PRELIMINARY AND INCOMPLETE: DO NOT CIRCULATE

Decomposing Wealth Differences: Earnings, Demographics, and Rates of Return*

Kartik Athreya† Grey Gordon‡ John Bailey Jones‡
Urvi Neelakantan‡

March 20, 2025

Abstract

Using a rich consumption-savings model, we decompose cross-sectional and life-cycle differences in the wealth distributions of Black and White individuals. We calibrate the model to match the life-cycle wealth profiles of White individuals. When we allow the model's forcing processes to vary by race, the model reproduces 86% of the cross-sectional mean difference and 93% of the median difference in wealth. Our decomposition analyses suggest that differences in employment by race contribute most to wealth differences at the bottom of the distribution, while earnings and educational differences contribute more in the middle. Scale effects in the returns on wealth amplify differences at the top.

JEL codes: D31, E21, J15
Keywords: wealth heterogeneity, life cycle, saving

*We are grateful to Mariacristina De Nardi, Jeffrey Smith, and participants at several conferences for helpful comments. Larissa Chan, Emily Emick, Simon Farbman and Luna Shen provided excellent research assistance. This research was supported in part through computational resources provided by the BigTex High Performance Computing Group at the Federal Reserve Bank of Dallas. The views expressed here are those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of New York, the Federal Reserve Bank of Richmond or the Federal Reserve System.

†Federal Reserve Bank of New York
‡Federal Reserve Bank of Richmond

1 Introduction

In this paper, we use a rich consumption-savings model to decompose cross-sectional and life-cycle differences in the wealth distributions of Black and White individuals. We calibrate the model's preference parameters to match the life-cycle wealth profiles of White individuals, while allowing our estimates of forcing processes—earnings, employment, incarceration, demographics, inheritances, and rates of return—to vary by race.[1] [2] Within the cohort we study, the model reproduces 86% of the difference in mean wealth (in the cross-section) and 93% of the difference in median wealth between Black and White individuals.

Among the factors that we consider, our decomposition analyses suggest a leading role for earnings. Differences in employment rates dominate at the bottom of the wealth distribution, while differences in earnings when employed, along with educational outcomes, matter more at the middle and top. We also consider several mechanisms that are potentially independent of race. Perhaps the most important of these are "scale effects" in rates of return, namely the positive correlation between wealth and portfolio rates of return. We find that scale effects amplify differences at the top of the wealth distribution, including the overall mean.

Figure 1 presents life-cycle wealth profiles for the cohort of individuals born between 1957 and 1964.[3] Beginning with medians, the smoothed profile in panel (a) shows that, at the "retirement" age of 65, the median wealth of Black individuals is just under $46,000 (2010 dollars), around one-fifth of the median for White individuals. The difference in medians of $174,000 is one-third of the $526,000 difference in means at the same age. Figure 1 also shows that median and mean wealth differences widen substantially after early adult life (around age 30) and stabilize after retirement. Turning to the left tail of the wealth distribution, the bottom panel of the figure shows the fraction of individuals with zero or negative net worth. The fraction of Black individuals with zero or negative net worth remains higher than the fraction of White individuals over the entire life cycle.

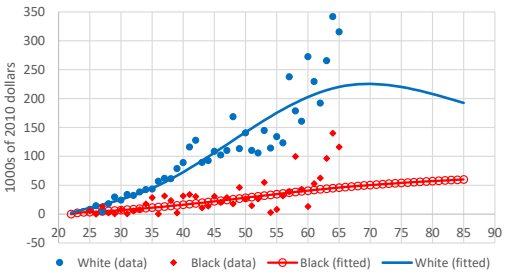
An alternative summary of how wealth differs by race is provided by a graphical

[1] Our analysis does not speak to the factors underlying the differences in these forcing processes.

