The Distributional Consequences of Belief Heterogeneity*

Nigel McClung
Bank of Finland
Research Economist

Tristan Nighswander
Reed College
Assistant Professor

This Version: September 10, 2020

Abstract

This paper studies a life-cycle model economy in which agents utilize idiosyncratic information to forecast personal income risk. Agents either inherit full information ("dynastic households") or learn to forecast lifetime income using personal employment data given incorrect initial beliefs ("non-dynastic households). We find that the distribution of initial beliefs regarding the transition probabilities into and out of unemployment greatly impact both aggregate savings and the stationary wealth distribution in a simple model economy. The wealth Gini coefficient in a model economy comprised solely of dynastic agents is 18%-52% lower than in the simulations with both dynastic and non-dynastic households. Optimistic households accumulate less wealth and are more likely to enter states of unemployment with little or no savings and pessimistic households over-accumulate wealth relative to households who inherit correct beliefs. Further, our framework shows how the welfare of rational agents and the value of holding correct beliefs depends on the economy-wide distribution of beliefs.

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, Bruce McGough, and Dan Tortorice for helpful comments and discussion. Corresponding author: Nigel McClung. nigel.mcclung[at]bof.fi. P.O. Box 160 Snellmaninaukio, Helsinki, FI 00101.

1 Introduction

A robust literature examines the determinants of the wealth distribution in life-cycle economies.[1] Throughout this literature, a key mechanism generating wealth inequality in model economies is the existence of uninsurable idiosyncratic income risk. Without insurance, agents amass wealth to guard against unanticipated unemployment spells, and therefore idiosyncratic employment experiences contribute to heterogeneous savings decisions over the life-cycle. As demonstrated in Huggett (1996), however, the presence of uninsurable risk is not by itself sufficient to match key features of the U.S. wealth distribution. As such, many researchers append various features to these life-cycle models to improve their empirical fit. Among these features, intergenerational ability transmission and voluntary bequests (De Nardi, 2004; De Nardi and Yang, 2014; Dynan et al., 2004), preference heterogeneity (Hendricks, 2007; Cagetti, 2003), and entrepreneurship (Quadrini, 1999; Cagetti and De Nardi, 2006) all appear to play a crucial role in moving the wealth distribution in model economies closer to empirical estimates.[2] The above mentioned papers, among others, assume full information rational expectations, and hence agents in these models form expectations of life-cycle income with complete information about their labor endowment process. This assumption precludes differences in beliefs about personal income risk and lifetime income despite a wealth of evidence of precisely this type of belief heterogeneity which we outline below. Moreover, little is known about the general equilibrium implications of this type of belief heterogeneity.

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 ("dynastic households") while the remaining

[1] See De Nardi and Fella (2017) for a survey of the literature aimed at generating realistic degrees of wealth inequality in heterogeneous agent life-cycle economies.

[2] A separate tract of the literature introduces heterogeneous rates of return on capital investments to explain the high degree of wealth inequality in the United States (see Benhabib et al., 2019; Gabaix et al., 2016,; and Kacperczyk et al., 2019).

newborn agents ("non-dynastic households") are born with incorrect beliefs about the persistence of idiosyncratic employment shocks 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. 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.

If non-dynastic households are optimistic with regard to their initial employment beliefs,[5] they will under-accumulate wealth relative to a dynastic household that experiences identical earnings over the life-cycle. If non-dynastic households are initially pessimistic, they will over-accumulate wealth relative to dynastic households. Thus, varying the distribution of optimists, pessimists, and dynastic agents in a general equilibrium model economy can result in a resource distribution with a large mass of agents with debt and long right tail (many dynastic, many optimists, many pessimists) or a relative uniform distribution of assets (many pessimists, few dynastic, few optimists). This constitutes the first contribution of this paper: a simple model economy with only two employment states, homogeneous preferences, and without retirement savings, bequest motives, or borrowing constraints that generates sensible, nontrivial changes in the wealth distribution. All we require is a modest deviation from full rationality.

Our analysis also suggests that the economy-wide distribution of initial beliefs helps to explain one's place in the wealth distribution when belief diffusion is studied in a general

[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.

[5]Optimism and pessimism are defined formally in Section 3. Informally, an optimistic (pessimistic) agent's initial beliefs lead them to overestimate (underestimate) the probability of employment in each period relative to the true exogenous employment probability.

equilibrium environment. In particular, we find extreme implications for the accumulation of assets and welfare of households with perfect information depending upon the distribution of beliefs amongst uninformed agents as the savings decisions of non-dynastic agents is an important determinant of the economy wide return on savings. Across the experiments we consider in Section 4, average wealth of dynastic households can fall by as much as 19.0% or rise by as much as 5.8% when dynastic agents occupy an economy populated with non-dynastic households relative to an economy populated with only fully informed, dynastic households. This finding has stark implications for the literature on rational inattention, which assumes agents pay a cost for information processing. Our results indicate that the *value* of acquiring full information is itself a function of the belief distribution of those around you. Therefore, this paper also contributes a framework that relates the value of learning to the distribution of beliefs in the population.

We conclude our analysis by asking our simple model economy to match several common statistical targets in the heterogeneous agent literature. We fix all preference and learning parameters and vary the proportion of households with correct, optimistic, and pessimistic beliefs. Our matching exercise indicates an extreme preference for non-rational households in the model economy. Across all target statistics, the best fit is established when all households are non-dynastic and the majority of households are optimistic under-savers. We stress that this final exercise is not a careful empirical exercise, but it illustrates how one can go a long way towards explaining features of the U.S. wealth distribution by adding a parsimonious specification of heterogeneous beliefs to an incredibly simple model. In future work we will embed these heterogeneous beliefs into richer, quantitative models.

Our view that labor market expectations are diverse and impacted by personal experience is supported by a growing empirical literature. For example, a 2019 College Pulse survey asked 7,000 students in the U.S. how much money they expect to make after graduation, and found that while the median salary for college graduates with 0-5 years experience is $48,000, the median expected salary was \$60,000.[6] Over-optimism is not unique to this survey: Jerrim (2015) shows that college-aged people overestimate life-time income by 40%, Alesina et al. (2018) provide evidence supporting the "temporarily embarrassed millionaire"[7] view held by Americans, as respondents in their study predict significantly higher upward mobility than data indicates, and Mueller et al. (2018) document over-optimism among unemployed workers with respect to their beliefs about their reemployment prospects. There is compelling evidence that individuals may suffer from over-pessimism as well. Tortorice (2012) shows that Michigan Survey of Consumers respondents are over-pessimistic about employment experiences after recessions. Rozyspal and Schlafmann (2019) corroborate this finding, noting that individuals are more likely to be pessimistic than optimistic when forecasting their personal income trajectory.

Finally, there is 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. SCE data suggests that less than .2% of variation in beliefs about the hiring rate are driven by changes in average beliefs, leading Ellison and Macauley to conclude there is "a great deal of idiosyncratic variation in *updates* to the hiring rate expectation". In light of this empirical evidence for optimism, pessimism, and general belief diffusion surrounding idiosyncratic earnings, we consider the implications of a wide range of belief calibrations in Section 4 prior to inferring the belief distribution from a moment matching exercise in Section 5. Further, our learning approach is consistent with Guvenen's findings as agents in the model

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

[7] John Steinbeck, *America and Americans*, 1966.

economy update incorrect initial beliefs about infrequently visited states (unemployment) more slowly than frequently visited states (employment).[8]

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 idiosyncratic earnings uncertainty to wealth inequality.[9] Further, 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[10] 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 HIP process itself is necessary to slow down the Bayesian learning process 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. Thus, we focus on a much simpler earnings process that is symmetric across all households regardless of age in order to establish the power of belief diffusion to reshape the economy wide

[8]See section 3 for a discussion of the mechanisms underlying our results.

[9]In an infinite-horizon model with uninsurable income risk, Giusto (2014) studies the income and wealth distribution when agents learn about the *aggregate* capital process.

[10]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.

asset distribution. As in Chang et al., Guvenen, and Guvenen and Smith, young agents in our model have greater forecast errors than older agents. 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.

