

Exploring Intergenerational Wealth Transfer Dynamics with Agent-Based Models

Osonde A. Osoba, Jonathan W. Welburn, Jonathan Lamb, Pedro Nascimento de Lima, Krishna B. Kumar

RAND Education and Labor

WR-A1259-8
March 2023
Prepared for RAND Ventures

RAND working papers are intended to share researchers' latest findings and to solicit informal peer review. They have been approved for circulation by RAND Education and Labor but have not been formally edited or peer reviewed. Unless otherwise indicated, working papers can be quoted and cited without permission of the author, provided the source is clearly referred to as a working paper. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors. RAND® is a registered trademark.



For more information on this publication, visit www.rand.org/t/WRA1259-8.

About RAND

The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.

© 2023 RAND Corporation

RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

Table of Contents

Figures	iv
Tables	iv
Abbreviations	v
Notation Summary	v
Chapter 1: Background and Context	1
Chapter 2: Outline of the Model’s Foundations	4
OLG Model Structure.....	4
Specifying Generations and their Constraints.....	4
Chapter 3: Model Implementation.....	7
Structure of the Wealth ABM	7
Household’s Financial Decision-making Behaviors	9
1. A Simple Financial Decision Model	9
2. An RL-based Financial Decision Model	10
Model Parameterization and Calibration Strategy	11
Structural Parameterization	11
Partially-Supervised Calibration	13
Computational Structure and Details for the Model	14
Chapter 4: Discussion of Some Experimental Results	16
Experimental Setup	16
Observed Sensitivities to Structural Parameters	16
Chapter 5: Future Work.....	19
Future Experiments	19
“Baby Bonds” aka Endowments at Birth:	19
Education subsidies	20
The Importance of Idiosyncratic Shocks and Insurance Schemes	20
Local Real Estate Dynamics	20
Conclusion.....	21
References	22

Figures

Figure 3.1: Schematic of ABM Household Life Course	7
Figure 3.2: Finite State Depiction of ABM Household States	8
Figure 3.3: Unrolled Depiction of ABM Households Evolving in Time	8
Figure 4.1: Estimated Influence of variation in Average Rent on the Black-<->Non-Black Wealth Gap	17
Figure 4.2: Estimated Influence of Variation in the Education Premium on the Racial Wealth Gap	18
Figure 4.3: Estimated Influence of Variation in Median Incarceration Length on the Racial Wealth Gap	19

Tables

Table 3.1: Mathematical Notation for Model Specification.....	10
Table 3.2: Recoding of SCF variables for synthetic population	12
Table 3.3: Model Parameters.....	13

Abbreviations

ABM	Agent-based Model
OLG	Overlapping Generations Model
NN	Neural Network
DP	Dynamic Programming
RL	Reinforcement Learning
REDGEN	the Race, Economic Disparities and GENERational wealth (REDGEN) model

Notation Summary

c_{it} : Consumption
 y_y : Income
 $\tau(y_y)$: Taxes on income
 $down_{it}$: Down payment
 δ_s : Appreciation factor e.g. (1+r)
 s_{it} : Savings (payment/expenditure to future self)
 W_t : Net household wealth process
 u_i : utility function
 θ_i : Education investment/consumption/cost
 V_{ht} : Value of the house
 ψ_h : fraction of home value required for down-payment.
 e_t : Education status Indicator
 σ_1 : Crime shock process
 R_{jit} : Transfers intra-household (from j->i)
 I_{ith} : Home-buying event Indicator
 H_{it} : Home-owning Indicator
 h_{it} : rent or mortgage payment

Chapter 1: Background and Context

Personal wealth is often viewed as a product of good decision-making. Choose a good education, a good job, earn a good income, invest wisely, take risks, start a business, innovate. However, the reality of wealth is quite different. Wealth is often passed down from generations through inheritance and in-vivo transfers. One's household's wealth is as much an expression of their opportunities, as it is the opportunities of their parents and their grandparents. This means that wealth is a way for the present to hang on to the past and to propel the future. For differences between Black and white households, wealth is a reflection of past inequity.

Longstanding inequities in the U.S. have worked against racial equity through myriad legal practices & exclusions: slavery, Jim Crow, the New Deal, FHA loans, the GI Bill's lack of Black inclusion, the *de jure* segregation of red lining, reverse red lining, racial covenants, and the persistent injustices of a criminal justice system that led to the death of George Floyd sparking a wave of protest and conversation (Rothstein 2017, Darity Jr and Mullen 2020). While no American alive bore firsthand witness to all of this history, most Americans have borne witness to some, and the accumulation of harms reverberates across the life course of all. The sociological, psychological and ethical costs of disparities and discrimination in access to education, health, jobs, housing, justice, and financing cannot be fully measured. There is however one indicator that reflects some portion of these inequitable costs: the Black-white wealth gap in American society today (Shapiro, Meschede et al. 2013, Fox 2016, Pfeffer 2018, Boshara 2020, Aladangady and Forde 2021, Toney and Robertson 2021). The wealth gap reflects income differences accumulated over time, and hence serve as a useful economic "stock" of disparities. While median white household income is \$68,000 and \$40,000 for Blacks, the total wealth of an average white family is *10 times* that of a Black family--\$170,000 v \$17,000 (Federal Reserve Board 2019).

This paper is motivated by addressing the history of unjust and, at times, racially motivated policy-making that is responsible for significant portions of the adverse racially disparate outcomes we can now measure. Given a long term policy challenge of identifying effective policies for addressing the wealth gap, this paper takes a first step in introducing a modeling environment that could be used for testing the long run impacts of policy interventions on differences between white and Black wealth.

The policy challenge of how to adequately close Black-white economic gaps is anything but simple. Some have proposed economic reparations for descendants of slavery, essentially a one-time wealth transfer to Black families (e.g., Coates (2015), Darity Jr (2021)). Others have proposed baby bonds—seed grants deposited by the government into investment accounts created for newborns in low-income families, possibly with additional yearly deposits (e.g., Hamilton and Darity Jr (2010)). While reparations are retrospective and aimed specifically at Blacks, baby bonds are prospective and technically race-neutral. But, given the higher rate of poverty of Blacks (20.8%, as opposed to 8.1% for whites), baby bonds would disproportionately go toward helping Blacks in accumulating wealth. They could also be targeted only at Black families. Another idea comes from a "rich uncle" policy that provides government guarantees for poor families when they borrow, essentially decreasing their collateral requirement and/or the interest rate or improving other lending terms. This replicates the process of intergenerational wealth transfers used by wealthier, often white households; these could also be targeted specifically at Black families. In addition to wealth-oriented policies, stricter enforcement of

anti-discrimination laws in hiring and lending, affirmative action, and better access to health and education could all decrease disparities.

The challenge is to turn these ideas into equitable, effective, and economically viable policies. To address this challenge, policymakers will need answers to key questions like: How do these candidate policies stack up against one another in reducing Black-white income and wealth gaps? What effect do they have on overall economic performance? How do the various costs and benefits involved accrue to different subgroups of Americans? In this paper, we develop and test a modeling foundation that could be used to create a unified and comprehensive framework requisite for answering these critical questions.

Specifically, we introduce a computational modeling approach that could be used to compare quantifiable outcomes of these proposed policies and provide actionable recommendations to policymakers and build on findings from the recent Black-white wealth gap discussion series from Edwards (2022), Paige (2022), Welburn, Lima et al. (2022). We anchor this model in a framework of economic stocks (net assets at $t=0$) and flows (annual income) and the relationships between the two given individual agent's employment and saving & borrowing decisions (including housing) from year to year for the simulation period.