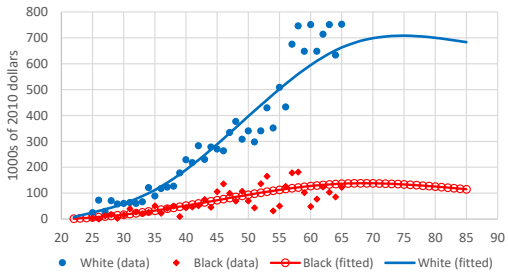
[2] In addition to race, we allow for heterogeneity by sex, marital status and education; we could easily use our methodology to decompose wealth differences along these dimensions.

[3] The data points in Figure 1 come from a synthetic panel of the 1957-1964 birth cohort constructed using the Survey of Consumer Finances (SCF). The smoothed (and extrapolated) profiles come from a set of regressions on individuals born between 1933 and 1988 with cohort controls. At each SCF wave we drop the top 0.5% of observations. Wealth is expressed on an individual basis, with the wealth holdings of couples split evenly. See Appendix A for further details.

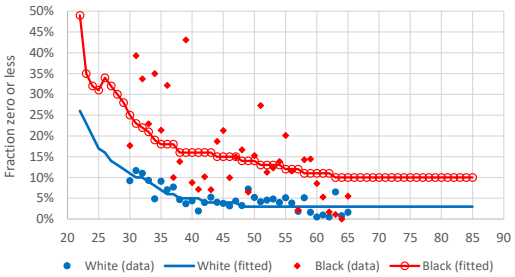
Figure 1: Wealth Differences over the Life Cycle



(a) Median



(b) Mean



(c) Fraction with Net Worth $\leq$ \$0

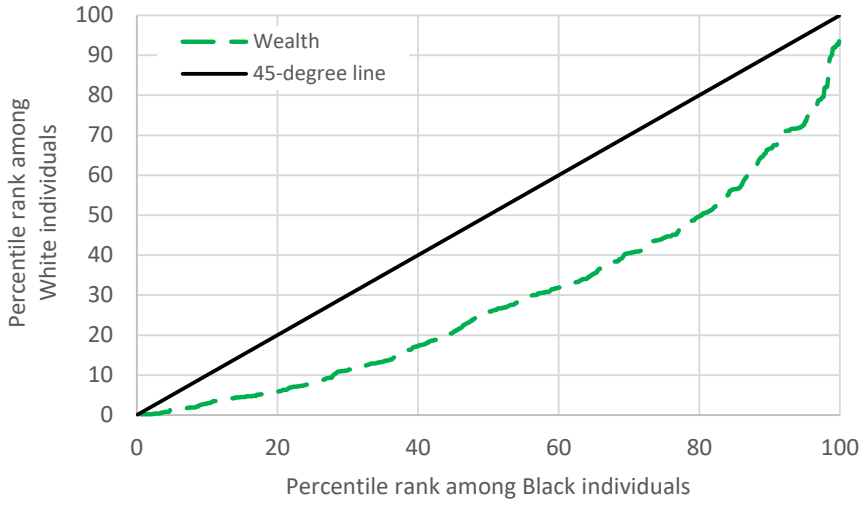
device based on rank.[4] Figure 2 is read as follows: each point on the dashed line maps where a given percentile of the Black wealth distribution places one on the White wealth distribution. For example, the median of the Black wealth distribution corresponds to the 25th percentile of the White wealth distribution. The 90th percentile of the Black wealth distribution corresponds to the 65th percentile of the White one. The differences do not close except at the very top. Figure 2 can also be distilled into a one-dimensional "rank Gini coefficient," the ratio of the area between the rank line and the 45-degree line to the entire area under the 45-degree line. We calculate a wealth rank Gini of 0.42.

The size and scope of the wealth differences found in the data suggest that we need to account for multiple factors. The substantial differences at the top of the wealth distribution indicate that factors important to large wealth holdings, such as high rates of return (Benhabib et al., 2019) and/or entrepreneurship (Boerma and Karabarbounis, 2023) could matter. The differences at the bottom of the distribution, on the other hand, suggest roles for differences in nonemployment and incarceration (Gordon et al., 2023).

To formalize these insights, we first carefully estimate a battery of forcing processes.