Our contributions complement several other recent papers in which the implication of belief heterogeneity is explored in similar modeling environments, particularly Cogley and Sargent (2008) and Ellison and Macauley (2019). Cogley and Sargent examine the impact of incorrect prior beliefs and learning about idiosyncratic employment processes on the savings decisions of a finite lived household. Although the learning mechanism underlying our results is quite similar, the focus of our analysis is considerably different. For one, we focus on the interaction of diffuse prior beliefs and idiosyncratic belief updating on the distribution of resources, whereas Cogley and Sargent are concerned with the similarities and differences between Bayesian learning and learning in an anticipated utility framework. Second, we study general equilibrium, and our primary results reflect the essential impact of general equilibrium price adjustments on economic outcomes. Ellison and Macauley show that diffuse beliefs about the hiring rate in an economy may persist due to a combination of poor prior beliefs and a belief updating process that is primarily driven by idiosyncratic, not aggregate, factors. In contrast to our work, Ellison and Macauley deviate from full information rational expectations by assuming rational inattention and they examine the implications of incomplete information about re-employment probability for the multiplicity of equilibria.

The remainder of this paper is organized as follows. Section 2 develops the model used throughout our analysis and Section 3 provides a high level discussion of the mechanisms at play in our model. Section 4 presents our benchmark results and explores the welfare effects of non-dynastic households via general equilibrium price adjustments on fully informed dynastic households. Section 5 offers a matching exercise in which we infer the distribution of initial beliefs by targeting common statistics in the wealth inequality literature. Section 6 considers robustness concerns and Section 7 concludes.

2 The Model

The modeling environment is a multi-period overlapping generations model in which households are subject to idiosyncratic earnings risk. A household is uniquely identified by the family dynasty they are born into ($i \in 1:N$), their age ($j \in 1:J$), and the time period, $t \in (1:T)$, into which they are born. Thus, a household of dynasty i born in time t is of age $j=1$ in period t, lives for J periods and upon death is replaced by a new household of dynasty i in period $t+J+1$. As such, the population of the economy at time t is given by $N_t = NJ$.

An agent is tasked with choosing a sequence of savings allocations, $a^{j,i}_{t+j-1}$, for $j = 1, \ldots, J$ with terminal assets $a^{J,i}_{t+J-1} = 0$, as our agents have no bequest motive and consume all available resources in the final period of life. An agent's lifespan is non-stochastic and known to each agent at the beginning of their life. Agents inelastically supply their labor in every period, however their labor productivity is subject to an idiosyncratic shock that each household must forecast. The asset allocations above are selected to solve the household's optimization problem in each period, outlined below for an arbitrary agent of dynasty i born at time t.

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

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

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 $\epsilon(L)$.

Furthermore, we assume that agents are born with no wealth (i.e. $a^{0,i}_{t-1} = 0$ for all t and i) and that 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:

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

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. In subsequent projects, non-rational agents will be tasked with learning other features of their employment process (e.g. a realistically calibrated economy wide age earnings profile, the dispersion of earnings over the life-cycle, etc). However, 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). Finally, to pin down each agent's dynamic programming problem, we need to specify the evolution of $P^{e,j,i}_{H,t+j-1}$ and $P^{e,j,i}_{L,t+j-1}$.

Assumption 1 *A proportion of each generation, ϕ, have true beliefs about the idiosyncratic income transition probabilities*

$$P^{e,j,i}_{k,t+j-1} = P_k$$

for $k = H, L$, for all t, and for all i, j in the ϕ proportion of households. Households belonging to this ϕ proportion are called ***dynastic households****.*

Dynastic households always have true beliefs about the persistence of employment and unemployment states. How might a household come to form dynastic beliefs? To illustrate one potential explanation, suppose that an adaptive learning agent estimates P_H and P_L using an infinite history of individuals' employment states. Since the employment state process is exogenous, this agent will surely learn the transition probabilities if they use, for example, the following recursive estimator: $P_{k,t} = t_k^{-1}\xi_t + (1 - t_k^{-1})P_{k,t-1}$ if $s_{t-1} = k$ where $t_k \leq t$ is the number of periods that state $k = H, L$ has been realized up until $t - 1$ and $\xi_t = 1$ if $s_t = k$ and 0 if $s_t \neq k$. When $s_{t-1} \neq k$ the recursive estimator is $P_{k,t} = P_{k,t-1}$. The learnability of the moments of exogenous processes is so taken for granted in the adaptive learning literature that most applications simply endow agents with complete information about exogenous processes. Generally, these applications involve infinitely lived agents who are learning about aggregate disturbances to the economy.

In our case, the learning agent would need to collect a lot of personal employment data from multiple finitely lived agents to learn these transition probabilities under general assumptions. How might an agent obtain this information? For one, a public institution could collect and publish these data, but it would be unreasonable to assume that all agents can access this data or would choose to utilize it once given access. Alternatively, we could assume that information is made available within dynasties, i, (e.g. by generations who pass down personal income data to the next generation). This is why we refer to households with always true beliefs about transition probabilities, p_H and p_L, as dynastic households.

Assumption 2 *A proportion of each generation, $1-\phi$, form beliefs according to*

$$\begin{aligned} P^{e,j,i}_{k,t+j-1} &= \xi^{i,j}_{k,t-j-1}\left(\xi^{i,j-1}_{k,t-j-2}(\gamma_k - \gamma_k P^{e,j-1,i}_{k,t+j-2}) + P^{e,j-1,i}_{k,t+j-2}\right) \\ &+ \left(1-\xi^{i,j}_{k,t-j-1}\right)\left(1-\gamma_k \xi^{i,j-1}_{k,t-j-2}\right) P^{e,j-1,i}_{k,t+j-2} \end{aligned} \quad (4)$$

*given $P^{e,1,i}_{k,t}$ for all $i, j > 2$ in the $1-\phi$ proportion of households, where $0 < \gamma_k < 1$, $\xi^{i,j}_{k,t-j-1} = 1$ if $s^{j,i}_{t-j-1} = k$ and 0 otherwise, and $k = H, L$. Households belonging to this $1-\phi$ proportion are called **non-dynastic households**.*

Two key features of non-dynastic expectation formation distinguish non-dynastic households from dynastic households. First, non-dynastic households do not know the true transition probabilities, but try to learn these payoff-relevant parameters using personal employment data (i.e. $s^{j,i}_{t+j-1}$) and the learning algorithm (4). Notice that (4) becomes the above mentioned recursive estimator of transition probabilities if we set γ_k equal to $1/t_k$ where k is the number of periods agent (i, j) experienced state k. Hence (4) can be viewed as a constant gain learning algorithm with gain parameter, γ_k, that delivers a weighted average of the transition probabilities.

Second, $P^{e,1,i}_{k,t} \neq P_k$ and $P^{e,j,i}_{k,t+j-1} \neq P_k$ are both possible under non-dynastic learning. This means that agents can overestimate or underestimate their expected future income stream over the life-cycle based on their initial beliefs, $P^{e,1,i}_{k,t}$, or based on personal experiences over the life-cycle. How are these initial beliefs assigned and how do they interact with the learning process?

Definition 1 *An individual (i, j) born at t is said to be an **optimist** or optimistic if $P^{e,1,i}_{H,t} \geq P_H$ and $P^{e,1,i}_{L,t} \leq P_L$, with strict inequality holding for at least one.*

An optimist overestimates the probability of being in the high employment state relative to the unemployment state. Since agents' expectation of future income depends positively (negatively) on $P^{e,j,i}_{H,t+j-1}$ ($P^{e,j,i}_{L,t+j-1}$) through (4), an optimist will overestimate their expected future income stream early in life relative to dynastic, realistic or pessimistic non-dynastic households.

Definition 2 *An individual (i, j) born at t is said to be a **pessimist** or pessimistic if $P^{e,1,i}_{H,t} \leq P_H$ and/or $P^{e,1,i}_{L,t} \geq P_L$, with strict inequality holding for at least one.*

A pessimist underestimates the probability of being in the high employment state relative to the unemployment state. Consequently, a pessimist will underestimate their expected future income stream early in life relative to dynastic, realistic or optimistic non-dynastic households.

Definition 3 *An individual (i, j) born at t is said to be an **realist** or realistic if $P^{e,1,i}_{H,t} = P_H$ and $P^{e,1,i}_{L,t} = P_L$.*

A realist will tend to behave similarly to a dynastic household early in life. However, a non-dynastic realist is distinct from a dynastic agent, since the former continually updates their beliefs about idiosyncratic employment risk over their life-cycle. Hence, a realist could become optimistic if they experience persistent employment spells, or they can become a pessimist if they experience frequent unemployment. Similarly, a pessimist could become an optimist, or an optimist a pessimist, as (4) interacts with personal life experiences to endogenously overturn the systematic biases that accompany initial beliefs.