Specifically, while Welburn, Lima et al. (2022) construct a model to analyze the impacts of a one-time transfer on the wealth gap within a single time period, this paper builds on widely used approaches in economics for modeling the economic trajectory of a dynasty across generations, an approach called overlapping generations or OLG models. At their simplest, OLG models include a young and an old household that overlap for a given period of time. In that time period (often assumed to be decades-long), the young can consume, earn income, and save while the old consumes its past savings before transferring its remaining wealth to the young. In the next period, the young household transitions to become the new 'old', and a new 'young' is born. While significant extensions have been made to OLG models which include simulations at annual time periods and the impact of taxes on inheritances, applications come short of modeling the processes (e.g., discrimination in the labor market and criminal justice) by which white wealth and Black wealth diverge. Despite their utility, standard OLG applications are limited in their ability to shed light on heterogenous populations with disparity in input parameters. Both heterogeneity in race and disparity of inputs including income of shocks are fundamental to identifying the differences in wealth outcomes we seek.

To understand how policy can shape the wealth gap, this paper extends the OLG literature through the introduction of a transfer-based, heterogenous extension to overlapping generations model which we call the Race, Economic Disparities and GENERational wealth (REDGEN) model. The REDGEN model is an empirical equilibrium model which makes novel use of computational methods from reinforcement learning (RL) and agent-based modeling (ABM) to capture the differences in incomes drawn from sample distributions portioned by race, as well as race-specific shocks, also drawn from distributions. Specifically, RL techniques provide significant computational improvements for solving dynamic programs within the OLG context while ABMs strengthen the ability to include the requisite realism of disparity within population dynamics.

The REDGEN model includes dynastic transfers, in-vivo, and inheritance, between overlapping young, middle, and old generations with annual time steps. Our objective in this paper is to enable the future study of the effect of novel intervention policies—such as reparations and baby bonds—for which it would be hard to find existing data and conduct econometric and other

statistical analyses. In comparison to standard OLG models which become computational expensive at the high dimensional state space, the combined use of RL and ABM¹ provides a more tractable way to study the evolution of income, household wealth, and intergenerational transfer over time among highly heterogeneous agents whose behavior is governed both by individual characteristics (e.g., race, inherited wealth, education level) as well as the policy environment (e.g., stronger or weaker redlining policies in real estate and banking). We aim to model spot markets such as labor and housing (and other forms of wealth) in equilibrium, for consistency with traditional Keynesian models while allowing the ABM to evolve dynamically in a non-equilibrium fashion.

¹ ABMs are simulations of agent behavior. One key difference between ABMs and standard equilibrium models (e.g., OLG) is that ABMS allow for interactions between agents.

Chapter 2: Outline of the Model's Foundations

Our work builds on a foundation of the overlapping generations (OLG) model described in the previous chapter. This chapter outlines the key aspects of the central OLG that we then extend into an ABM model aiming to overcome limitations to modeling heterogeneity within standard OLG models for this policy issue. The resulting ABM model enables more detailed analyses and experimentation on the dynamics of generational wealth accumulation. The baseline OLG model specification comes next.

OLG Model Structure

Our OLG basic structure consists of three concurrent generation (young, middle, old) making decisions around consumption, saving, intergenerational transfers, as well as a set of life-stage conditional decisions that exert influence on income and wealth.

Specifying Generations and their Constraints

Young

At time t , the young generation, g , consumes $c_{g,t}$, chooses whether to attend post-secondary education $e_{g,t} = \{0 \text{ if no, } 1 \text{ if yes}\}$ at cost $\theta_{g,t}$, pays housing rent $h_{g,t}$ saves $s_{g,t}$ and pays a net intergenerational transfer (i.e., the difference $R_{gj,t} - R_{jg,t}$). Young earns income from an interest earning endowment received in the prior period, $\omega_{g,t-1}$, and pre-tax labor income $y_{g,t}$ where $\tau(y_{g,t})$ is a tax rate function leading to the following budget constraint:

$$c_{g,t} + e_{g,t}\theta_{g,t} + \tilde{h}_{g,t} + s_{g,t} + \sum_{j \neq g} (R_{gj,t} - R_{jg,t}) \leq \omega_{g,t} + \tilde{y}_{g,t} (1 - \tau(y_{g,t})). \quad (1)$$

Middle

At time t , the middle generation, $g - 1$, consumes $c_{g-1,t}$, chooses whether to buy a house $I_{g-1,t} = \{0 \text{ if no, } 1 \text{ if yes}\}$, with down payment $\eta_{g-1,t}$, pays housing rent $h_{g,t}$ if they do not own a home (defined by the binary indicator $H_{g-1,t} = \{0 \text{ if no, } 1 \text{ if yes}\}$), saves $s_{g-1,t}$ and pays a net intergenerational transfer. Middle earns income from interest earning savings from the prior period, $s_{g-1,t-1}$, and pre-tax labor income $y_{g-1,t}$ as follows

$$c_{g-1,t} + I_{g-1,t}\eta_{g-1,t} + h_{g-1,t} + s_{g-1,t} + \sum_{j \neq g-1} (R_{g-1j,t} - R_{jg-1,t}) \leq (1 + r)s_{g-1,t-1} + y_{g-1,t}(e) (1 - \tau(y_{g-1,t})) \quad (2)$$

where the down payment is an exogenously imposed percentage ψ_i of home value V_t

$$\eta_{g,t} = \psi_i V_t \quad (3)$$

for a given household i and home value V_t a function of future rents and interest as follows:

$$V_t = \kappa \sum_{n=t}^{\infty} \delta^n h_n, \delta = \frac{1}{1+r}; t \rightarrow \infty \Rightarrow V_t = \kappa \frac{h_t}{r} \quad (4)$$

Old

Finally, at time t , the old generation, $g - 2$, consumes $c_{g-2,t}$, pays housing rent $h_{g,t}$ if they do not own, and pays a net intergenerational transfer. Old earns income from interest earning savings from the prior period, $s_{g21,t-1}$, and pre-tax income $y_{g-2,t}$ (i.e., pension, social security) as follows

$$\begin{aligned} & c_{g-2,t} + (1 - H_{g-2,t})h_{g-2,t} + \sum_{j \neq g-2} (R_{g-2j,t} - R_{jg-2,t}) \\ & \leq y_{g-2,t}(e) (1 - \tau(y_{g-2,t})) + (1+r)s_{g-2,t-1} + H_{g-2,t}V_t \end{aligned} \quad (5)$$

In each period for each generation, the stochastic income stream is a function of the median income for generation g at period t , μ_{gt} , an income shock, σ_{gt} , conditional on education, e_g and a set of demographic characteristics, D_g , (e.g., race, location) defined as follows where σ_t is a shock

$$y_{g,t} = F_y(\mu_{gt}, \sigma_{gt} | e_g, D_g). \quad (6)$$

Furthermore, wealth is defined as follows

$$W_{g,t} = s_{g,t} + H_{g,t}V_t. \quad (7)$$

Finally, we introduce a home ownership constraint which states that a given household is a homeowner ($H_{g,t} = 1$) if the household chooses to buy a house ($I_{g,t} = 1$) in the current period t or prior as follows:

$$H_{g,t} \geq \epsilon \sum_{t=1..t} I_{g,t}, \quad H_{it} \in \{0,1\} \quad (8)$$

Within each period each household has Constant Relative Risk Aversion (CRRA) preferences such that

$$u_{g,t}(c_{g,t}) = \beta c_{g,t}^{1-\rho} \quad (9)$$

and agents exhibit warm glow preferences in their utilities


$$U_{g,t} = \sum_{t=g}^{g+N} \left(u_{g,t}(c_{g,t}) + \phi_1(R_{gg+1,t}) + \phi_2(R_{gg-1,t}) \right) \quad (10)$$

where $\alpha \in [0,1]$ is a coefficient of altruism and $\phi_1, \phi_2 \in [0,1]$ are weights reflecting preference towards generations $gg + 1$ and $gg - 1$.

