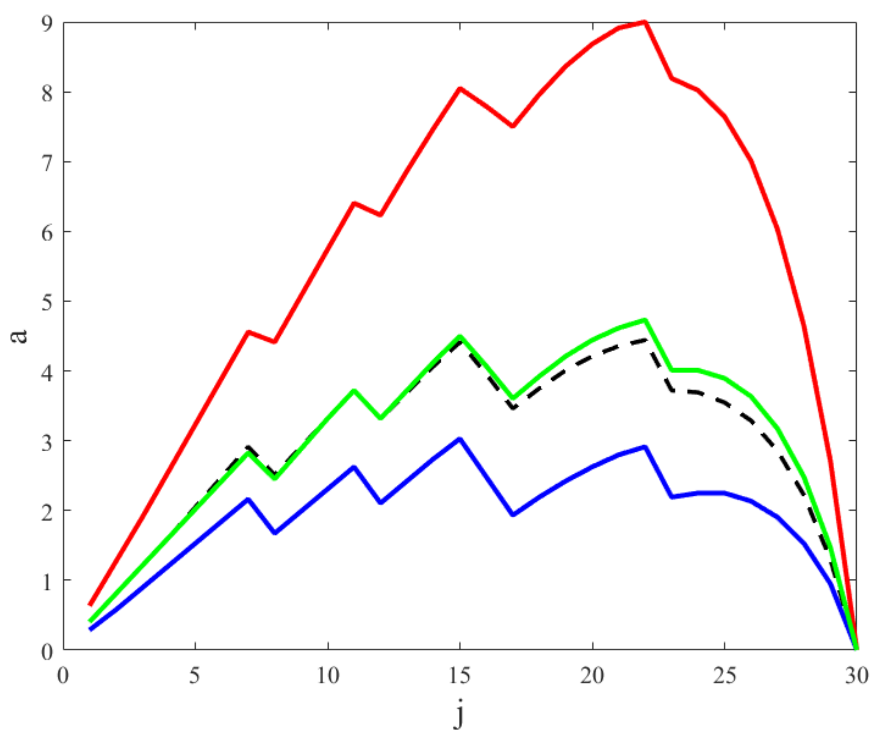

Figure 3: Asset Evolution

In our simple framework, agents' beliefs about the employment transition probabilities greatly impact their savings decisions. Figure 3 displays the savings decisions of each agent type featured in Figure 1 over their lives given fixed arbitrary values of R and w. The pessimist drastically over-accumulates assets relative to the other agents, holding more than twice the wealth of the FIRE agent in the simulation. This overaccumulation is entirely driven the pessimist's dismal expectations of lifetime income and correspondingly strong precautionary motive. The opposite is true of the optimist, particularly early in their life before the experience of unemployment causes their optimism to diminish endogenously. Agents with correct initial beliefs and FIRE agents have similar savings schedules, despite the fact that a non-FIRE realist's beliefs are relatively volatile and the FIRE agent's beliefs are fixed. Because FIRE agents and non-FIRE agents with correct initial beliefs are almost indistinguishable from the perspective of measuring wealth accumulation, we abstract from realists altogether in the remainder of the paper. The fact that pessimists over-accumulate and optimists under-accumulate relative to FIRE agents will have implications for aggregate savings (and therefore equilibrium interest rates and wages), and wealth inequality. We explore these implications in section 5.

4 The Income-Wealth Relationship

Benhabib et al. (2017) provide evidence that the earnings Gini coefficient is a poor predictor of the wealth Gini coefficient using data from 9 countries. However, they conclude that their lack of statistical relationship is merely suggestive due to the small sample size they consider. Here we provide an updated account of this relationship utilizing the 2019 Credit Suisse Global Wealth Report, which harmonizes wealth data for a larger number of countries. We combine all wealth data labeled as satisfactory or above by Credit Suisse with OECD data (2021) on post tax and transfer income inequality to generate a dataset of 28 countries[8].