Our specification of heterogeneous beliefs implies a rich distribution over agent's types. As in other simple models, an agent's type is pinned down by their age, j, their asset holdings, a, and their employment status, ϵ, but agents are also distinguished by their beliefs, $P^{e,j,i} = (P^{e,j,i}_{H,t+j-1}, P^{e,j,i}_{L,t+j-1})$ which are entirely determined by the agent's initial beliefs, their age, and their employment histories. Importantly, the distribution of beliefs in the economy is time-invariant (stationary) because the distributions over initial beliefs, employment, and the economy's age structure is also time-invariant. This allows us to write the time t distribution over agents' types as $\Lambda(x, P^{e,j,i})$.

To close the model, we assume a standard aggregate production function: $Y_t = K_t^{\alpha} H_t^{1-\alpha}$, where H_t is the effective labor supplied to the market in period t, K_t is the aggregate capital stock at t, and Y_t is aggregate output. Factor prices are determined by profit maximization in perfectly competitive markets such that

$$\begin{aligned} R_t &= \alpha K_t^{\alpha-1} H_t^{1-\alpha} + 1 - \delta && (5) \\ w_t &= (1-\alpha) K_t^{\alpha} H_t^{-\alpha}. && (6) \end{aligned}$$

where δ is the rate of capital depreciation. The economy consists of three markets that need to clear in each period. First, the labor market clears if and only if the number of effective hours worked, H_t, equals the number of labor inputs in the economy at time t:

$$H_t = \sum_{j=1}^{J} \sum_{i=1}^{N} \epsilon(s_t^{j,i}) \tag{7}$$

Second, the asset market clears if and only if the capital stock, K_{t+1}, equals the sum of household savings at t:

$$K_{t+1} = \sum_{j=1}^{J} \sum_{i=1}^{N} a_t^{j,i} \tag{8}$$

Finally, the goods market clears by Walras' Law.

A stationary, competitive equilibrium is a collection of aggregate quantities (K, H, Y), prices (r, w), continuation values $(V(x))$, policy functions $(a(x),\ c(x))$, and a distribution of agents' types $(\Lambda(x, P^e))$ such that: 1. policy functions and value continuation functions solve the household's optimization problem; 2. firms maximize profits; 3. the distribution over household types, $\Lambda(x, P^{e,j,i})$, is stationary; 4. prices are given by (4) and (5); 5. markets clear.

3 Evolution of Beliefs: Inspecting the Mechanism

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

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 for the agents. 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 agent, 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 both realistic (green)

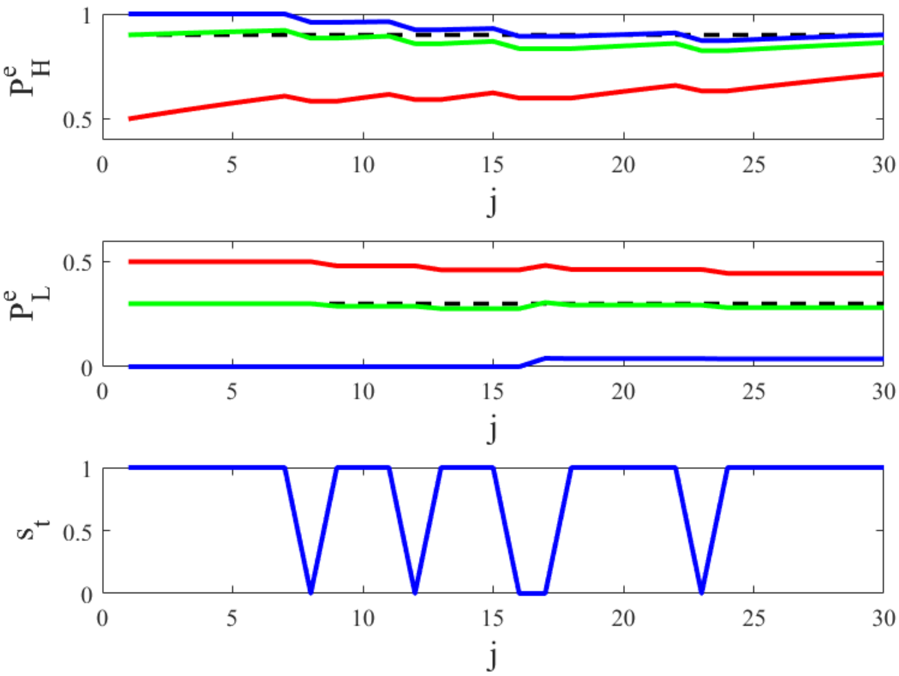

Figure 1: Belief Evolution- 30 year economic life

and pessimistic (red) agents until period 8 when the low employment state is experienced for the first time. Following a one period experience of low employment, 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 $J = 100$.

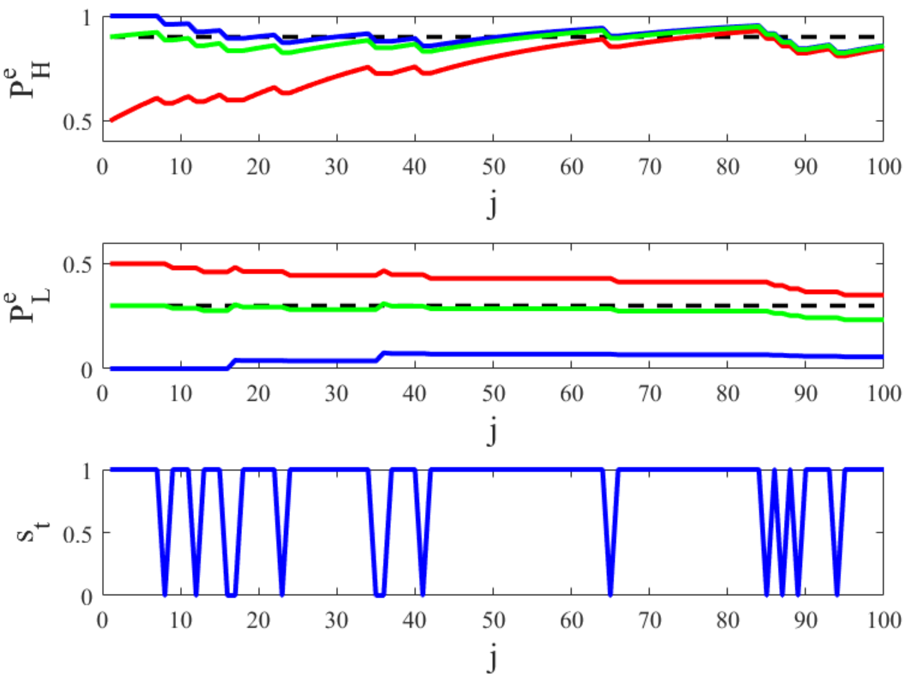

Figure 2: Belief Evolution- 100 year economic life

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 dynastic households). However, when a finite life with imperfect initial information is imposed, individuals will be unable 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.

But is this mechanism realistic? Do households have different beliefs based on their idiosyncratic experiences? Many papers discussed in the introduction support the answer "yes." But we can complement the above mentioned scientific studies with anecdotal examples. Consider an archetype which many of us are already familiar: the Depression era saver. Those who experienced job loss associated with the large, negative shock of the Great Depression at a young age and were subsequently frugal for the remainder of their working lives. The Great Recession of 2009 and the Covid-19 pandemic and related recession are likely generating similar patterns now. In a recent study outlining the impact of macroeconomic experiences and risk, Malmendier and Nagel (2011) show that preferences for risk-taking in the asset market are closely tied to the return of risky assets over one's lifetime, with a 5% point increase in experienced stock returns leading to an increased probability of stock market participation by 10% points. They argue that a portion of this observed risk taking behavior can be rationalized by a model in which agents ignore some historical data and instead predicate their beliefs on their own personal experiences. In a rational expectations framework, long-term savings behavior (and beliefs over any horizon) should be unaltered by short-run business cycle fluctuations or idiosyncratic experience. If this is true, why is it that we all likely have family members with their own economic scars who will ardently argue that the 1981-1982 Recession was dramatically worse than the Great Recession of 2009? The learning dynamics in Figures 1 and 2 are generated by a parsimonious model of the kinds of scarring effects that rational expectations frameworks lack.