[4]Perhaps the best-known application of rank analysis to Black-White differences is Bayer and Charles (2018), who examine earnings. Jones and Neelakantan (2022b) document Black-White rank differences in wealth, extending the approach used by Sabelhaus and Volz (2024).

Figure 2: Rank Differences in Wealth



These include earnings, employment, incarceration, demographics, inheritances and inter-vivos transfers, initial wealth, and rates of return on wealth. We allow the processes to differ by education, sex/marital status,[5] and race. A strength of our analysis is that we track a single cohort—those born between 1957 and 1964—over time. This allows us to use panel data from the National Longitudinal Survey of Youth 1979 (NLSY79) to estimate precisely the earnings and demographic processes facing individuals in our cohort. Of particular importance is the joint stochastic process for nonemployment, earnings and incarceration, which extends our earlier work in Gordon et al. (2023). These estimates, along with our detailed estimates of portfolio shares, are of interest in their own right.

We then use our estimated processes as inputs in a rich life-cycle consumption-savings model. We calibrate the model's preference parameters to match the wealth profiles of White individuals, with distinct targets for each age-education-sex/marital status cell. Using such a rich set of targets, along with highly disaggregated forcing processes, results in the parameters being tightly identified. We assume that, conditional on other correlates, Black individuals have the same preferences as Whites.

Using our cohort to construct a cross-section, we find that the model matches well the wealth distributions we wish to study. It reproduces 86% of the mean wealth difference, 93% of the median difference, 81% of the wealth rank Gini, and 118% of the difference in those with non-positive wealth. This means that we can distribute most of the cross-sectional wealth heterogeneity found in the data among the endowments and forcing

[5]We differentiate between single men, single women, and married individuals.

processes that we feed into our model. We do this using the Shapley-Owen decomposition (Shorrocks, 2013). We find that the most important force is differences in earnings and employment, as measured by the joint process for the two. Disaggregating further, we find that Black-White differences in employment, reflecting in part differences in incarceration, contribute more to model-predicted wealth differences than do differences in the earnings of the employed. This is especially the case at the bottom of the distribution. For example, if Black individuals had the same employment rates as White individuals, the Black-White difference in the fraction of individuals with non-positive net worth would shrink by 62%. Differences in employment likewise account for 44% of the median difference and 30% of the mean difference. Differences in the earnings of the employed play a smaller but still substantial role, accounting for 22% and 20% of the median and mean differences, respectively.

It bears noting that the earnings and employment differences eliminated in our experiments are those that remain after controlling for educational outcomes. The contributions of education itself, mediated in part through earnings and employment but also through demographic outcomes and preference parameters, is largest at the top of the distribution—equalizing educational outcomes would reduce the mean wealth difference by 26%. In combination, earnings, employment and education account for the bulk of the wealth differences in our model. Forces such as demographic differences, intergenerational transfers, or race-specific differences in rates of return contribute relatively little.

Because Black-White differences in wealth are, in proportional terms, substantially larger than the corresponding differences in earnings, the model can replicate the wealth differences only if it contains several amplification mechanisms. These include bequest motives, as well as taxes and transfers, but arguably the most important are scale effects in the returns to wealth. If, for example, all households earned a fixed return of 3%, the mean wealth difference would shrink by 60%. Rates of return are by far more important at the top of the wealth distribution: equalizing returns results in a smaller change in the median difference and an even smaller change (16%) in the difference in non-positive wealth. Because the scale effects are for the most part independent of race, removing them has a fairly modest effect (21%) on the rank Gini, even when they have large effects on level (dollar) differences. In contrast, earnings and employment heterogeneity have substantial effects on both rank and level differences.[6]

[6]In the language of Bayer and Charles (2018), the scale effects induce "distributional" differences, changing level differences even when they do not change relative rankings, much as multiplying every individual's wealth holdings by the constant $x \in [0, 1)$ would reduce level differences without changing relative ranks. In contrast, Black-White differences in endowments induce "positional" differences, changing

The rest of the paper proceeds as follows. Section 2 provides a literature review. We lay out our life-cycle model in section 3. We describe our main data sources in section 4 and in section 5 we describe how we estimate the model's forcing processes and calibrate its parameters. In section 6 we assess the model's fit and decompose Black-White wealth differences among the model's forcing processes and endowments. We examine the model's mechanisms more closely in section 7 and conclude in section 8.