Figure 4: Income and Wealth Gini

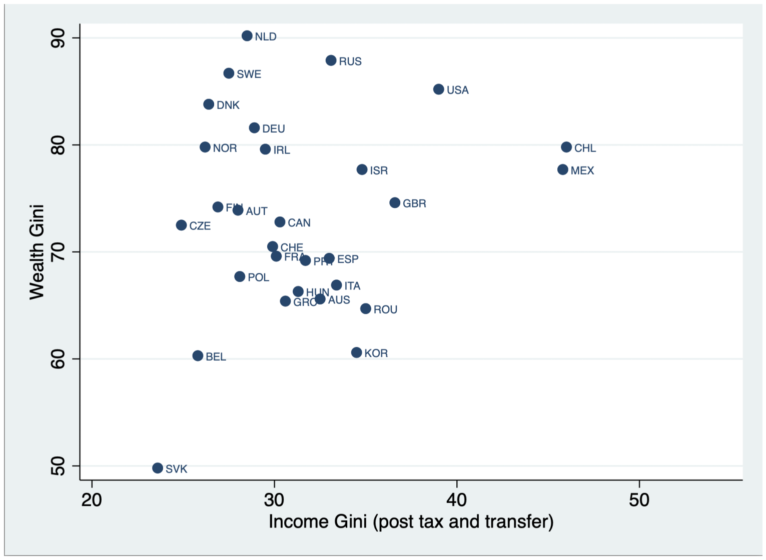

The Benhabib et al. result is preserved in our much larger sample. Just considering countries with an income Gini coefficient near 30, we see that the wealth Gini coefficient ranges between 65 (Montenegro and Hungary) and 90 (the Netherlands). This large range is difficult to explain in the context of a quantitative life-cycle model, as wealth inequality in this class of models is heavily correlated with inequality in the income process fed into the model[9] The slope coefficient associated with Figure 1 is just .34 with a standard error of .33, and the R^2 of this regression is .04. Thus, we conclude that income inequality is not an adequate predictor of wealth inequality

5 Calibration and Results

In this section, we provide an overview of the model calibration and present results from several calibration exercises in which we vary the proportion of FIRE agents (ϕ) and the

[8]As many wealth surveys (like the US Survey of Consumer Finances) are not conducted annually, we match the most recent OECD income data (2017-2019) with the wealth data reported by Credit Suisse spanning the same time frame.

[9]See De Nardi and Fella (2017) for an overview of this literature.

proportion of non-FIRE agents with pessimistic initial beliefs (λ).

5.1 Calibration

Table 1: Calibration

Parameter	Value	Interpretation
β	0.96	Discount Rate
σ	2.0	IES
α	0.33	Capital Share
δ	0.025	Depreciation rate of capital
P_L	0.3	$Pr(\epsilon' = \epsilon(L) \mid \epsilon = \epsilon(L))$
P_H	0.9	$Pr(\epsilon' = \epsilon(H) \mid \epsilon = \epsilon(H))$
P_L^O	0	Optimist Initial Belief: $Pr(\epsilon' = \epsilon(L) \mid \epsilon = \epsilon(L))$
P_H^O	1	Optimist Initial Belief: $Pr(\epsilon' = \epsilon(H) \mid \epsilon = \epsilon(H))$
P_L^P	.5	Pessimist Initial Belief: $Pr(\epsilon' = \epsilon(L) \mid \epsilon = \epsilon(L))$
P_H^P	.5	Pessimist Initial Belief: $Pr(\epsilon' = \epsilon(H) \mid \epsilon = \epsilon(H))$
$\epsilon(H)$	1	Payoff if employed
$\epsilon(L)$	0.15	Payoff if unemployed
γ_H	0.04	Gain parameter learning on high state
γ_L	0.04	Gain parameter learning on low state
J	62	Length of agent's life
j_R	45	Retirement age

Now we calibrate our simple model and study the model-implied aggregate wealth distribution. Throughout this section, we hold fixed β, σ, firm and labor market parameters (i.e. capital share, α, depreciation rate, δ, and P_L, P_H, $\epsilon(H)$, $\epsilon(L)$), learning gain parameters (γ_H and γ_L), initial beliefs (P_L^O, P_H^O , P_L^P , P_H^P), and demographic parameters (J, j_R) (see Table 1 for details). Note that by fixing these parameters, we are holding the income distribution constant throughout this section.