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 featured in Figure 1 in each period of 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 dynastic 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. The realist and the dynastic agent have similar savings schedules, despite the fact that one's labor market expectations are relatively volatile and the other's beliefs about employment probabilities are fixed. Because dynastic agents and realistic agents are almost indistin-

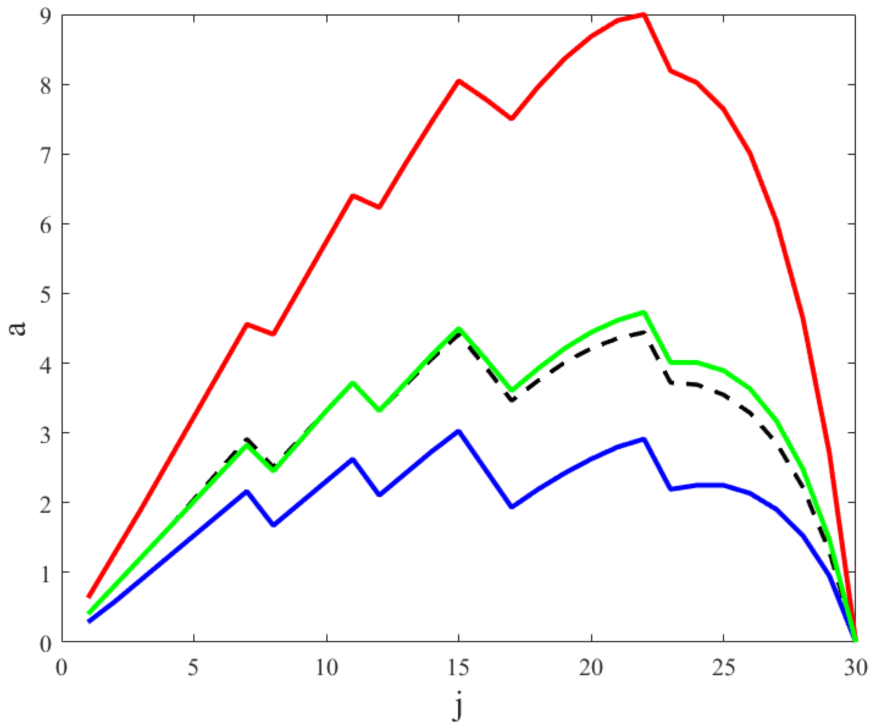

Figure 3: Asset Evolution

guishable 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 dynastic agents will have implications for aggregate savings (and therefore equilibrium interest rates and wages), and wealth inequality. We explore these implications in the following section.

4 Calibration and Results

In this section, we present results from three different experiments, distinguished by the proportion of agents in each experiment with correctly specified beliefs (ϕ). In experiment 1, $\phi = 75\%$, meaning three quarters of households are dynastic and a quarter of all households are non-dynastic. In experiment 2, $\phi = 50\%$, and in experiment 3, $\phi = 25\%$. Statistics from each experiment are compared to a baseline in which all agents are dynastic ($\phi = 1$).

Table 1: Calibration

Parameter	Value	Interpretation
β	0.96	Discount Rate
σ	2.0	IES
α	0.33	Capital Share
δ	0.075	Depreciation rate of capital
P_L	0.3	$Pr(\epsilon' = \epsilon_L \| \epsilon = \epsilon_L)$
P_H	0.9	$Pr(\epsilon' = \epsilon_H \| \epsilon = \epsilon_H)$
P_L^{opt}	0	Optimist Initial Belief: $Pr(\epsilon' = \epsilon_L \| \epsilon = \epsilon_L)$
P_H^{opt}	1	Optimist Initial Belief: $Pr(\epsilon' = \epsilon_H \| \epsilon = \epsilon_H)$
P_L^{pes}	.5	Pessimist Initial Belief: $Pr(\epsilon' = \epsilon_L \| \epsilon = \epsilon_L)$
P_H^{pes}	.5	Pessimist Initial Belief: $Pr(\epsilon' = \epsilon_H \| \epsilon = \epsilon_H)$
$\epsilon(H)$	1	Payoff if employed
$\epsilon(L)$	0.1	Payoff if unemployed
γ_H	0.04	Gain parameter learning on high state
γ_L	0.04	Gain parameter learning on low state
J	30	Length of agent's life

Across each experiment, we hold all preference parameters (β and σ), firm and labor market parameters (α, δ, P_L, P_H, $\epsilon(H)$, and $\epsilon(L)$), learning parameters (γ_H and γ_L), initial beliefs (P_L^{opt}, P_H^{opt} , P_L^{pes} , P_H^{pes}), and demographic parameters (J) fixed. Within an experiment, we vary the proportion of non-dynastic households with pessimistic beliefs (λ) and optimistic beliefs ($1 - \lambda$). The calibration used for our primary analysis is outlined in Table 1.

Our choice of preference parameters, learning parameters, and α are all standard in the literature. The transition probabilities governing the labor earnings process (P_L and P_H) were calibrated to match estimated employment transition probabilities from PSID data in Ashman and Neumuller (2019).[11] Given these transition probabilities, the likelihood of back-to-back years of unemployment is low, but not highly improbable (.3) and the vast majority of agents will experience at least one bout of unemployment in their working lives (probability of always being employed = $P_H^J = .9^{30} = .042$). Payoffs in each employment state, ($\epsilon(H)$ and $\epsilon(L)$) were selected so that the high employment state corresponds to receiving the

[11] 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 $P_L \in (.09, .50)$.

economy-wide wage and the low-employment state corresponds to receiving 10% of said wage. We chose a non-zero payout in our low employment state to mimic the low, but non-zero, replacement wage associated with a year of unemployment. Our calibration of the terms governing optimism and pessimism (P_L^{opt}, P_H^{opt}, P_L^{pes}, and P_H^{pes}) 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] We discuss alternative model calibrations in Section 6 and the appendix details our numerical approach to simulating the model economy.

Section 4.1 examines how variation in the belief distribution affects the aggregate wealth distribution. In Section 4.2, we turn our attention to the asset distribution of optimistic, pessimistic, and dynastic households to argue that although initial beliefs are a large component of agent sorting, they do not entirely determine one's place in the wealth distribution. Section 4.3 studies the welfare of dynastic households across different calibrations. In this section we are able to see the general equilibrium consequences of belief diffusion as the wealth of perfectly informed households is influenced as we vary the distribution of beliefs amongst non-dynastic agents.

4.1 Wealth Effects of Imperfect Information

As outlined above, our experiments are defined by the proportion of agents born with correct information about their idiosyncratic earnings process (ϕ). In each experiment, we vary the proportion of non-dynastic agents who begin their economic lives with pessimistic expectations (λ) from $\lambda = .75$ (High Pessimism or "PessH") to $\lambda = .5$ (Medium Pessimism or "PessM") to $\lambda = .25$ (Low Pessimism, or "PessL"). All other non-dynastic agents are assigned optimistic beliefs. Figures 4-6 display the aggregate wealth distribution in each experiment as λ varies. Figure 4 outlines the wealth distribution in experiment 1, figure 5

[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.

shows the wealth distribution for experiment 2, and figure 6 shows the wealth distribution for experiment 3.