2 Related Literature

A useful way to organize the literature on wealth differences is to work off a basic asset accumulation equation:

$$a_{t+1} = (1 + r_t)a_t + e_t + x_t - c_t, \tag{1}$$

where: a_t denotes wealth at time t; r_t denotes the returns on those assets; e_t denotes labor earnings; x_t denotes other forms of income and expenses, including taxes, transfers, medical expense and received inheritances; and c_t denotes consumption. Most papers seeking to explain wealth differences consider at least one of these elements. In the context of Black-White wealth differences, several recent papers (Ashman and Neumuller 2020; Aliprantis et al. 2022) emphasize earnings, e_t.

As a distinct topic, the literature on racial differences in earnings is more extensive than that on wealth (Altonji and Blank 1999; Lang and Lehmann 2012). The key points are that earnings differences are large and in the past few decades have been closing slowly, if at all (Bayer and Charles, 2018). Even with similar returns and saving rates, this immediately implies large and persistent differences in wealth.

Earnings are essential to our analysis as well. In our model, savings respond to a stochastic process for nonemployment, earnings and incarceration estimated from the NLSY79. Building on our earlier work in Gordon et al. (2023), this joint process accounts for incarceration and other forms of nonemployment—both of which are substantially higher for Black individuals—and the tendency of individuals with high rates of nonemployment (or incarceration) to have lower earnings when employed. This allows our quantitative model to distinguish the consequences of differences in employment from the consequences of differences in the earnings of the employed.

In proportional terms, however, Black-White heterogeneity in wealth tends to be considerably larger than the corresponding heterogeneity in income (e.g., Wolff 2018, Table 1).

relative ranks as well as level differences.

It is also well known that Black households are less heavily invested in high-yielding assets such as stocks (e.g., Derenoncourt et al. 2022). These findings motivate us to allow financial portfolios (and therefore rates of return, r_t) to vary with a number of covariates, including race.

We find that much of the difference in portfolio composition can be predicted by total wealth without reference to race: wealthier individuals of either race are more heavily invested in high-return assets. This may reflect the ability of wealthier individuals to hire better financial advisors or have more of the resources required to participate in high-yielding investments (e.g., pay fixed costs or make a house down payment). If so, scale effects in rates of return will sustain and reinforce initial differences in wealth. On the other hand, if earning high returns through activities such as entrepreneurship is essential to accumulating large amounts of wealth (Quadrini 1999, Cagetti and De Nardi 2006), limited access to those investments may be one reason why Black individuals have less wealth (Boerma and Karabarbounis, 2023). While our framework is most consistent with the first interpretation, the race effects in our portfolio share functions may reflect differential access.

Although we allow for portfolio allocations to differ by race, data limitations lead us to assume that within most classes of assets and debts, rate of returns do not differ by race.[7] The one exception is housing, where we impose the return differential estimated by Kermani and Wong (2021), who report that Black homeowners earn lower returns due to a higher risk of distress sales. In other work, Sakong (2022) finds lower returns due to systematic differences in the timing of purchases and sales. Wolff (2022) reports similar findings and provides an extended literature review.

Another potential source of wealth differences lies in differences in consumption, c_t. A number of papers have found that households with higher incomes save at higher rates (Dynan et al. 2004; Bozio et al. 2017). There are a number of potential reasons why financial wealth might be a luxury good. These include warm glow bequest motives that are stronger for the rich (De Nardi, 2004), means-tested transfers and progressive Social Security benefits that discourage saving by poorer households (Hubbard et al. 1995; Huggett and Ventura 2000), and the higher returns available to richer households. Our model contains all of these mechanisms. As with returns that rise with wealth, any tendency for savings to be a luxury good will reinforce existing wealth differences.[8]

[7] We allow rates of return to differ by sex and education, following Fagereng et al. (2020), and allow for a "pure" scale effect independent of the portfolio composition, following Fagereng et al. (2020) once again. None of these adjustments are explicitly race-based.