Our choice of $\alpha, \beta, \sigma, \gamma_H$ and γ_L are standard in the literature. The transition probabilities (P_L and P_H) are calibrated to match estimated employment transition probabilities from PSID data in Ashman and Neumuller (2019).[10] Payoffs, $\epsilon(H)$ and $\epsilon(L)$, were selected so that

[10]Ashman and Neumuller provide estimates of the semi-annual transition probabilities into and out of unemployment broken down by race, education, and family structure. Their estimates indicate that an annualized $P_H \in (.79, .995)$ and an annualized $P_L \in (.09, .50)$.

the high employment state corresponds to receiving the economy-wide wage and the low-employment state replaced 15% of annual income[11]. Our calibration of the terms governing optimism and pessimism (P_L^O, P_H^O, P_L^P, and P_H^P) were chosen so that optimists believe they will always be employed until experience causes them to update beliefs and pessimists think there is a 50-50 probability of unemployment next period regardless of their current employment state.[12] Although giving optimists 100% certainty that they will be employed next period may on its face seem like a large departure from rational expectations, their beliefs about P_H^0 are just 10% higher than the true value of P_H.

5.2 Results

Table 2 highlight our results from several model calibrations in which we vary the proportion of agents with pessimistic preferences ($\lambda \in [0, 1]$) while fixing the fraction of informed agents ($\phi = 0$). Results are compared to a benchmark economy comprised of FIRE agents ($\phi = 1$). Table 2 displays the Gini coefficient on wealth (Gini_W) and income (Gini_I) along with the market clearing interest rate (r) for each calibration.

Table 2: Wealth Statistics in Model Economies

				$\phi = 0$		
	FIRE	$\lambda = 0$	$\lambda = .10$	$\lambda = .50$	$\lambda = .90$	$\lambda = 1$
Gini_W	41.2	44.9	44.7	42.3	38.2	36.9
Gini_I	25.0	25.0	25.0	25.0	25.0	25.0
r	3.0	3.6	3.4	2.6	2.1	1.9

It is clear that when the economy is comprised of many optimists ($\lambda < .5$), the low savings propensity of optimistic households bids up the market interest rate and leads to

[11] We chose 15% as our replacement rate as unemployment insurance replaces between 30% and 50% of lost earnings (Gorman 2021) for 16-30 weeks, depending on the state (Center on Budget and Policy Priorities 2021). Increasing (decreasing) the value of $\epsilon(L)$ decreases (increases) the wealth differences between pessimistic and optimistic societies.

[12] We selected a 50-50 split for pessimists' beliefs in light of recent research by Enke and Graeber (2019) which argues that agents faced with uncertain binary environments are likely to gravitate towards 50-50 probabilities as their default. Alternative calibrations result in qualitatively similar wealth distributions.

greater inequality than a model comprised of FIRE or a majority of pessimistic agents. This increased inequality occurs as many optimistic agents hold low or 0 wealth while the high market interest rate provides pessimistic agents and optimists who have experienced negative income shocks (and thus shed their optimism) with high returns relative to a model comprised of rational agents.

Our results indicate that belief diffusion may be an important mechanism for rationalizing the weak empirical relationship between income and wealth inequality outlined in Section 4. Each model economy has an identical income process (Gini_I=25), however the wealth Gini coefficient ranges from 9.0% higher than Gini_W in the FIRE economy (44.9 when $\lambda = 0$) to 10.4% lower (36.9 when $\lambda = 1$). As noted in Section 4, for countries with an income Gini coefficient of roughly 30 the highest wealth Gini coefficient (90) is just over 38% higher than the lowest wealth Gini (65). In our simplified model that eliminates all differences between model economies outside of differences in employment beliefs, we can generate a range of wealth Gini coefficients for which the highest Gini coefficient is nearly 22% higher than the lowest. We conclude that belief heterogeneity is capable of explaining over half of the variation in observed wealth inequality for a given degree of income inequality.