Figure 4: Wealth Distribution in Experiment 1: $\phi = .75$

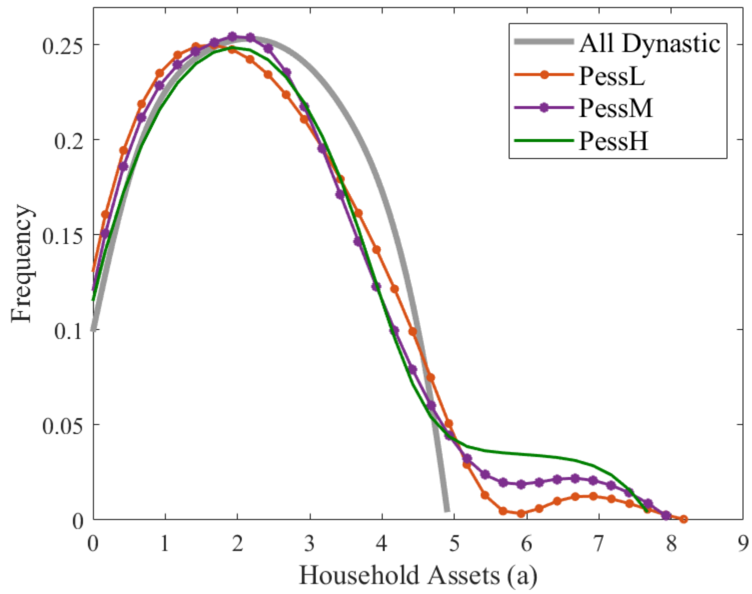

Figure 5: Wealth Distribution in Experiment 2: $\phi = .5$

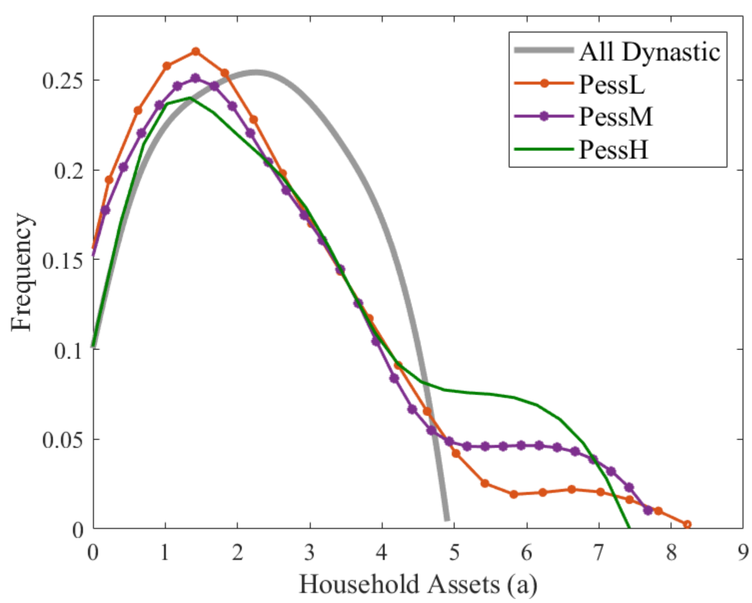

Figure 6: Wealth Distribution in Experiment 3: $\phi = .25$

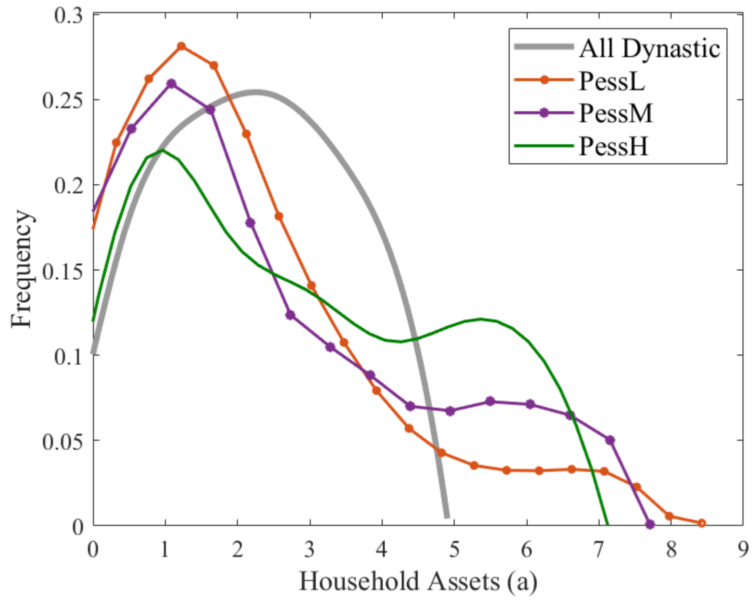

Beginning with figure 4, we see that the high precautionary savings motives of pessimistic households generates a long right tail in the aggregate wealth distribution. As the proportion of pessimists (λ) increases, this tail grows thicker. However, when non-dynastic households make up just 25% of the population, the shape of the wealth distribution is only mildly impacted by the increased proportion of pessimists amongst the population of poorly informed non-dynastic households. As we move to figure 5, we see the proportion of pessimists begin to dramatically re-shape the aggregate distribution of wealth. There is a considerable rise in the fraction of households holding debt (see the increased vertical intercept of the wealth distribution) as λ decreases due to the proportional rise in optimists amongst the non-dynastic population. Further, the rise in λ leads to a distribution which appears more bi-modal than unimodal with a long right tail.

Figure 6 shows just how dramatically the wealth distribution can be re-shaped by the distribution of non-dynastic beliefs. With many optimists and few pessimists, the wealth distribution is similar to figures 4 and 5. However, when pessimists make up the majority of the population ($\lambda(1 - \phi) > .5$), the aggregate resource distribution flattens considerably and a long right tail is no longer observed. This occurs as the presence of many pessimists (with a high preference for savings) drives down the market rate of interest. Thus, the amassing of extreme fortunes is less likely due to lower returns on accumulated wealth but more households occupy a position in the distribution above the mean due to their high savings desire.

Table 2: Model Economy Wealth Statistics

	K	r	% borr	%ND borr	Gini	90-10	90-50	50-10
Benchmark: 100% Dynastic								
Baseline	2.46	10.6%	1.5%	NA	31.1	6.3	1.8	3.5
Experiment 1: 75% Dynastic								
PessH	2.65	9.7%	2.2%	3.7%	37.8	7.4	2.0	3.7
PessM	2.52	10.3%	2.6%	5.5%	37.4	7.7	2.0	3.9
PessL	2.42	10.8%	3.0%	7.4%	37.2	8.5	2.0	4.3
Experiment 2: 50% Dynastic								
Pess H	2.85	8.9%	3.2%	4.5%	40.1	10.9	2.5	4.3
PessM	2.60	9.9%	4.0%	6.2%	42.7	12.9	2.5	5.1
PessL	2.37	11.0%	4.4%	7.3%	42.9	13.8	2.3	6.0
Experiment 3: 25% Dynastic								
PessH	3.06	8.1%	5.0%	5.9%	42.8	17.1	2.4	7.2
PessM	2.67	9.6%	5.4%	6.7%	47.4	18.4	3.1	5.9
PessL	2.33	11.2%	5.4%	6.7%	47.1	15.2	2.7	5.6

Table 2 outlines a number of summary statistics for each of our three experiments and provides baseline data for a model economy comprised entirely of dynastic agents. We see that average wealth holding (K) and the equilibrium interest rate (r) are both highly sensitive to the fraction of high-saving pessimistic and low-saving optimistic households (changes in ϕ or λ). Across all three experiments, aggregate wealth falls and the interest rate rises as the proportion of pessimists decreases, either holding ϕ fixed and varying λ or holding λ fixed and varying ϕ. In either case, this is a function of fewer households with a high desire to save in order to insure against unrealistically high expectations of future employment loss. Further, the fraction of households borrowing (% borr) and the fraction of non-dynastic households in debt (%ND borr) increase as the fraction of pessimists decreases. This occurs despite the fact that high pessimism lowers the general equilibrium interest rate and makes borrowing more attractive, which should increase the fraction of households in debt. However, decreased pessimism corresponds to increased optimism for a given experiment and optimists are more likely to borrow at any market interest rate than pessimists or dynastic households.

The one exception to this result is in experiment 3, PessL. In this experiment, the majority of households are optimistic (75% non-dynastic households of which 75% are optimists) which should lead to higher rates of borrowing. However, equilibrium prices are primarily driven by the majority group of optimists and their low savings desire. Thus, the high interest rate associated with a highly optimistic society leads to a constant fraction of borrowers between the PessM and PessL simulations. This reinforces the importance of tracking general equilibrium price adjustments when analyzing life-cycle model economies. Prices are highly sensitive to agent beliefs, which means the economic impact of more optimism or pessimism in the vacuum of partial equilibrium will mis-state the true effect of these forecast errors.

Inequality in the model economy, measured by the Gini coefficient, 90-10, 90-50, and 50-10 ratios, is increasing in the proportion of non-dynastic households $(1-\phi)$ for any fixed value of pessimism (λ). Relative to the baseline in which only dynastic households are modeled, the Gini coefficient of the wealth distribution is between 18% (Experiment 1 PessH) and 52% (Experiment 3 PessM) higher across our experiments. Although the magnitude of this inequality change is quite large, the qualitative change is unsurprising as a higher proportion of non-dynastic households means higher diffusion of beliefs about the attractiveness of the equilibrium interest rate and thus more diffuse capital accumulation preferences. However, inequality has a non-linear relationship with the proportion of pessimistic agents in any given experiment with a fixed ϕ. This occurs as increasing λ both drives down the return to savings (leading to a more condensed distribution of non-pessimistic households) and also increases the mass of agents above the mean in the distribution (either pessimists or dynastic/optimistic households that have received consistently high labor endowment draws).