[8] As with returns that rise with wealth, the tendency of households with higher wealth to save at higher

Our model also allows for expenditures to translate to consumption in ways that depend on family structure (marriage and children). This motivates us to introduce a demographic transition process in our model that captures Black-White differences in household size and composition. Our approach has many parallels with Ashman and Neumuller (2020); one difference is that we explicitly account for co-residing children in our flow utility function and allow their number to vary.

In addition to affecting consumption, non-earnings sources of funds, x_t, and initial wealth, a_0, also affect wealth directly. In this context, inheritances and inter-vivos transfers loom large, as White individuals are far more likely to receive such transfers and, conditional on receiving them, to receive larger amounts (Jones and Neelakantan, 2022a). Sabelhaus and Thompson's (2023) decomposition analysis implies that inheritances and intervivos transfers explain about 14-16% of racial wealth differences, while Ashman and Neumuller's (2020) quantitative model attributes 26% of the median difference to intergenerational transfers. Additionally, Brandsaas (2024) finds that inter-vivos transfers explain 30% of the Black-White homeownership difference. For our analysis, we estimate directly from the data the distribution of transfers that individuals in our model receive, and allow it to depend on race and education. We likewise take the initial distribution of wealth for new adults (age 22) from the data for our cohort. This allows us to capture, among other things, differences in student loans (Houle, 2014). We also allow for uncertain medical expenses and, as noted above, taxes and government transfers.

Relative to the existing structural literature, the contribution of our paper lies in its breadth and detail. We consider a wide range of mechanisms, measure them precisely, and calibrate the model to a rich set of targets. We take seriously the possibility that different dimensions of wealth heterogeneity reflect different mechanisms, and show that this is in fact the case. It is likely that with alternative calibrations, none deeply implausible, any single summary measure of wealth heterogeneity could be replicated by multiple mechanisms. To sort out the competing explanations we calibrate a model with multiple mechanisms to multiple wealth targets.

Beyond connecting to existing work on racial differences, our analysis features many of the mechanisms that appear in the voluminous literature on wealth heterogeneity. In recent surveys, De Nardi and Fella (2017) conclude that "the transmission of human capital and bequests across generations, entrepreneurship, and medical-expense risks" play

rates may also reflect reverse causality: individuals who prefer to save are more likely to become wealthy (Krusell and Smith, 1998). We allow our model's preference parameters to vary by education but not by race or at the individual level.

leading roles, while Benhabib and Bisin (2018) emphasize the importance of heterogeneous rates of return in generating a thick right tail in the distribution of wealth. These findings inform the construction of our model.

3 Model

Our model consists of heterogeneous individuals choosing consumption, saving and bequests. Individuals differ ex-ante by race, sex and education. They differ ex-post in their realizations of marriage and fertility, lifespan, initial wealth and received inheritances, earnings, medical expenses, and rates of return. These processes are exogenous, and in most cases depend on an individual's race, sex, education and age. Individuals pay taxes to and receive transfers from the government. They can borrow subject to limits. The composition of their wealth portfolios, and thus their rates of return, vary mechanically with their total wealth in a way that is consistent with the data.

3.1 Demographics

Individual i is characterized by their race, sex and education, which we collect in the vector d_i. A person's age is indexed by t, which is discrete and measured in years. Individuals enter the model at age t_0 and live up to age T. Survival is uncertain and given by the rate s_{t,d_i}.

While we allow for marriage, the unit of analysis in our model is the individual. We assume that married individuals keep "separate books": couples split their income evenly, but each spouse chooses their own level of consumption and holds their own assets. Marriage, spousal income and co-residing children are all modelled as exogenous stochastic processes. These assumptions eliminate the need to compute equilibria in the marriage market, which simplifies our model considerably.