Overlapping Generations in a Wealth Agent-Based Model

We extend the OLG approach by employing an agent-based model (ABM) in a reinforcement learning environment with numerical results. In the model, households follow a birth-death process which, while modeled using annual time steps, can be segmented into three 20-year epochs that organize major financial decisions as represented in Figure 3.1 according to the traditional young, middle, and old stages of life in OLG models. For simplicity we assume all new households that follow the initial generation enter at age 20 with high school education and renter status. During the young epoch (ages 20-39), households choose whether to pursue a college degree to receive a stochastic positive level change in earnings. In the middle epoch (ages 40-59), households choose to buy housing or rent in perpetuity. In the old and final epoch (age 60-79), households are assumed to retire, live off savings, and leave a bequest to descendants upon death at age 80. Households are assigned to a demographic group D representing race which is constant across all epochs.

Figure 3.1: Schematic of ABM Household Life Course



After a household is born and enters the young epoch, the default state is regular employment. In this state, households earn gross income starting at a stochastic base wage drawn from a log-normal distribution conditional on race, $Y_{gt} \sim LN(\mu_{gt}, \sigma_{gt} | D)$. Unemployment is subsumed within the stochastic income where periods of unemployment result in fewer days worked and lower incomes in each period t . Annual gross incomes get a static tax rate τ applied to them. Households also spend a stochastic portion of annual wages on housing h . This rent expense is stochastically drawn from a parameterized log-normal distribution but also upper-bounded by disposable income. Wages grow annually at a rate parameterized by each 20-year epoch, with the greatest growth in young years and declining in advanced years. Households in the young epoch who pursue college are assumed to have depressed earnings while enrolled in college in addition to tuition payments. After graduating, they receive a new base wage that is the minimum of their prior wage and a random draw from an educated worker income distribution, $\min[y_{old}, Y_{ed}]$.

Thus, for each period t after birth, households are in one of four states – employed, incarcerated, in school, or sick/injured – shown in Figure 3.2. A household's stochastic income stream is conditional on its state, and the transition probabilities, shown through arrows Figure 3.2 are parameterized by race. We assume the probability of transitioning to school is zero in the middle and later epochs.

However, households may also experience shocks in which they transition to a different state within a given period. Guided by first-order impacts to financial well-being and racial inequity, we introduce two types of idiosyncratic shocks: incarceration and medical. Incarceration

temporarily removes a household from the workforce and interrupts consumption for one or more periods. In addition to missed earnings, individuals who have ever been incarcerated receive a one-time negative level change to base annual wages representing both lower hourly rates and a reduction in average days of employment. Medical events encompass serious illness or injuries requiring hospitalization and have the effect of interrupting earnings for the year while other expenses continue.

Figure 3.2: Finite State Depiction of ABM Household States

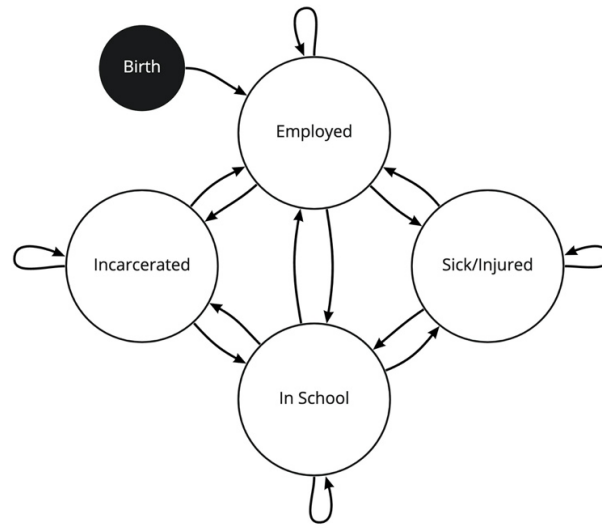
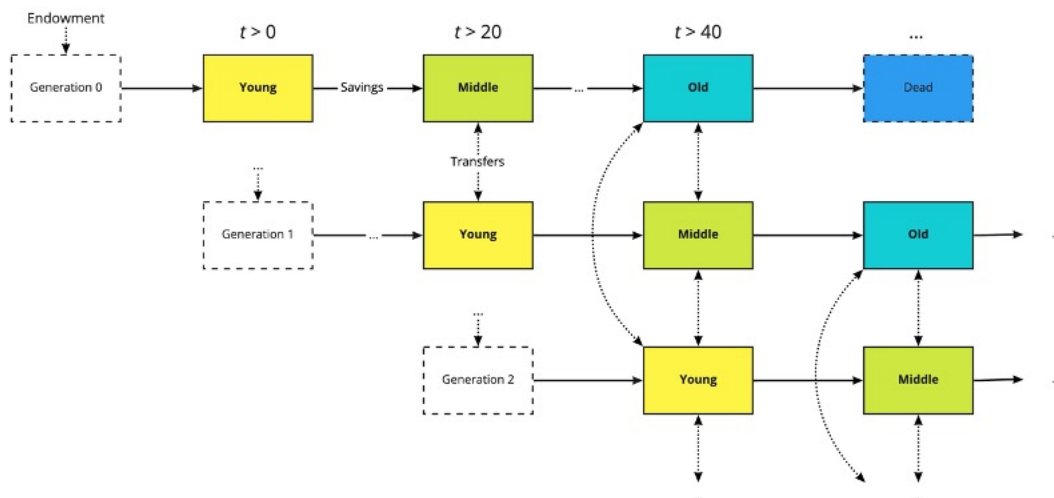


Figure 3.3: Unrolled Depiction of ABM Households Evolving in Time



Upon entering the middle epoch, a new young agent is generated with an initial endowment $\omega_{g,t} \in (0, \infty)$ and linked to the parent, forming a dynasty as illustrated in Figure 3.3 (note: we assume that the endowment given at birth is strictly positive while wealth in future periods can be negative). Each period agents allocate current liquid wealth between savings for future consumption, *in vivo* transfers to other active generations in the dynasty, and in the final period a

bequest to be divided between descendants (included in transfers in the figure). Accumulated savings grow with a fixed savings rate defined by r_s and so evolves according to:

$$s_{t+1} = (1 + r_s)(s_t + \tilde{s}_{new,t}) \quad (11)$$

If agents own their home, we assume the asset remains non-liquid until the agent's death and appreciates following a simple geometric Brownian process with drift equal to a constant real-estate interest rate and a small variance:

$$V_{t+1} \sim GBM(\mu = r_e, \sigma^2 \ll 1) \quad (12)$$

Net worth W_t at the start of any period is calculated as the sum of savings and home value $s_t + V_t$. Note that agents are allowed to carry negative savings. For simplicity we assume a symmetric interest rate on positive and negative holdings. At death, a household's wealth is liquidated and distributed to the remaining currently living members of the dynasty in a final statically specified bequest action. Future experiments may model this action endogenously and adaptively to identify estate-management behaviors and policies that are influential on wealth disparities.