As λ falls, the 90-10 ratio increases in the majority of cases outlined above. However, the 90-10 ratio decreases in experiment 3 with extreme optimism (PessL) as the high market interest rate does more to encourage savings at the bottom end of the distribution than the top. Recall, a lower λ means a longer right tail, but a very low fraction of pessimists means the density in that tail is low. Across experiments 1 and 2, the 50-10 ratio rises as λ is decreased, but in experiment 3 the 50-10 ratio falls in response to decreased pessimism. This is likely driven by the dramatic change in the market interest rate across each calibration of experiment 3

In light of the conflicting empirical support for optimism, pessimism, and belief diffusion regarding idiosyncratic earnings forecasts, we do not take a stand on a "preferred" model. However, the dramatic general equilibrium implications of belief diffusion on the aggregate wealth distribution indicates that a better understanding of precisely how households form beliefs could go a long way towards explaining the empirical distribution of wealth.

4.2 Wealth Distribution by Subpopulation

We now turn our attention to examining the impact of changing λ and ϕ on the wealth of each subpopulation (optimist, pessimist, dynastic) in order to rationalize our findings in Section 4.1. Table 3 displays average capital for the model economy as a whole ($\bar{K}$), average wealth for each sub-population ($\bar{K}_{opt}$, $\bar{K}_{pess}$, $\bar{K}_{dyn}$ for optimists, pessimists, and dynastic agents, respectively), and each sub-population average relative to $\bar{K}$ in the model economy. As argued in Section 4.1, on average optimists under-accumulate wealth (15%-57% of $\bar{K}$ across experiments), pessimists over-accumulate wealth (144%-214% of $\bar{K}$ across experiments), and dynastic households are heavily influenced by the presence and belief distribution of non-dynastic households (65%-111% of $\bar{K}$ across experiments).

When there are many pessimists (Experiment 3, PessH), the market interest rate is driven down by the high savings incentive of these agents. Thus, the average wealth of pessimists is just 144% of the average wealth in the economy, corresponding to the flattened, nearly bi-modal distribution displayed in figure 6 in red. As the proportion of pessimists

Table 3: Capital by Agent Type

	$\bar{K}$	$\bar{K}_{opt}$	$\frac{\bar{K}_{opt}}{K}$	$\bar{K}_{pess}$	$\frac{\bar{K}_{pess}}{K}$	$\bar{K}_{dyn}$	$\frac{K_{dyn}}{K}$	$var(K_{dyn})$
Experiment 1: 75% Dynastic								
PessH	2.44	0.90	37%	4.30	177%	2.09	86%	1.30
PessM	2.32	1.00	43%	4.41	190%	2.19	94%	1.40
PessL	2.22	1.11	50%	4.50	203%	2.31	104%	1.57
Experiment 2: 50% Dynastic								
PessH	2.62	0.67	25%	4.15	159%	1.95	74%	1.14
PessM	2.39	0.94	39%	4.33	181%	2.13	89%	1.33
PessL	2.17	1.14	52%	4.52	208%	2.36	108%	1.64
Experiment 3: 25% Dynastic								
PessH	2.81	0.44	15%	4.04	144%	1.83	65%	.99
PessM	2.45	0.87	35%	4.29	175%	2.07	84%	1.29
PessL	2.14	1.22	57%	4.57	214%	2.39	111%	1.67

falls, market interest rates are driven by either optimists or dynastic households, both of whom have considerably lower preferences for savings. Pessimists accumulate their highest average capital and their highest relative wealth (214% of the economy wide average) when the majority of agents are optimistic (Experiment 3, PessL), corresponding to the green distribution in figure 6 wherein the distribution has the longest, fattest tail out of any of the model economies considered.

Now that we have established some of the intuition behind the economy-wide distributions displayed in figures 4-6, we turn to the following question: Is one's place in the aggregate wealth distribution entirely driven by initial beliefs? Table 3 provides strong evidence that across all three experiments, average wealth is quite different between each sub-population. However, it tells us nothing about the likelihood of an optimist acquiring more wealth than a pessimist, or a dynastic agent acquiring more debt than an optimist. In order to address this issue, consider figures 7-9. In each figure, we display the wealth distribution for experiment 2 for PessH (figure 7), PessM (figure 8) and PessL (figure 9)[13].

As these figures show, although it is more likely for pessimists to accumulate wealth above

[13] See On-line Appendix Q for the full distribution associated with experiments 1 and 3

Figure 7: Wealth Distribution by Sub-population: $\phi = .50$; $\lambda = .75$

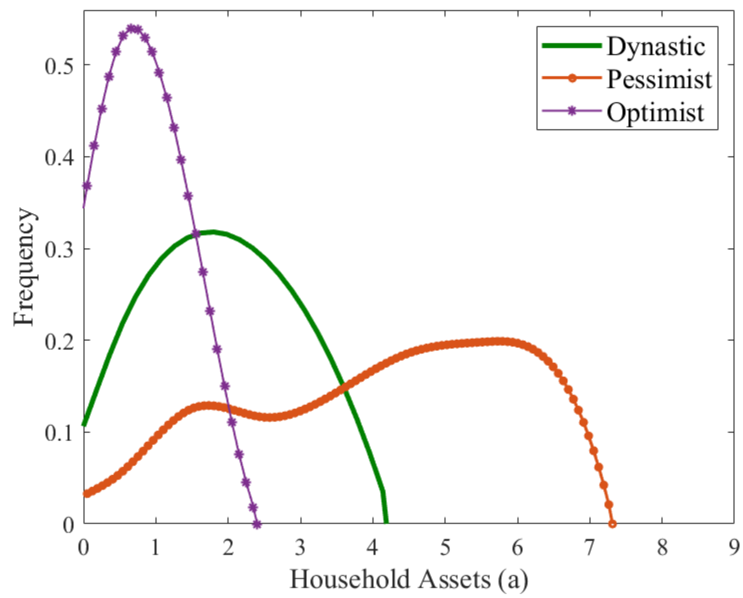

Figure 8: Wealth Distribution by Sub-population: $\phi = .50$; $\lambda = .5$

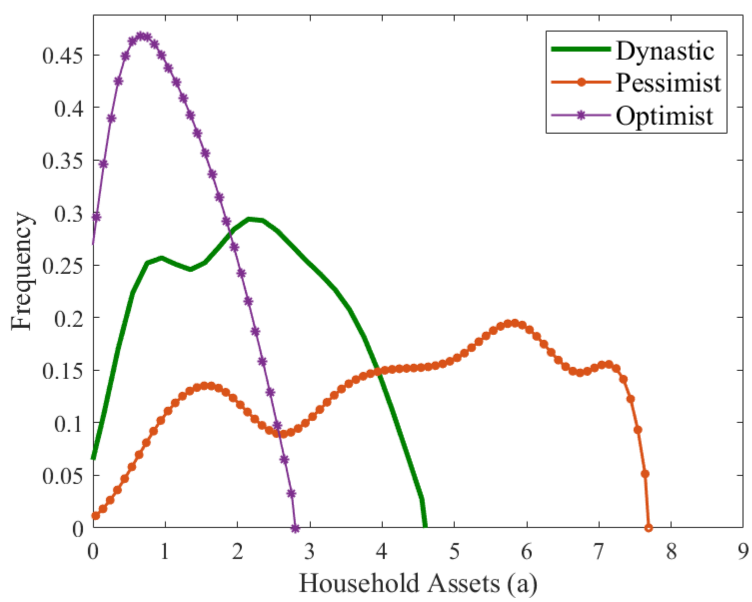

Figure 9: Wealth Distribution by Sub-population: $\phi = .50$; $\lambda = .25$

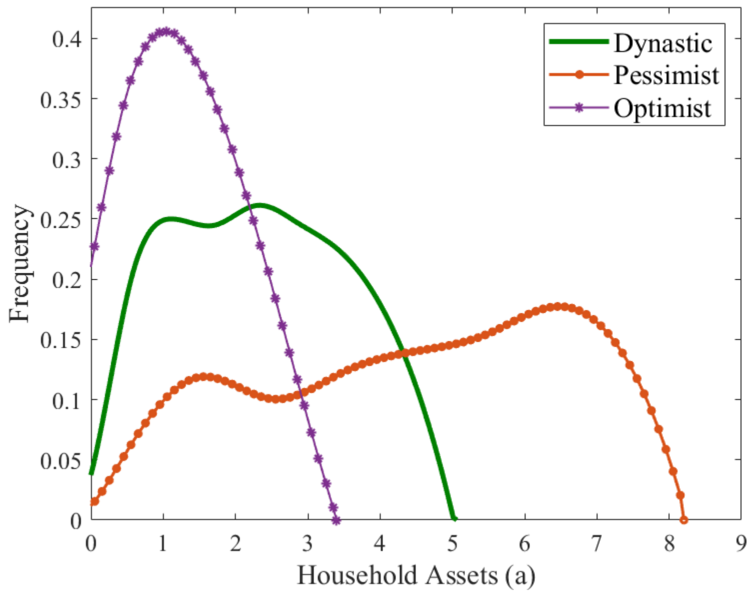

the economy-wide average and for optimists to hold less wealth than average, one's initial beliefs are not a sink. Many agents who are born optimistic manage to enter the top half of the wealth distribution, while many pessimistic agents fall well below the wealth holdings of dynastic and optimistic households. This movement throughout the distribution can occur in multiple ways. An optimist could acquire above average wealth by either receiving a series of exceptional employment draws (leading to high lifetime income but never shedding their optimism) or by receiving a series of poor employment draws early in life (leading to a pessimistic world-view and thus a high savings propensity). Similarly, a pessimist with consistently low employment draws could fail to acquire the high wealth of most pessimists due to insufficient lifetime income or they could fail to move up in the wealth distribution due to a series of positive employment shocks that turn their pessimism to optimism (thus their high savings propensity to low savings propensity). Because one's world view evolves in conjunction with one's labor market experiences, there are numerous channels through which initial optimism or pessimism can interact with labor market experiences in order to generate lifetime outcomes.