Let the vector $f_{i,t}$ denote the individual's marital status, $m_{i,t} \in \{0, 1\}$, and number of co-residing children (up to 3), following a Markov chain described below. We will use $N_{i,t} = N(f_{i,t})$ to denote consumption-equivalized household size, and

$$n_{i,t} = n(f_{i,t}) \coloneqq \frac{N(f_{i,t})}{1 + m_{i,t}}, \tag{2}$$

to denote the "effective size" of person i's household.

3.2 Earnings and financial income

The individual's earnings are given by $e_{i,t} = e(d_i, t, f_{i,t}, f_{i,t-1}, \lambda_{i,t}, \mu_{i,t})$. Earnings follow a latent Markov chain (see, e.g., Bartolucci et al. 2010), where $\lambda_{i,t}$ denotes the unobserved latent state and $\mu_{i,t}$ controls the mapping from this latent state to the observed outcome $e_{i,t}$. Loosely speaking, $\lambda_{i,t}$ and $\mu_{i,t}$ can be viewed as persistent and transitory earnings shocks, respectively. There are eight possible earnings outcomes: nonemployment, incarceration, and six positive earnings values.

Our model of earnings follows closely the one developed in Gordon et al. (2023). One way in which we diverge from our earlier work is that here we account for marriage: if total household earnings are $e^*_{i,t}$, the wife and husband receive $\varsigma e^*_{i,t}$ and $(1-\varsigma)e^*_{i,t}$, respectively. Earnings transitions also depend on the number of children, and whether any of them are new (captured by $f_{i,t}$ and $f_{i,t-1}$). We assume that the causality flows strictly from demographics to earnings: earnings shocks do not affect the probability of marriage or the arrival of children. This simplifies our calibration process greatly, as it allows us to estimate the demographic process in isolation, and then take its parameters as given when estimating the process for earnings.

The individual's net financial income is given by $r_{i,t}a_{i,t}$, where $a_{i,t}$ denotes net worth (wealth), and $r_{i,t} = r(a_{it}, d_i, f_{i,t}, t, \rho_{i,t})$ is the rate of return. It is well known that household portfolios vary systematically across the wealth distribution, with richer households holding higher shares of high-yielding assets such as stocks and business equity (Campbell, 2006). We capture this in our model by allowing $r_{i,t}$ to vary systematically with wealth, combining portfolio shares estimated from the data with asset-specific returns. We also allow for transitory return shocks, $\rho_{i,t}$.

3.3 Budget constraints

The individual's total financial resources, cash-on-hand or $coh_{i,t}$ (àla Deaton 1991), are given by

$$coh_{i,t} = (1 + r_{i,t})a_{i,t} + e_{i,t} + ss_{i,t} + b_{i,t} - mx_{i,t} - \tau_{i,t}. \tag{3}$$

In addition to wealth, returns on wealth, and earnings, households may receive Social Security benefits ($ss_{i,t}$) and inheritances/inter-vivos transfers ($b_{i,t}$). They pay taxes ($\tau_{i,t}$) and out-of-pocket medical expenses ($mx_{i,t}$).

The law of motion for net worth is then

$$a_{i,t+1} = coh_{i,t} + mtt_{i,t} - c_{i,t}, \tag{4}$$

where $c_{i,t}$ is consumption and $mtt_{i,t}$ denotes means-tested transfers. These transfers provide a consumption floor (Hubbard et al., 1995):

$$mtt_{i,t} = \max\left\{n_{i,t}^{fpl}\,\underline{z} - \max\{coh_{i,t}, \delta coh_{i,t}\},\ 0\right\}, \quad \delta \in (0,1]. \tag{5}$$

The consumption floor itself is $n_{i,t}^{fpl}\,\underline{z}$: z denotes the per-capita floor, which depends on age and whether the individual is incarcerated (captured by $\lambda_{i,t}$); and $n_{i,t}^{fpl} = n^{fpl}(f_{i,t})$ is the equivalence scale implied by the US Federal Poverty Limit divided by $1 + m_{i,t}$. In a standard consumption floor set-up, $\max\{coh_{i,t}, \delta coh_{i,t}\}$ is replaced by $coh_{i,t}$. When this is the case, means-tested transfers top up poorer individuals' cash-on-hand, allowing them to purchase the minimum consumption level $n_{i,t}\,\underline{z}$ without going into debt. This generates perverse incentives, however, in that individuals with negative net worth, who are common in the data, would have their debt retired instantly. We therefore limit the amount of net debt (negative cash-on-hand) that can be offset by the transfers to a fraction of the total, $\delta coh_{i,t}$, where δ is the debt discharge rate.