Household's Financial Decision-making Behaviors

In each time step, households must make constrained financial decisions on consumptions, savings, *in vivo* transfers, education, housing. The constraints are largely due to current or expected near-term income for individual households as well as within a currently active dynasty. The ABM is designed to be modular with respect to the actual model of financial behavior that agents use to navigate these annual constrained decisions. Behavioral models select decisions based on solutions to constrained optimization programs running at the level of individual agents.

Two behavioral models are available for experiments: a basic financial decisionmaking model and a reinforcement learning (RL) based model. We detail each below.

1. Basic Financial Decisionmaking Model

The basic model solves an instantaneous static optimization program. It identifies consumption, savings, and in vivo transfers that maximize the agent's linear warm-glow-augmented utility function:

$$U_{g,t} = u_{g,t}(c_{g,t}, s_{g,t}) + \phi(R_{gg+1,t} + R_{gg-1,t}) \quad (13)$$

where $u_{g,t}(c_{g,t}, s_{g,t})$ is the Cobb-Douglas utility $c^\alpha \times s^\beta$ with parameters set around the values of the average propensities for consumption and saving.

The most salient restriction or limitation of this approach is the optimization is restricted to the current period. It does not account for inter-temporal influences except indirectly via utility from current saving for future consumption. Furthermore, the solutions under this formulation of household utility are likely to statically account for real variability in financial decisionmaking. We use this decisionmaking model primarily to provide a baseline for development.

2. RL-based Financial Decisionmaking Model

We advance upon the basic decisionmaking model through the use of Q-learning, often used within the reinforcement learning literature, to solve the Bellman equation for the Markov Decision Process (MDP) specified by our ABM. This results in an agential policy function that enables each agent to endogenously decide on consumption, savings, transfers, as well as education and house-purchasing investments. To elaborate on this approach, we introduce the following notation:

Table 3.1: Mathematical Notation for Model Specification.

Notation	Meaning
$s_t \in \mathcal{S}$	A household's state description. Includes all relevant household variables (demographics, income, savings, housing equity, education status, etc.)
$a_t \in \mathcal{A}$	A household's vector of decisions or actions made at time t
$(s, a) \in \mathcal{S} \times \mathcal{A}$	Sample decision scenarios (s, a) in the product space $\mathcal{S} \times \mathcal{A}$
$u = \mathcal{R}(s_t, a_t) \in R^+$	Rewards or utility to the household for responding to state s_t with action a_t .
$\mathcal{R}^*(s_t)$	Best possible reward achievable in state s_t with the best possible response e.g., $\operatorname{argmax}_{a_t \in \mathcal{A}} \mathcal{R}(s_t, a_t)$.
$Q(s_t, a_t)$	Q-function. The true utility value of an action given a current state
$\pi(s): \mathcal{S} \rightarrow \mathcal{A}$	Household's response policy function i.e., a function that captures the agent's financial behavior or strategy for responding to any state $s \in \mathcal{S}$

The Q-function is a pivotal mathematical entity for learning optimal control in an MDP. We can express the Q-function as a Bellman Equation

$$Q^*(s_t, a_t) = u(s_t, a_t) + \gamma \cdot \max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1}). \quad (14)$$

There are several approaches to solving this Bellman equation in various traditions. The AI/ML/RL tradition has a solution based on function approximations for Q. We model the Q-function as a neural network approximator which we tune iteratively by applying learning algorithms to experience samples (s, a, u) extracted by interacting or exploring the ABM. One effective learning algorithm is called SARSA (for "on-policy" control):

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [u_{t+1} + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)] \quad (15)$$

The learning process “bootstraps” itself to increasingly better approximations of Q^* using sample outcomes from repeated interactions with the environment. Once properly trained, we derive an optimal policy based on the trained Q-function as:

$$\pi(s_t) = \operatorname{argmax}_{a \in \mathcal{A}} Q(s_t, a_t). \quad (16)$$

We train separate baseline policy models for Black and white households to capture the fact that both kinds of households are interacting with very different financial environments based on the different structural parameters facing the different households.

Model Parameterization and Calibration Strategy

The model is equipped with diverse collection of structural parameters implemented as described in the preceding section. The value of the ABM for analysis and policy experimentation depends crucially on how well the model represents and/or abstracts key dynamics in the real-world. We address this representativeness goal with two strategies:

1. **Empirical Structural Parameterization:** we refer to a variety of economic data to find empirically-observed estimates for structural model parameters that have clear real-world analogues.
2. **Partially Supervised Model Calibration:** We treat model parameters that have no *clearly-observable* real-world analog as degrees of freedom which we fit under a supervised learning approach. This is essentially equivalent to solving an optimization problem on an objective defined in terms of how well candidate models fit real outcomes on observable dimensions.

We address each path in detail below.

Structural Parameterization

The model uses data and parameter estimates primarily from two sources. The Survey of Consumer Finances (SCF) is a survey of U.S. households conducted with a nationally representative cross section of 4,500 to 6,500 families every three years and includes a detailed hierarchy of assets and debt balances (Federal Reserve Board 2019). The Panel Study of Income Dynamics (PSID) covers a variety of income, expenditure, and activity items from a nationally representative panel of families spanning up to several generations from 1968 to the present (Johnson et al. 2018). Additional parameters are drawn from Federal Reserve Economic Data (FRED), the Bureau of Justice Statistics (BJS), and previous literature.

There are several cautions in linking these datasets. The SCF is a repeated cross-section that oversamples high net worth individuals. The PSID by contrast was designed to oversample low-income households and is a panel design. However, comparisons have shown that both similarly reflect macroeconomic aggregates and when properly weighted and align closely, particularly for the bottom 95th percentile (Pfeffer et al. 2016; Bosworth and Smart 2009) which is our main demographic of interest. We also leverage different strengths of the data, as SCF is referenced for initial account balances and distributional benchmarks, while we leverage PSID to estimate flows and shocks over time. Because of coding differences across and within datasets, throughout the model and in supporting analysis we distinguish between Black and non-Black households. Since the latter includes other historically disadvantaged groups, differences in parameters and outcomes are likely conservative approximations of the Black-white wealth gap. We also where applicable use data attributed to the head of household or reference person (RP).

The model is initialized with a synthetic population drawn from a weighted SCF sample using provided replication weights. State variables for the initial synthetic population are recoded from

SCF responses including age, race, income, education, housing information, and savings, as detailed in Table 3.2. We approximate liquid savings by taking the value of all financial assets, which includes bank accounts, stocks, bonds, certificates of deposit, and other similar instruments, and subtract outstanding installation loans and credit card balances. Years of education is not included in SCF responses so random increments are generated for observations reported as “some college.”

Table 3.2: Recoding of SCF variables for synthetic population

Variable	Domain	SCF	Recode Procedure
Age	$(0, \infty)$	age	Direct value
Race	$\{0,1\}$	race	1 if “Black/African American”, 0 otherwise
Wage	$(0, \infty)$	wageinc	Direct value
Educated	$\{0,1\}$	edcl	1 if “college degree”, 0 otherwise
Education Years	$(0,4)$	edcl	4 if “college degree”, ~Uniform(1,3) if “some college”, 0 o.w.
Tenure	$\{0,1\}$	housecl	1 if “owns ../house/condo/coop/etc.”, 0 otherwise
House Value	$(0, \infty)$	houses	Direct value
House Equity	$(-\infty, \infty)$	houses, nh_mort	houses - nh_mort
Savings	$(-\infty, \infty)$	fin, install, ccbal	fin - install - ccbal

Most model parameters as well as values for newly created agents are derived from PSID, summarized in Table 3.3. Using Stata and the package *psidtools* (Kohler, 2015) we constructed an unbalanced panel of individuals indicated as the reference person or head of household from years 1989 to 2019. Because of the complex sample design combining three stratified probability samples, variables are provided for sampling error stratum and clusters, as well as both cross sectional and longitudinal household weights, which are used as appropriate to each analysis. All monetary variables are adjusted to 2020 dollars using the “Consumer Price Index for All Urban Consumers” calculated by the St. Louis Federal Reserve.