The one group insulated from the interaction of idiosyncratic experience and beliefs are the dynastic households, born with complete knowledge of the evolution of employment draws. However, this perfect knowledge of the exogenous employment distribution does not insulate dynastic agents from the impact of belief evolution on their own consumption and savings decision. In Section 4.3, we argue that the wealth accumulation and welfare of dynastic agents depends crucially on the beliefs and actions of their non-dynastic peers via the general equilibrium impact of imperfect information on prices.

4.3 Savings and Welfare of Dynastic Households

The extent to which well-informed (dynastic) households suffer from occupying an economy with uninformed (non-dynastic) agents is a function of both the fraction of non-dynastic households and the distribution of non-dynastic agents across optimistic and pessimistic beliefs. Pessimists prefer relatively high market interest rates to amass wealth for fear of future unemployment and optimists prefer relatively low market interest rates given their preference for borrowing to smooth their consumption out of high anticipated future income. Thus, the fraction of misinformed agents and the distribution of beliefs over these agents has large implications for the equilibrium interest rate and wage, which determine the expected lifetime utility of dynastic households. In particular, we find that for any proportion of non-dynastic households, dynastic agents are made worse off (lower average wealth holdings) as the fraction pessimists in the uninformed population increases. Straightforwardly, this is due to the fact that pessimists lower the market interest rate.

Table 4 displays average wealth and variance in wealth for dynastic agents when all individuals have full information ($\bar{K}^*$ and $var(K^*)$) [14], average wealth and variance in wealth for dynastic households in each experiment discussed in Section 4.2 ($\bar{K}_{dyn}$ and $var(K_{dyn})$), and the ratio of these statistics. We include the variance of wealth given the common assumption of CRRA utility in the literature which we adhere to in this paper. Higher variance

[14] This model corresponds to the baseline model displayed in Table 2.

Table 4: Dynastic Capital Holdings

	$\bar{K}^*$	$\bar{K}_{dyn}$	$\frac{\bar{K}_{dyn}}{\bar{K}^*}$	$var(K^*)$	$var(K_{dyn})$	$\frac{var_(K_{dyn})}{var(K^*)}$
Experiment 1: 75% Dynastic						
PessH	2.26	2.09	93%	1.49	1.30	87 %
PessM	2.26	2.19	97%	1.49	1.40	94%
PessL	2.26	2.31	102%	1.49	1.57	105%
Experiment 2: 50% Dynastic						
Pess H	2.26	1.95	86%	1.49	1.14	76%
PessM	2.26	2.13	94%	1.49	1.33	89%
PessL	2.26	2.36	104%	1.49	1.64	110%
Experiment 3: 25% Dynastic						
PessH	2.26	1.83	81%	1.49	0.99	66%
PessM	2.26	2.07	92%	1.49	1.29	86%
PessL	2.26	2.39	106%	1.49	1.67	112%

in wealth for a fixed average represents lower economy wide welfare, as households have a preference for certainty under this utility specification. When pessimistic households dominate the uninformed distribution, informed agents suffer as the high desire to save of these non-dynastic households drives down the market interest rate and thus lowers the returns to savings of dynastic agents. However, when optimistic agents dominate the uninformed distribution, dynastic households accumulate more wealth on average when occupying an economy with uninformed individuals, but this comes at the expense of increased variance in the distribution of dynastic wealth.

Consider the transition between the highly uninformed (25% dynastic) economy to the highly informed (75% dynastic) economy. With high pessimism, average K is 1.83, just 81% of average wealth in the all dynastic economy and 2.09 in the highly informed economy, or 93% of average wealth in the all dynastic economy. This 14% increase in average capital occurs as the proportion of non-dynastic agents who are pessimistic is held constant, but the proportion of uninformed agents decreases. However, with high optimism, moving from the highly uninformed to the highly informed economy results in a considerable increase in wealth for dynastic households. Dynastic agents go from making 102% to 106% of the average wealth in an all dynastic society. However, in each experiment in which average dynastic wealth increases relative to the all dynastic economy (PessL across all three experiments), the variance of wealth for dynastic agents increases above the variance in the all dynastic world. Although the high interest rates associated with many optimists drives up wealth accumulation, it also increases inequality in the model economy which negatively effects the welfare of dynastic households due to the heightened uncertainty they face.

Focusing on experiment 2 in which the proportion of dynastic agents is fixed ($\phi = .5$), as we vary the proportion of pessimistic households from .75 to .50 to .25 average capital holdings amongst dynastic households increases from 1.95 to 2.13, to 2.36. As a fraction of average societal wealth in an economy occupied exclusively by dynastic households, dynastic agents vary from owning 86%, to 94%, to 104%, representing a 21% increase in average wealth holdings as the proportion of pessimists decreases. Thus, the value of holding correct beliefs is a function of both the degree to which your peers are well-informed *and* the distribution of beliefs of said peers. Connecting to the literature on rational inattention in which agents weight the cost of information acquisition prior to engaging in learning, our results imply a channel through which the returns to information acquisition are determined by the beliefs of other agents. Our results more generally indicate that the value of being well-informed is highly influenced by the beliefs of those around you.

5 Stylized Exercise: Choosing λ and ϕ

In Section 5, rather than varying the proportion of dynastic, optimistic, and pessimistic households and studying the resulting wealth distribution, we target moments in US data by varying the proportion of each agent type (governed by ϕ and λ). All other model parameters are held fixed at the values reported in Table 1. Given the simplified model economy we study, this exercise is not meant to perfectly match moments of the US wealth distribution but rather to identify parameterizations that move our baseline model, which falls well short of matching any target statistics, closer to empirical reality. Our aim is similar to that of Maliar and Maliar (2006) in which quasi-geometric discounting is introduced to the standard neoclassical growth model. Although the introduction is insufficient to fully capture the degree of inequality observe in the data, as in Maliar and Maliar our exercise generates a considerable improvement over a baseline model in which all agents forecast rationally.

The moments we target are the Gini coefficient, wealth holdings of the top 1%, top 10%, and bottom 40%.[15] The first line of Table 5 reports estimated values of the relevant wealth statistics in SCF data as reported in Hendricks (2007). The second line of Table 5 displays the same statistics for the benchmark model which, as in Section 4, is an economy comprised of entirely dynastic households ($\phi = 1$). The remaining rows of Table 5 correspond to calibrations of ϕ and λ that were selected to best fit the five target statistics. We highlight the value of each statistic we aim to match along the diagonal of the table.

Table 5: Moment Matching Exercise

	Gini	Top 1%	Top 10%	Bot. 40%	ϕ	λ
Data	.80	34.7	78.1	1.0	N/A	N/A
Benchmark	.31	2.1	19.2	18.1	1.0	N/A
Gini Match	.53	3.8	32.1	7.5	0.0	0.3
Top 1% Match	.47	4.4	28.9	10.1	0.0	0.1
Top 10% Match	.50	4.0	32.4	9.0	0.0	0.2
Bot. 40% Match	.52	3.0	27.0	6.2	0.0	0.5

Several findings immediately stand out from our analysis. First, the benchmark model comprised of dynastic households provides the worst fit to US data by a wide margin. This calibration accounts for a small fraction of the wealth inequality implied by US data (model Gini=0.31, data implied Gini=0.80) with both too little wealth held by the wealthiest members of society and far too much wealth held by the poorest 40% of households. Second, the model strictly prefers calibrations in which *no* households are dynastic. Whether we

[15]These statistics are common targets in the quantitative life-cycle literature aimed at re-creating the empirical wealth distribution.

are targeting the Gini coefficient, moments of the wealth distribution, or the consumption output ratio, in every case the model prefers a calibration in which $\phi = 0$. Third, no matter what statistic we chose to target, the presence of non-dynastic agents leads to all wealth statistics in the model economy moving closer to the wealth statistics associated with US data.