Individuals not receiving means-tested transfers face a borrowing limit proportional to their expected non-financial income:

$$a_{i,t+1} \geq -\phi_t(\bar{e}_{i,t} + ss_{i,t} - \tau_{i,t}), \tag{6}$$

$$\bar{e}_{i,t} \coloneqq E(e_{i,t} \mid d_i, t, f_{i,t}, f_{i,t-1}, \lambda_{i,t}). \tag{7}$$

The quantity $\bar{e}_{i,t}$ is an expectation taken over the transitory shock $\mu_{i,t}$; using expected rather than realized earnings allows us to remove the transitory shock from the individual's state vector.[9] When transfers are positive ($mtt_{i,t} > 0$), equation (6) does not hold. We assume instead that positive transfers imply that $c_{i,t} = n_{i,t}\,\underline{z}$ and $a_{i,t+1} = \min\{0, (1-\delta)coh_{i,t}\}$, consistent with the notion that the transfers are in kind (e.g., Medicaid) and must be consumed immediately. Incarcerated individuals have their consumption restricted in the same way.

3.4 Preferences

The individual's flow utility function is

$$u(c_{i,t}, h_{i,t}, n_{i,t}) = \frac{1}{1-\sigma} n_{i,t} \left(\frac{c_{i,t}^{\zeta}(h_{i,t} + \underline{h})^{1-\zeta}}{n_{i,t}} \right)^{1-\sigma}, \quad \sigma \geq 0, \quad \zeta \in (0,1), \tag{8}$$

[9] If the borrowing limit depended on realized earnings, transitory earnings shocks would affect the consumer's budget constraints in a way distinct from their contribution to cash-on-hand.

where $c_{i,t}$ and $n_{i,t}$ denote the individual's consumption and effective household size, respectively, and $h_{i,t} + \underline{h}$ denotes the flow of housing services. The parameter ζ gives the weight on consumption within the consumption-housing composite, while σ is the coefficient of relative risk aversion over the composite.

The service flows from owner-occupied housing are given by

$$h_{i,t} = \max\big\{h(a_{i,t+1}, d_i, f_{i,t}, t+1),\, 0\big\}, \tag{9}$$

where $h(\cdot)$ is the individual's imputed housing stock, found using the equations used to predict portfolio shares in the rate of return calculations. Individuals thus lack direct control over their housing wealth, but can adjust it indirectly through saving.[10] The term $\underline{h} > 0$ accounts for the housing services consumed by renters. Because we include imputed rents in our measure of housing returns, the inclusion of $h_{i,t}$ in the utility should be viewed as capturing the *additional* utility from homeowning as opposed to renting.[11]

Individuals value bequests according to

$$\upsilon_b(a) = \frac{\kappa_b^\sigma}{1-\sigma}(\kappa_b \chi_b + a)^{1-\sigma}, \tag{10}$$

following De Nardi (2004) and Lockwood (2018), among others. The parameter κ_b determines the strength of the bequest motive, while χ_b determines the extent to which bequests are a luxury good.

3.5 Recursive formulation

The individual discounts future utility at the rate β. Eliminating i and t subscripts—except in the case of direct age dependencies—and using primes to denote one-period-

[10] We model time-t housing flows as a function of time-$t+1$ wealth for computational reasons. In particular, the state space for our model uses cash-on-hand, which combines assets with several other variables, rather than assets themselves. Allowing time-t housing to depend on time-t assets would require us to add assets to the state space.