Average initial income \bar{y} and the effects of education and age are predicted by regressing log labor income of “earners,” defined as observations where RPs reported as both employed and with positive income, on degree, marital status, and age group, with individual, region, and year fixed effects. Wage growth rates r_y are predicted by regressing the change in income $\log y_{it} - \log y_{it-1}$, requiring observations with at least two consecutive earning periods, on race, degree, age group, and a quadratic function of income with year fixed effects. Housing costs h are calculated as the weighted average of reported rent as a proportion of income. As a proxy for the probability of medical shocks π_m we calculate the weighted average by race of individuals who reported overnight hospital stays in a year. There is no directly reported costs of individual medical events in PSID, but total spending on hospitals and nursing homes was included since 1999. We estimate the average cost of medical shocks ξ by regressing hospital costs on an indicator for hospitalization, sex, marital status, income, and individual and year fixed effects. The average home appreciation rate r_e is calculated as the longitudinally weighted average change in reported home value as a percentage $\log V_t - \log V_{t-1}$ for Black and non-Black households.

Savings in the model are assumed to represent a portfolio of investments with a composite rate of return r_s , for which we use the geometric average of real annual² returns reported in FRED from 2010 to 2020 for the Wilshire 5000, money market funds, and non-jumbo savings accounts. The probability π_i and length λ of incarceration are derived from BJS as the number of admissions to federal prison as a percentage of total population and average sentence length respectively (a known underestimate), segmented by race. Since this does not include state or local prison admissions, π_i is adjusted to reflect cumulative lifetime risk in line with other studies (Western et al. 2021; Bonczar and Beck 1997). A one-time shock to base wage as a percentage of current income ϵ is based on a Pew Charitable Trusts report and incorporates both change in wage rate and reduction in days of employment for individuals who were ever incarcerated.

Table 3.3: Model Parameters

Parameter	Description	f	Non-Black	Black
\bar{y}	Average initial log income with no degree	~Lognormal	10.58 [‡]	10.18 [‡]
\bar{y}_e	Average initial log income with degree	~Lognormal	10.86 [‡]	10.52 [‡]
ϵ_e	Wage shock from degree completion	Constant	0.18 [‡]	0.19 [‡]
$r_{y,young}$	Annual base wage growth rate in young epoch	Constant	0.045 [‡]	0.045 [‡]
$r_{y,middle}$	Annual base wage growth rate in middle epoch	Constant	0.00 [‡]	-0.00 [‡]
$r_{y,old}$	Annual base retirement income growth rate in old epoch	Constant	-0.10 [‡]	-0.10 [‡]
h	Share of income spent on rent or mortgage	Constant	0.27 [‡]	0.27 [‡]
r_s	Annual savings (real) rate of return.	Constant	0.0268 [†]	0.0268 [†]
\bar{r}_e	Average annual real estate rate of return.	~Normal	0.021 [‡]	0.008 [‡]
π_i	Annual probability of incarceration	Constant	0.002 ^{††}	0.005 ^{††}
λ	Average incarceration length	Poisson	4 ^{††}	6 ^{††}
ϵ_p	Wage shock from incarceration as percentage	Constant	-0.52 [×]	-0.44 [×]
π_m	Annual probability of medical event	Constant	0.11 [‡]	0.13 [‡]
ξ	Direct cost of medical event	Constant	819 [‡]	819 [‡]

[†] Weighted mix of Wilshire 5,000, money market, and savings interest rate returns reported in FRED

[‡] Author's analysis of the Panel Study of Income Dynamics

^{††} Bureau of Justice Statistics admissions to federal prisons

[×] The Pew Charitable Trusts. 2010. "Collateral Costs: Incarceration's Effect on Economic Mobility," 1-40.

Partially-Supervised Calibration

First, some notation to frame the calibration approach. The ABM may be conceived as a dynamic function, \mathcal{W}_t , mapping a high-dimensional vector of parameters, \vec{v} , into model output dimensions that vary over simulation time.

$$\mathcal{W}_t(\vec{v}): T \times \mathbb{R}^p \rightarrow \mathbb{R}^d \quad (17)$$

² We follow the principles outlined Spizman (2008) in using the geometric mean for average growth rates.

\mathcal{W}_t is intended to approximate the dynamics/evolution of a real-world process $\widetilde{\mathcal{W}}_t$ for the same output dimensions. Given this basic framing, we treat the calibration process as an optimization on functions of \mathcal{W}_t and $\widetilde{\mathcal{W}}_t$ over sub-dimensions of the parameter space. An implicit assumption here is that the best parameter for imitating real-world dynamics is the observed value of that parameter in the real-world. Thus, we mainly restrict the optimization process to dimensions of parameter space corresponding to unobserved or unmeasured parameters.

The optimization program for the calibration process may be written as:

$$\arg \min_{\vec{v}} d(\mathcal{W}_t(\vec{v})|\widetilde{\mathcal{W}}_t) \quad (18)$$

where $d(\mathcal{W}_t(\vec{v})|\widetilde{\mathcal{W}}_t)$ is a comparison or distance function estimating how similar the ABM simulation traces are to the real-world traces on selected output dimensions for a given parameter value. Solving this program is a supervised learning problem since we are comparing and selecting arguments based on how well the models they parameterize match with a source of ground truth (the real world) on specified outcomes.

The implementation of the optimization program as well as the choice of distance functions are design choices. Our initial work applies the following comparison function to the quartiles of the wealth distributions (disaggregated by race) at the terminal date of the simulation (2019):

$$d(\mathcal{W}_t(\vec{v})|\widetilde{\mathcal{W}}_t) = \log \left(\sum_{q \text{ in quartiles}, n \text{ in race}} \frac{(\mathcal{W}_{t|q,r}(\vec{v}) - \widetilde{\mathcal{W}}_{t|q,r})^2}{\widetilde{\mathcal{W}}_{t|q,r}} \right). \quad (19)$$

This equation is the logarithm of the Chi-squared goodness-of-fit statistic on the terminal wealth distributions for the simulated and real populations, after disaggregating by both race & wealth quartiles. The logarithm helps control the magnitude of the statistic in a smooth fashion given that the wealth samples have a high dynamic-range. The initial implementation of the optimization program involves grid-search over hypercubes of the parameter space. We use the best set of parameters found under this approach subject to given budget on computation time.

Computational Structure and Details for the Model

This chapter ends with some discussion of computational details for the implemented model. The model and all its ancillary calibration and experimental structures are implemented in the Python language. Python is an open-source language. So that implementation choice is less likely to impose software licensing barriers on potential adoptees.

This open-source language choice also has strong implications for ease of doing calibration and experiment runs. In single run mode, the model is not egregiously intensive to run (depending on the population size and financial behavioral model in use). But the model requires multiple Monte Carlo runs to produce useful samples for calibration as well as for characterizing effect sizes in any synthetic experiments. Doing sequential Monte Carlo runs quickly inflate that computational budget even if the base model was perfectly efficient. Monte Carlo runs are eminently parallelizable. And parallelized computation is easiest to setup and run on scalable

compute resources (e.g. Amazon Web Services) when the base model is in an open-source language like Python.