Although the values of Gini_W in our model economy fall well below the value of the wealth Gini in U.S. data ($\approx$ 80), this is primarily a function of the low degree of income inequality we feed into our model. As our focus is outlining the impact of belief heterogeneity on economy-wide outcomes and not matching specific moments of the wealth distribution, we chose a simpler 2-state earnings process than what is typically fed into life-cycle models. This 2-state earnings process provides a clear way of assigning optimism and pessimism with respect to perceived employment probabilities (see Section 3) whereas a more robust earnings process would require a less intuitive learning environment. Future work will be aimed at directly calibrating the belief distribution in the model economy using Survey of Consumer Expectation (SCE) data in modeling environment with much greater skew in the imposed income distribution.

Table 3: Wealth Statistics in Model Economies

	$\phi = .25$				$\phi = .75$		
	FIRE	$\lambda = 0.1$	$\lambda = .5$	$\lambda = .9$	$\lambda = 0.1$	$\lambda = 0.5$	$\lambda = 0.9$
Gini_W	41.2	44.0	42.1	39.4	42.1	41.6	40.9
Gini_I	25.0	25.0	25.0	25.0	25.0	25.0	25
r	3.0	3.3	2.7	2.2	3.1	2.9	2.7

Table 3 highlights the same variables (Gini_W, Gini_I, and r) for our model economy, but rather than assuming that all agents are non-FIRE ($\phi = 0$) as in Table 2, we vary the proportion of FIRE agents along with the proportion of pessimists (λ). As in Table 2, we see that economies comprised of a higher proportion of optimistic agents ($\lambda < .5$) generate a higher degree of wealth inequality than those in which the majority of non-FIRE agents are pessimists. Further, as the proportion of FIRE agents increases, the spread of wealth Gini coefficient is reduced as the proportion of pessimists in the economy is varied.

Figure 5 provides evidence of increased right skew in the wealth distribution when agents hold biased beliefs. We plot the wealth distribution in the FIRE economy along with the wealth distribution when $\phi = 0$ and $\lambda = .1, .5$, and $.9$, corresponding to panels (a), (b), and (c)[13]. In each panel, the wealth distribution from the FIRE economy is plotted in red alongside the wealth distribution for the corresponding value of λ in blue. For each model economy with non-FIRE agents, the wealth distribution displays a longer right tail, indicating a subset of agents who acquire more wealth than any individual in the purely rational FIRE economy.

When $\lambda = .1$, meaning there are many optimists and a few pessimists, the tail behavior of the wealth distribution bares little resemblance to the tail behavior of the FIRE model economy. Rather than a quick tapering off of the mass of agents holding high wealth when all agents are rational, a small number of agents amass wealth well above the majority of households in the economy. This occurs as the majority of optimists drive the market interest rate up and provide enormous incentives for pessimists to save and accumulate high wealth.

[13] These graphs correspond to the wealth moments outlined in Table 2.