Consider the row titled "Gini Match" above in which ϕ and λ were jointly determined by matching the model Gini coefficient on wealth to its empirical analogue. The closest our model can come to the US wealth Gini of 0.80 is 0.53, which we achieve when $\phi = 0$ and $\lambda = 0.3$, meaning no agents are dynastic with 70% of households born optimistic and 30% of household born pessimistic. Our benchmark model generates just 39% of the wealth inequality reported in the data, whereas our model of best fit is able to account for 66% of the observed Gini coefficient. Not only does our fit improve with respect to the Gini coefficient, but the 1% and 10% share of wealth nearly double relative to the benchmark model and the bottom 40% share decreases by nearly 9% points, all representing improvements of our model fit relative to an all dynastic economy.

A similar pattern emerges as we consider other calibration targets, with $\phi = 0$ for all models and λ=0.1, 0.2, 0.5 as we match the top 1%, top 10%, and the bottom 40%, respectively. These values of λ are consistent with the empirical literature in which many individuals are overly optimistic about their employment outlook and others report pessimism. As noted in Section 4, the presence of a few pessimists considerably increases wealth inequality as many optimists drive the market interest rate up and a few pessimists are able to amass a large amount of wealth thanks to their predisposition for saving interacted with the high market return.

While the improvement in fit may appear modest, we stress that these changes are non-trivial given the nature of our modeling environment. In a model with no bequest motives, no need for retirement savings, no age-earnings relationship, and no borrowing constraints, the presence of belief diffusion regarding a simple exogenous earnings process is sufficient to more than double the wealth holdings of the top 1% while simultaneously cutting the wealth holdings of the bottom 40% in half. The inability of model economies to replicate wealth holdings at the top of the distribution is a pernicious issue in the life-cycle literature, and the results from our simplified model indicate that the previously unexplored avenue of belief diffusion may be a key driver of this gap between model and reality.

6 Sensitivity to Alternative Model Calibrations

Section 6 briefly discusses the sensitivity of our results to alternative model calibrations and assumptions about how boundedly rational learning agents make decisions given their forecasts. We argue that changes in the parameter calibrations affect results in sensible ways, without impacting the main qualitative conclusions of our analysis. We find that the following changes all increase the qualitative gaps between non-dynastic and dynastic households reported in Section 4: increasing (decreasing) P_L (P_H), decreasing (increasing) ϵ_L (ϵ_H), decreasing J, increasing σ, making optimists more optimistic (decreasing P_L^{opt} or increasing P_H^{opt}), or making pessimists more pessimistic (increasing P_L^{pes} or decreasing P_H^{pes}).

Increasing (decreasing) P_L (P_H) or decreasing (increasing) ϵ_L (ϵ_H) impact the model in the same way: each makes the unemployed state worse relative to the employed state and thus exacerbates differences between agents who disagree about the likelihood of said states occurring. Similarly, making optimists more optimistic or pessimists more pessimistic amplifies the biases of each agent type and increases the number of employment states an agent must experience before shedding their initial, incorrect beliefs. Decreasing J amounts to shortening the lives of our agents which, as discussed in Section 3, means the beliefs of non-dynastic agents are less likely to converge to the truth. Symmetrically, if we increase J, the wealth gap between agent types are minimized as eventually non-dynastic households are able to learn the exogenous employment process and only young agents who have not out-lived their poor initial beliefs remain biased.

Heterogeneity in γ_H and γ_L do very little to change our reported results, qualitatively or quantitatively. However, decreasing γ does increase the differences between the non-dynastic and dynastic economies. This occurs as a lower γ corresponds to less rapid learning, thus a higher persistence of initial (biased) beliefs.

7 Conclusion and Future Work

We show that a modest, realistic deviation from the standard assumption of rational expectations can lead to a dramatic reshaping of the economy wide distribution of wealth. Our mechanism is capable of generating a more realistic degree of skewness in the aggregate wealth distribution than a standard model comprised of fully informed rational agents. Further, we outline how the welfare of fully informed agents in a simple life-cycle model can be negatively impacted by the beliefs of mis-informed agents with whom they occupy the economy. The mechanism we employ is grounded in both a theoretical literature exploring alternatives to rational expectations and an empirical literature outlining a great deal of diffusion in beliefs regarding labor market earnings. We chose to abstract from the sophistication typically leveraged in life-cycle models aimed at matching the empirical wealth distribution in order to fully elucidate the mechanism driving our results.

Many questions remain which we intend to pursue in subsequent research. We outline several below. First, we aim to extend our modeling approach to a more sophisticated quantitative life-cycle model in which beliefs are calibrated utilizing the Survey of Consumer Expectations in order to understand the extent to which belief heterogeneity can rationalize 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 (2019). Wealth Distribution and Social Mobility in the US: A Quantitative Approach. *American Economic Review, 109(5),* 1623-1647.

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*.

Cagetti, M. (2003). Wealth Accumulation over the Life-cycle and Precautionary Savings. *Journal of Business and Economic Statistics, 21(3),* 339-353.

Cagetti, M, and M. De Nardi (2006). Entrepreneurship, Frictions and Wealth. *Journal of Political Economy, 114(5),* 835-870.

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

Cogley, T. and T. Sargent (2008). Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making. *International Economic Review, 49(1),* 185-221.

De Nardi, M. (2004). Wealth Inequality and Intergenerational Links. *The Review of Economic Studies, 71(3),* 743-768.

De Nardi, M. and F. Yang (2014). Bequests and Heterogeneity in Retirement Wealth. *European Economic Review 72,* 182-196.

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

Dynan, K., Skinner, J., and S. Zeldes (2004). Do the Rich Save More? *Journal of Political Economy, 112(2),* 397-444.

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.

Gabaix, X, Lasry, J.M, Lions, P.L., and B. Moll (2016). The Dynamics of Inequality. *Econometrica, 84(6),* 2071-2111.

Giusto, A. (2014). "Adaptive Learning and Distributional Dynamics in an Incomplete Markets Model," *Journal of Economic Dynamics and Control, 140*, 317-333.

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.

Hendricks, L. (2007). How Important is Preference Heterogeneity for Wealth Inequality? *Journal of Economic Dynamics and Control, 31(9),* 3042-3068.

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.*

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.

Kacperczyk, M, Nosal, J, and L. Stevens (2019). Investor Sophistication and Capital Income Inequality. *Journal of Monetary Economics, 107,* 18-31.

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.

Maliar, L. and S. Maliar (2006). The Neoclassical Growth Model with Heterogeneous Quasi-Geometric Consumers. *Journal of Money, Credit, and Banking, 8(3)*, 646-654.

Malmendier,U. and S. Nagel (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking? *Quarterly Journal of Economices, 126(1),* 373-416.

Marcet, A. and T. Sargent (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.*

Quadrini, V. (1999). The Importance of Entrepreneurship for Wealth Concentration and Mobility. *The Review of Income and Wealth, 45(1),* 1-19.

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

Steinbeck, J. (1966). America and Americans. Penguin Random House.

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

Appendix. Outline of Equilibrium Solution Algorithm:

1. Propose a candidate steady state capital stock, K_0, and corresponding interest rate, R, and wage, w, that solve the firm optimization problem.

2. Solve the household problem ($c(x)$ and $a(x)$) given prices R and w.

   We use a standard value function iteration approach to solve the household optimization problem. Our value function iteration approach proposes a fine grid over savings allocations, a, and a fine grid over transition probability beliefs, P_H^e and P_L^e. In each period of each non-dynastic agent's life, beliefs evolve according to (4) and are then rounded to nearest transition probability pair (P_L^e, P_H^e) in the transition probability grid.

3. Using the optimized capital decision rule, $a(x)$, compute individual savings decisions for the stable distribution of households.

4. Compute aggregate capital, K_1, the capital stock in the model economy given preferences. If K_1 is sufficiently close to K_0, stop. Otherwise, set $K_0 = K_1$ and repeat steps (1)-(4) until convergence is achieved.