The model does make use of third-party open-source libraries like Jax (for adaptive behavioral modeling) and Pandas (for data management). These dependencies may impose some brittleness on the implementation depending on the evolution and continued maintenance of those libraries.

Chapter 4: Discussion of Some Experimental Results

We implemented the model described in the preceding chapters to enable experimentation for policy insight. This chapter discusses some of our early findings from the implemented model. The first discussion below provides important context for parsing any results from the ABM. After that we present some model sensitivity findings to give insight into how much influence different structural parameters have on wealth propagation dynamics as represented by our ABM.

Experimental Setup

The ABM is a *stochastic* simulation of *approximated* (or abstracted) real-world dynamics. Therefore, there is some appropriate process and structuring needed to perform experiments and interpret findings in this synthetic policy sandbox. It is useful to identify that there are at least two kinds of errors we need to control for when performing experiments:

- **Structural Errors:** These are errors due to improperly approximated dynamics and influences. The calibration process, the use of real-world statistical relationships, & the careful specification of agent behavioral models aim to control these errors as the preceding chapters outline. The misspecification of behavioral models is a key influential factor determining structural errors. Mis-specified behaviors are often also the hardest correct because faithful quantitative modeling of adaptive behaviors without anchoring constraints is challenging³.
- **Estimation Errors:** These are errors due to inherent stochasticity in the model. Even if the model is perfectly representative of real-world dynamics & interaction, the simulations are being driven by stochasticity from the population sample and randomized model parameters (e.g., instantaneous rent, wage growth, etc.). So, the simulation outputs trace a random sample path of unknown probabilistic distribution. We control for estimation errors by using Monte Carlo procedures when estimating simulation outcomes.

Monte Carlo estimation comes with theoretical guarantees on how the precision of estimates increase.⁴ So we can get decently precise estimates of average outcomes, provided we can pay the sampling & computational cost. And with modern scalable computing infrastructure, these costs are largely directly translated into monetary costs.

All results of experiments (sweeps) on this ABM are, thus, averaging over multiple sample-runs to address inherent model stochasticity. We will present estimated confidence bars for results as well to help communicate outputs clearly considering this inherent stochasticity.

Observed Sensitivities to Structural Parameters

A key goal for constructing this model was to produce a representative sandbox model of wealth dynamics that can be used in future research to test hypothetical policies and counterfactual

³ Davis, Paul K., Angela O'Mahony, Timothy R. Gulden, Osonde A. Osoba, and Katharine Sieck. *Priority challenges for social and behavioral research and its modeling*. RAND Corporation Santa Monica United States, 2018; Osoba, Osonde, and Paul K. Davis. "An artificial intelligence/machine learning perspective on social simulation: New data and new challenges." *Social-Behavioral Modeling for Complex Systems* (2019): 443-476.

⁴ Every n -fold increase in the Monte Carlo sample set yields a \sqrt{n} -fold increase in the precision of the estimator.

scenarios. This paper presents a model that is not fully calibrated and leaves a full calibration as a subject for future research. But, given the careful specification of the model structure (see parametrization discussion in the preceding chapter), the model in this paper can provide useful policy insight into the relative influence of various structural parameters on the wealth gap.

To extract information on relative parameter influence from the model, we run the ABM holding constant all structural parameters, except the parameter under investigation. This goes on for multiple runs in parallel. Each run provides an estimate of the stochastic outcome measure (the wealth gap) for the specific value of the target parameter. We then sweep the target parameter over a range of values in a race-neutral manner (same parameter value). The aim is a computational estimate of the association between changes in the input target parameter and changes in the output measure, a very crude estimate of a derivative function. Parameters that induce higher variability in output measures may be said to have more influence on the wealth gap (or any other selected outcome measure in the model), all other factors held constant. The figures below plot examples of the influence of various parameters on mean net worth disaggregated by race (Black vs. non-Black).

The first figure, Figure 4.1, depicts the influence of our average rent factor on mean net worth. It shows that as average rent increases, the wealth gap widens. While the net worth of both groups declines in absolute terms in response to higher rents, the reduction is far larger for Black. The gap is much smaller at lower average rent levels and the influence plot exhibits a significant phase change. This figure highlights the potentially distortionary impact of inflationary pressures where increases in core price levels can expand the wealth gap. This is mostly due to the interest on debt as households go into debt to pay housing costs.

Figure 4.1: Estimated Influence of variation in Average Rent on the Black-<->Non-Black Wealth Gap

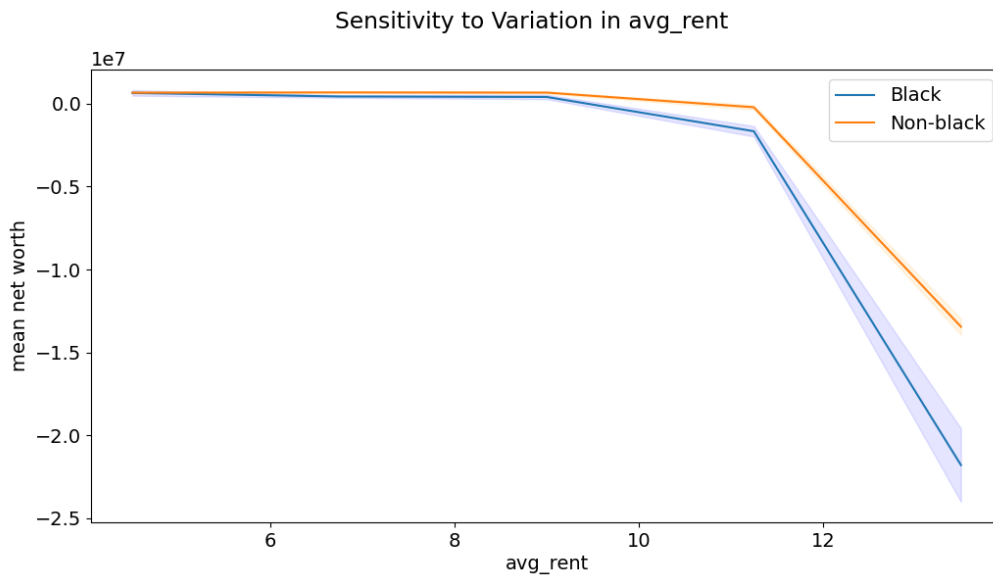


Figure 4.2 depicts the influence of the education premium factor on mean net worth. Recall that this premium captures the average positive wage shock that the status of being “sufficiently

educated”⁵ imparts to a wage earner. The figure shows that variation in that factor within our model has little to no influence on average net worth. One rationalizing intuition for this pattern is that any average change in agents’ education choices would occur non-differentially across race since everyone is responding to the same incentive and all “sufficiently educated” agents compete in a common pool for jobs. It is worth noting that our analysis of PSID indicates a slight differential in this wage education premium factor, with the Black demographic experiencing a slightly higher premium for getting educated. Further analysis is needed to examine how such a differential might influence the wealth gap in contrast to the nil influence estimated under a non-differential premium factor.