Figure 5: Wealth Distribution Across Modeling Environments

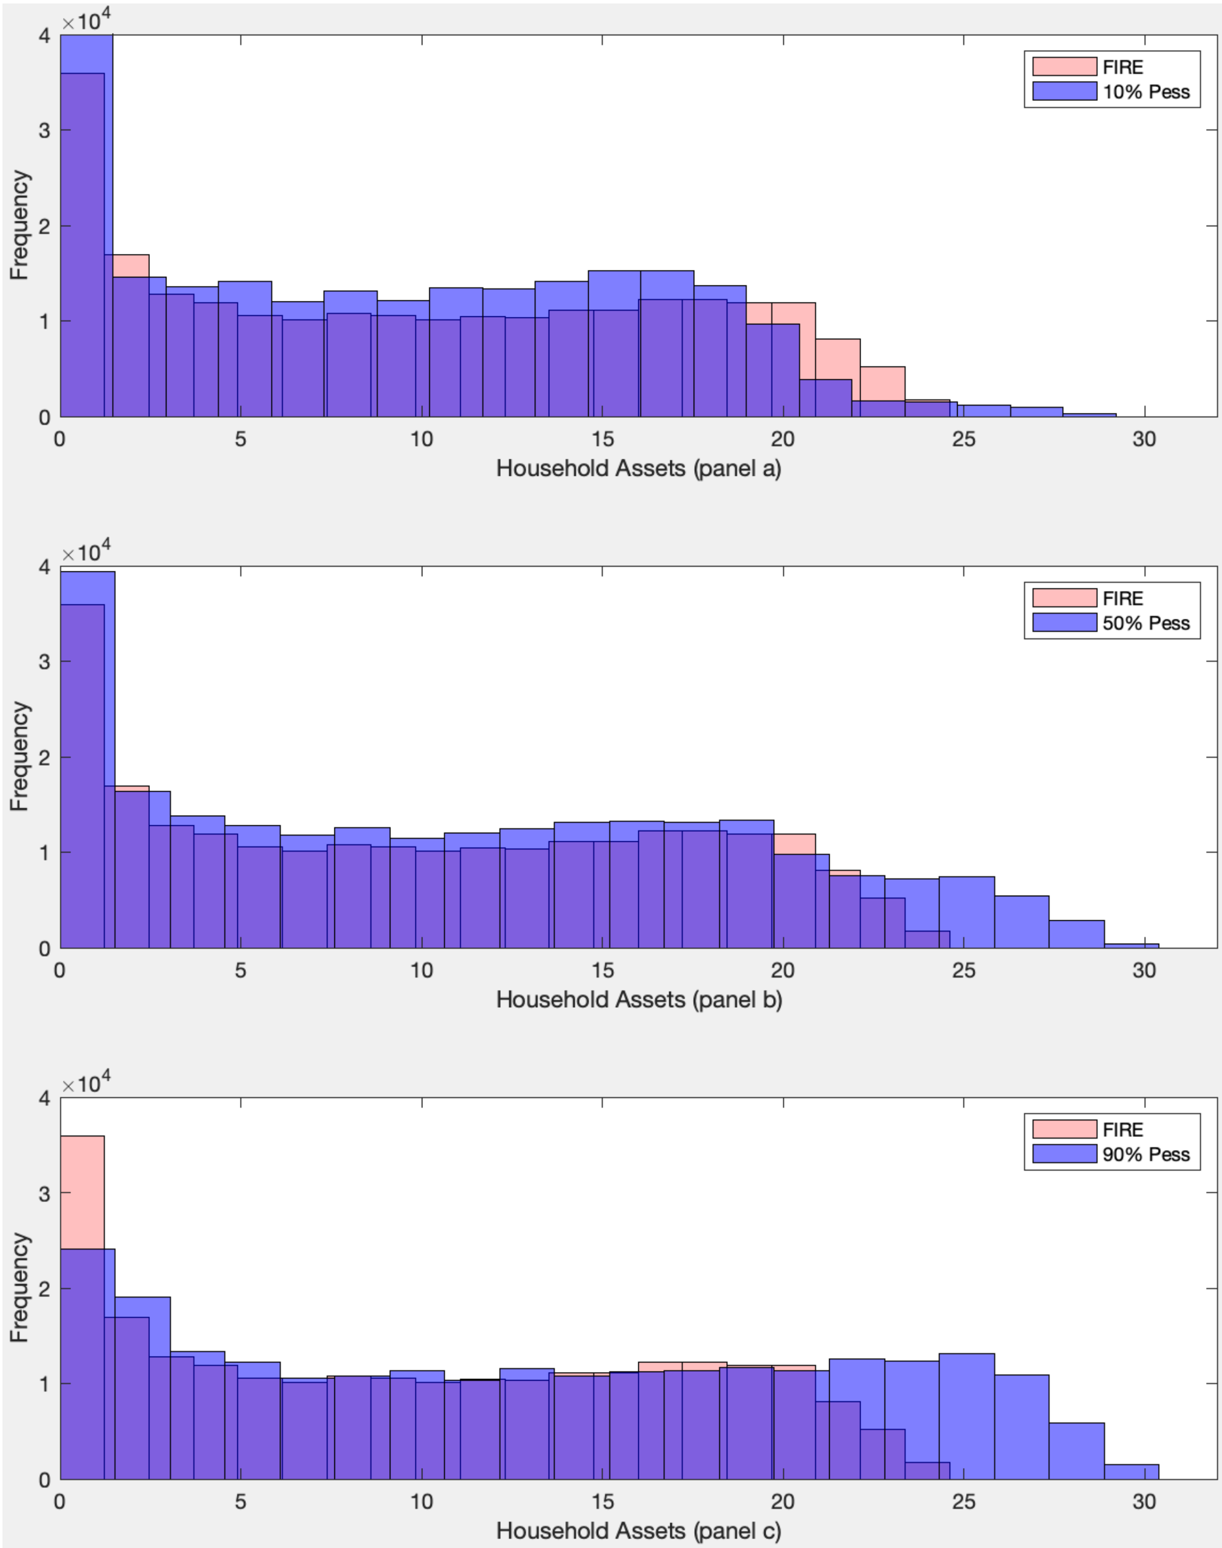

As λ increases to .5 (panel b), the tail of the wealth distribution becomes thicker as more pessimistic agents with a high savings propensity accumulate high wealth relative to the FIRE economy. However, the presence of as many pessimists as optimists adds additional weight to the right tail. Thus, even though there are more agents accumulating high wealth, the gini coefficient decreases as these high wealth agents stand out less relative to the average wealth in the economy.

When λ approaches 1 (panel c) and the vast majority of agents are pessimistic, meaning the return on savings is low, the wealth distribution looks nearly uniform. Further, although the wealth distribution is shifted right relative to the FIRE economy, the tail behavior is fairly indistinguishable from an environment in which all agents are rational. We conclude that belief heterogeneity may be a strong predictor of the tail behavior of the wealth distribution. As there is considerable interest surrounding matching the top-end of the wealth distribution in calibrated life-cycle models, we believe this mechanism merits further exploration.

6 Conclusion and Future Work

We show that a modest, realistic deviation from full information rational expectations can dramatically reshape the economy-wide wealth distribution. Our mechanism is capable of generating a more realistic degree of skewness in the aggregate wealth distribution than a standard model with FIRE. Further, heterogeneous beliefs about personal employment risk may help rationalize our finding that the income Gini coefficient is a poor predictor of the wealth Gini coefficient in 28 countries. In our model, the same income process with a Gini coefficient of 25.0 is capable of generating a wealth distribution with Gini coefficients between 36.9 and 45.0.

Many avenues for future research remain. First, these beliefs could be calibrated utilizing the Survey of Consumer Expectations and embedded in a quantitative life-cycle model in order to study whether belief heterogeneity accounts for the high degree of inequality observed in US wealth data. Second, we hope to extend our model of idiosyncratic learning to a model of learning within networks. In such an environment, agents will utilize personal information as well as information from network members to formulate forecasts.

References

Alesina, A., Stantcheva, S., and E. Teso (2018). Intergenerational Mobility and Preferences for Redistribution. *American Economic Review, 108(2)*, 521-554.

Ashman, H. and S. Neumuller (2020). Can Income Differences Explain the Racial Wealth Gap? A Quantitative Analysis. *Review of Economic Dynamics, 35*, 220-239.