Figure 4.2: Estimated Influence of Variation in the Education Premium on the Racial Wealth Gap

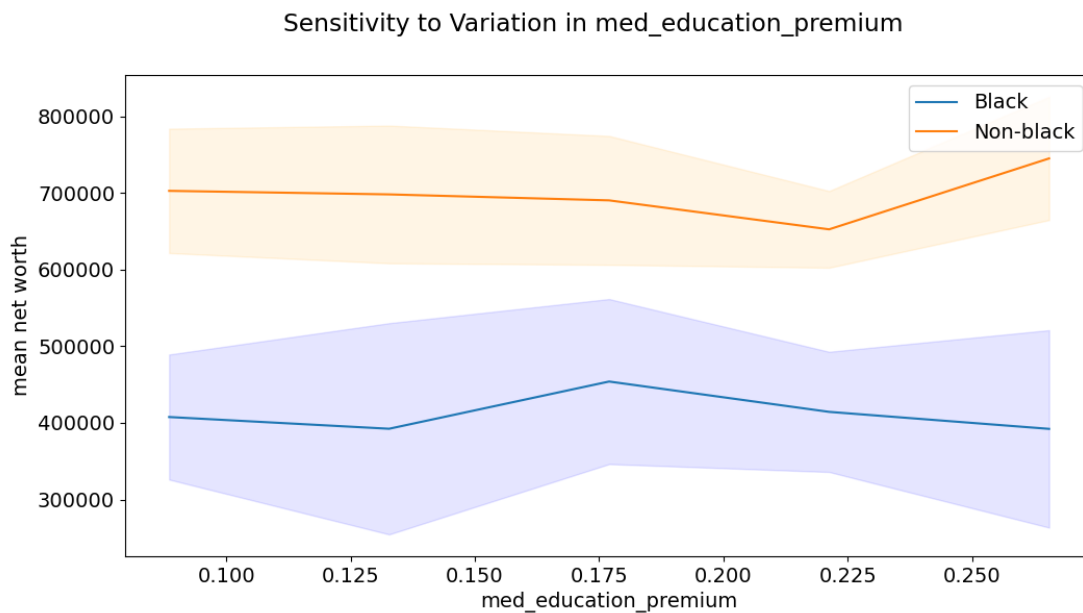
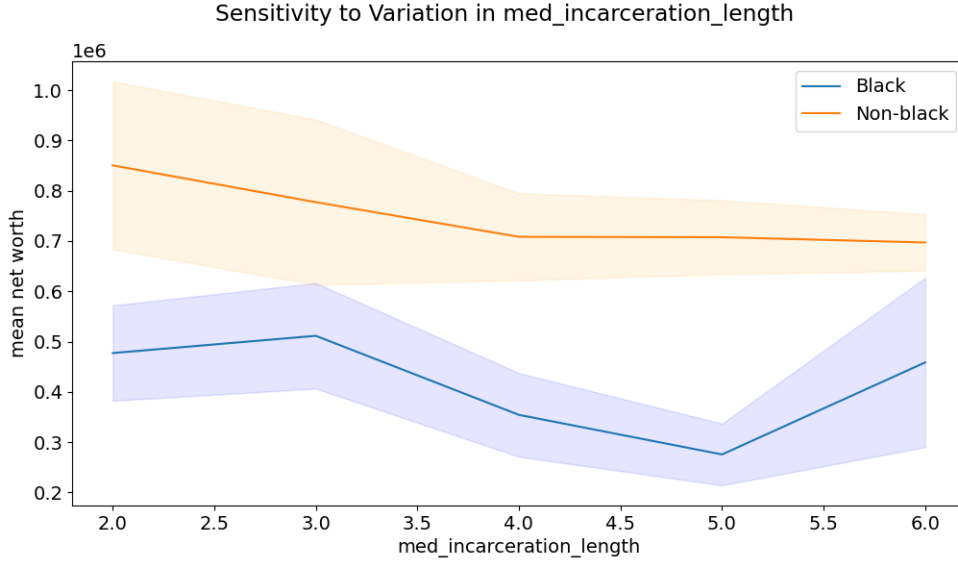


Figure 4.3 explores the influence of idiosyncratic criminal justice shocks (rate and intensity) on the wealth gap. Figure 4.3 highlights the influence of median incarceration length on the wealth-gap. Longer sentences seem to be associated with a slight widening of the wealth gap. Although it also has a depressing effect on net worth across the board. Further analyses are needed to characterize this effect in more depth.

⁵ We use the threshold of four or more years of college education to determine this parameter from PSID statistics.

Figure 4.3: Estimated Influence of Variation in Median Incarceration Length on the Racial Wealth Gap



Chapter 5: Future Work

This work builds on a long legacy of OLG models used for policy experimentation within economics. Specifically, this paper introduces and implements a novel computational model that leverages reinforcement learning to provide the requisite levels of heterogeneity to examine how the racial wealth gap responds to various scenarios. Of course, this model can and should be expanded by future work. Specifically, for robust and defensible policy experimentation, further calibration work using the approaches described in this paper is likely necessary. However, before this step, the model introduced within this paper provides a methodological advancement through its novel application of advanced computational tools in an OLG setting while stimulating useful policy insight. In this final chapter, we discuss the topic of future work enabled by our model. We primarily focus on experiments that can uncover key interventions and trade-offs for policies seeking to address the wealth gap.

Future Experiments

There are two initial experiments we aim to perform. These are motivated and informed by current policy conversations (e.g., by Darity et al. 2021) as well as the policy conversations around student debt). There are also other policy experiments that are less salient in current policy conversations whose outcomes seem particularly interesting based on our review of important dynamics that determine the wealth gap.

In order of priority, we discuss several policy experiments that could be performed by future work using a calibrated version of the model in this paper.

“Baby Bonds” aka Endowments at Birth:

The first experiment will evaluate the effectiveness and impact of policy interventions that provide agents with invested endowments (or baby bonds) at birth. The experiment can have a

few dimensions. For example, endowments may be unrestricted in use or restricted to apply to specific expenses. They may be disbursed in a race-dependent or race-neutral manner. They may also be disbursed in a mean-tested or unconditional manner. These salient experiment dimensions already specify 8 distinct kinds of experiment that are readily implementable in the model since we can simulate any intervention that is conditional on recorded household characteristics in the model. As an early hypothesis on the effect, one might expect the expected effects of such endowments to be muted if they are not sufficiently large to outweigh the cost of gaining “sufficient education” either at birth or when the investment is spent

Education subsidies

Following up on the pivotal influence of sufficient education, another experiment could explore the relative influence of race-based as well as mean-tested subsidies for education expenses. This experiment could evaluate how interventions that blunt the impact of runaway tertiary education cost may affect the wealth gap. Again, as an early hypothesis on expected effects based on the analysis within this paper, one might expect that even a race-neutral educational subsidy would improve the long-term financial well-being of Black students relative to non-Black peers reducing the wealth gap. This expectation is mostly informed by the disproportionate representation of Black citizens in the lower percentiles of wealth. This expectation is also informed by the relative importance of positive wage and income shocks in the lower wealth percentiles (compared to the higher relevance of capital gains and other non-wage income at higher wealth percentiles).

The Importance of Idiosyncratic Shocks and Insurance Schemes

Negative idiosyncratic income or wealth shocks may seem negligible on a large fraction of individual wealth paths. But when considering whole populations, the shocks can have rather strong inhibitory effects on average wealth outcomes. This is due to spiky nature of the costs imposed on households undergoing these shocks. We propose experiments that either limit the intensity of these shocks (e.g., by controlling the costs of medical events, reducing excessive bail, modifying minimum sentencing guidelines), controlling their occurrence rates (likely only possible for criminal justice shocks), or by implementing efficient insurance schemes that effectively smooth out the impact of such shocks (e.g., an efficient or well-subsidized universal healthcare scheme).

The alternative angle on shock impact is the forecasted effect of climate change on real estate markets and the knock-on effects on the wealth-gap. This is particularly salient to the USA because of the *cultural trend* of accumulating and locking-up wealth in real estate which can be devastated more readily by changing climate patterns.

Local Real Estate Dynamics

The history of segregation and its effects on education outcomes suggest that a better accounting of the wealth gap will need to account for micro-local housing choices and patterns (Shapiro, Meschede et al. 2013, Rothstein 2015). The effects of local real estate could easily be extreme given that wealth accumulation dynamics are most strongly determined by gains from investments (usually in the form of real estate) and positive income shocks (usually only achievable via education outcomes). The quality of education services often also depends on local real estate markets. Unsurprisingly, there are multiple paths of influence linking real estate dynamics to wealth accumulation outcomes. This suggests our model may benefit strongly from more detailed representation of local real estate markets.

Further experiments to explore may include examining modified tax policies, the role of credit (personal and business), entrepreneurship, as well as potential race-correlated financial behavioral patterns.

Conclusion

This paper continues the momentum of several recent efforts to inform on historic challenges and potential policy solutions (e.g., the recent racial wealth gap discussion series from Edwards (2022), Paige (2022), Welburn, Lima et al. (2022)). However, these discussions often highlight the intellectual challenges of understanding long run impacts of policy intervention that inherently involve significant population heterogeneity and model stochasticity. This paper responds to the need to act to address historic inequity, and to policy evaluation of what interventions are effective or efficient at alleviating the adverse wealth gap outcomes. We hope that future researchers will use the advancements introduced within this paper to furnish policy-makers with the requisite tools for identifying & prioritizing policy interventions.

References

- Aladangady, A. and A. Forde (2021). "Wealth Inequality and the Racial Wealth Gap."
- Bonczar, Thomas P., and Allen J. Beck. 1997. "Lifetime Likelihood of Going to State or Federal Prison." *Bureau of Justice Statistics Special Report*, no. March: 1–13.
- Boshara, R. (2020). "Examining Black-White Wealth Gaps." Federal Reserve Bank of St. Louis, Center for Household Financial Stability.
- Bosworth, Barry, and Rosanna Smart. 2009. "Evaluating Micro-Survey Estimates of Wealth and Saving." *SSRN Electronic Journal*. January.
- Chen, Siyan, and Saul Desiderio. "Calibration of Agent-Based Models by Means of Meta-Modeling and Nonparametric Regression." *Computational Economics* (2021): 1-22.
- Coates, T.-N. (2015). *The case for reparations*, Columbia University Press.
- Davis, Paul K., Angela O'Mahony, Timothy R. Gulden, Osonde A. Osoba, and Katharine Sieck. *Priority challenges for social and behavioral research and its modeling*. RAND Corporation Santa Monica United States, 2018.
- Darity Jr, W. A. (2021). "The True Cost of Closing the Racial Wealth Gap."
- Darity Jr, W. A. and A. K. Mullen (2020). *From here to equality: Reparations for Black Americans in the twenty-first century*, UNC Press Books.
- Edwards, K. A. (2022). *Accounting for Black-White Wealth Differences: A Stylized Model of Wealth Accumulation*. Santa Monica, CA, RAND Corporation.
- Federal Deposit Insurance Corporation, National Rate on Non-Jumbo Deposits (less than \$100,000): Savings (DISCONTINUED) [SAVNRNJ], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/SAVNRNJ>, September 24, 2021.
- Federal Deposit Insurance Corporation, National Rate on Non-Jumbo Deposits (less than \$100,000): Money Market (DISCONTINUED) [MMNRNJ], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/MMNRNJ>, September 25, 2021.
- Federal Reserve Board (2019). *Survey of Consumer Finances (SCF)*. B. o. G. o. t. F. R. System.
- Fox, L. E. (2016). "Parental Wealth and the Black-White Mobility Gap in the U.S." *Review of Income and Wealth* 62(4): 706-723.
- Hamilton, D. and W. A. Darity Jr (2010). "Can 'Baby Bonds' Eliminate the Racial Wealth Gap in Putative Post-Racial America?" *The Review of Black Political Economy* 37(3-4): 207-216.

- Johnson, David S., Katherine A. McGonagle, Vicki A. Freedman, and Narayan Sastry. 2018. "Fifty Years of the Panel Study of Income Dynamics: Past, Present, and Future." *Annals of the American Academy of Political and Social Science* 680 (1): 9–28.
- Loury, Glenn C. "Is Equal Opportunity Enough?." *The American Economic Review* 71, no. 2 (1981): 122-126.
- Loury, Glenn C. "Intergenerational transfers and the distribution of earnings." *Econometrica: Journal of the Econometric Society* (1981): 843-867.
- Osoba, Osonde, and Paul K. Davis. "An artificial intelligence/machine learning perspective on social simulation: New data and new challenges." *Social-Behavioral Modeling for Complex Systems* (2019): 443-476.
- Panel Study of Income Dynamics, public use dataset. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI 2019.
- Paige, J. W. (2022). Examining the Loss of Wealth and Downward Mobility of African Americans: A Review of Challenges for the Black Middle Class. Santa Monica, CA, RAND Corporation.
- Pfeffer, F. T. (2018). "Growing Wealth Gaps in Education." *Demography* 55(3): 1033-1068.
- Pfeffer, Fabian T., Robert F. Schoeni, Arthur Kennickell, and Patricia Andreski. 2016. "Measuring Wealth and Wealth Inequality: Comparing Two U.S. Surveys." *Journal of Economic and Social Measurement* 41 (2): 103–20.
- Rothstein, R. (2015). "The Racial Achievement Gap, Segregated Schools, and Segregated Neighborhoods: A Constitutional Insult." *Race and Social Problems* 7(1): 21-30.
- Rothstein, R. (2017). *The color of law: A forgotten history of how our government segregated America*, Liveright Publishing.
- Shapiro, T., T. Meschede and S. Osoro (2013). "The roots of the widening racial wealth gap: Explaining the black-white economic divide."
- Spizman, L. W., Marc A. (2008). "A Note on Utilizing the Geometric Mean: When, Why and How the Forensic Economist Should Employ the Geometric Mean Papers." *J. Legal Econ.* 15: 43.
- The Pew Charitable Trusts. 2010. "Collateral Costs: Incarceration's Effect on Economic Mobility," 1–40.
- Toney, J. and C. L. Robertson (2021). "Intergenerational Economic Mobility and the Racial Wealth Gap." *AEA Papers and Proceedings* 111: 206-210.
- Kohler, Ulrich. 2015. "PSIDTOOLS: Stata module to facilitate access to Panel Study of Income Dynamics (PSID)," *Statistical Software Components S457951*, Boston College Department of Economics, revised 06 Jul 2021.
- Welburn, J. W., P. N. d. Lima, K. B. Kumar, O. A. Osoba and J. Lamb (2022). *Overcoming Compound Racial Inequity: Policies and Costs for Closing the Black-White Wealth Gap*. Santa Monica, CA, RAND Corporation.
- Western, Bruce, Jaclyn Davis, Flavien Ganter, and Natalie Smith. 2021. "The Cumulative Risk of Jail Incarceration." *Proceedings of the National Academy of Sciences of the United States of America* 118 (16): 1–6. <https://doi.org/10.1073/pnas.2023429118>.

Wilshire Associates, Wilshire 5000 Total Market Index [WILL5000IND], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/WILL5000IND>, September 24, 2021.