Benhabib, J., Bisin, A., and M. Luo (2017). Earnings Inequality and Other Determinants of Wealth Inequality. *American Economic Review Papers and Proceedings, 107(5),* 593-597.

Branch, W.A., Evans, G.W., and B. McGough (2013). Finite Horizon Learning. In Sargent, T. and Vilmunen, J., editors, *Macroeconomic at the Service of Public Policy,* chapter 8, 141-163. Oxford University Press, Oxford.

Brown, M. (2019). Expectations vs. Reality: Early career Salaries. *LendEDU report, 02/12/2019*.

Center on Budget Policy and Priority. How Many Weeks of Unemployment Compensation are Available? *Policy Basics, 09/6/2021.*

Chang, Y., Hong, J., and M. Karabarbounis (2018). Labor Market Uncertainty and Portfolio Choice Puzzles. *AEJ: Macroeconomics, 10(2),* 222-262.

Credit Suisse Global Wealth Report (2019). Wealth Inequality (Gini Coefficient). (Accessed on 10 March 2021)

De Nardi, M., and G. Fella (2017). Savings and Wealth Inequality. *Review of Economic Dynamics, 26,* 280-300.

Ellison, M. and A. Macaulay (2019). A Rational Inattention Unemployment Trap. *CEPR Discussion Paper No. 13761.*

Enke, B. and T. Graeber (2019). Cognitive Uncertainty. *NBER Working Paper No. 26518.*

Evans, G.W. and S. Honkapohja (2001). Learning and Expectations in Macroeconomics. Princeton University Press, Princeton.

Gorman, L. (2020). Unemployment Benefit Replacement during the Pandemic. *NBER The Digest, 07/2020.*

Guvenen, F. (2007). Learning Your Earnings: Are Labor Income Shocks Really Very Persistent? *American Economic Review, 97(3),* 687-712.

Guvenen, F. and A. Smith. (2014). Inferring Labor Income Risk and Partial Insurance. *Econometrica, 82(6),* 2085-2129.

Huggett, M. (1996). Wealth Distribution in Life-Cycle Economies. *Journal of Monetary Economics, 38(3)*, 469-494.

Hunt, E. (2019). Life-cycle Horizon Learning, Social Security Reform, and Policy Uncertainty. *Working Paper.*

Jantti, M., E. Sierminska and T. Smeeding (2008). The Joint Distribution of Household Income and Wealth: Evidence from the Luxembourg Wealth Study. *OECDSocial, Employment and Migration Working Papers, no. 65.*

Jerrim, J. (2015). Do College Students Make Better Predictions of their Future Income than Young Adults in the Labor Force? *Educational Economics, 23(2)*, 162-179.

Kreps, D. (1998). Anticipated Utility and Dynamic Choice. *1997 Schwartz Lecture,* in Frontiers of Research in Economic Theory *edited by D.P. Jacobs, E. Kalai, and M. Kamien,* Cambridge University Press, Cambridge England.

Marcet, A. and Sargent, T. (1989). Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory, 48(2)*, 337-368.

Mueller, A., Spinnewijn, J., and G. Topa (2018). Job Seekers' Perceptions of Employment Prospects: Heterogeneity, Duration Dependence and Bias. *NBER Working Paper 25294.*

OECD (2021), Income inequality (Gini Coefficient). doi: 10.1787/459aa7f1-en (Accessed on 06 May 2021)

Piketty, T. (2014). Capital in the Twenty-First Century. *Cambridge Massachusetts: The Belknap Press of Harvard University.*

Rozyspal, F. and K. Schlafmann (2019). Overpersistence Bias in Individual Income Expecataions and its Aggregate Implications. *CEPR Discussion Paper 12028.*

Tortorice, D. (2012). Unemployment Expectations and the Business Cycle. *The B.E. Journal of Macroeconomics (Topics), 12(1),* 1-47.

Vermeulen, P. (2016). Estimating the Top Tail of the Wealth Distribution. *American Economic Review, 106(5),* 646-650.