[11] The non-pecuniary benefits of homeownership include opportunities for customization and emotional attachment to one's property and/or neighborhood (Nakajima and Telyukova, 2017).

ahead values, we can write the individual's Bellman equation as

$$
\begin{aligned}
V_t(coh, f, \lambda, d) = & \\
& \max_{c\geq 0, a'} u(c,h,n) + \beta\Big(s_{d,t}\mathbb{E}\big(V_{t+1}(coh', f', \lambda', d) \,|\, a', f, \lambda, d, t\,\big) + (1 - s_{d,t})\upsilon_b(a')\Big) \\
& \quad \text{s.t. equations (3), (5) and (9)}, \\
& \qquad c = \begin{cases} coh - a' & \text{if } mtt = 0 \text{ (and not incarcerated)} \\ n^{fpl}\,\underline{z} & \text{if } mtt > 0 \text{ (or/and incarcerated)} \end{cases}, \\
& \qquad a' \geq -\phi_t(\bar{e} + ss - \tau) \text{ if } mtt = 0, \\
& \qquad\quad = \min\{0, (1-\delta)coh\} \text{ if } mtt > 0.
\end{aligned}
$$

4 Data

While we rely on a variety of data sources, our three principal datasets are the NLSY79, the Current Population Survey (CPS), and the SCF. We use the NLSY79 to estimate the stochastic processes for demographic transitions and earnings. We use CPS data to estimate earnings quantiles at each age for each race, sex, education and marital status combination. We use SCF data to construct the life-cycle wealth profiles that we target in our estimation and the portfolio shares that go into our rate of return process and our imputed housing service flows. We also use the SCF to measure the inheritances and inter-vivos transfers that individuals receive over their lives.

The NLSY79 is a nationally representative panel survey of young men and women born between 1957 and 1964. From 1979-1994, respondents were interviewed every year; since 1994, interviews occur every other year. The NLSY79 collects information about education, employment, family, and finances, and it is one of the few nationally representative surveys that enables us to observe an individual's incarceration status. Our sample selection criteria are the same as in Gordon et al. (2023), but we have extended the sample to 2020, allowing us to track individuals through their early 60s.

Although the NLSY79 is our principal data source for estimating the earnings process, to calculate the cutoffs that delineate earnings quantiles, we use the larger sample available in the CPS (downloaded from IPUMS; Ruggles et al. 2020). We use data from 1976 to 2023. Since the CPS is not a panel, we employ a synthetic cohort approach. Our sample selection criteria and synthetic panel construction are the same as in Gordon et al. (2023). We

also use the CPS to construct within-quantile mean earnings, the earnings an individual receives when landing in a particular earnings bin.

Appendix C describes how we construct our earnings, employment and incarceration measures from the NLSY79 and CPS datasets. One complication is accounting for marriage. Earnings are measured at the household level and assumed to be split evenly between spouses. Employment is for the most part also defined at the household level, as an individual enjoying significant spousal income is categorized as employed.

The SCF is a survey conducted by the Federal Reserve Board that includes unique information on the balance sheets, income, and demographics of U.S. families. A notable feature of the SCF is its heavy oversampling of the rich, allowing it to provide accurate estimates of the upper tail of the wealth distribution. The SCF is a repeated cross-section, with new data collected every 3 years. As with the CPS, we utilize this data as a synthetic panel. Our data cover the years 1983-2022.

5 Parameterization

We parameterize our model in three stages. First, we set some parameters at standard values or by reference to established work. Second, we estimate the processes for demographic transitions (marriage and fertility) and earnings from the NLSY79, and the process for portfolio shares from the SCF. Third, we estimate a number of structural parameters using a rich array of wealth-related moments for White households from the SCF.

Throughout our empirical analysis, we group people into four education bins: no high school degree (less than high school, or LHS); high school degree but no college experience (HS); college attendance without a bachelor's degree (some college, or SC); and bachelor's degree or higher (college graduate, or CG).

5.1 First stage: Parameters taken from other analyses

We set the coefficient of relative risk aversion, σ to 2, a standard value. We set the borrowing constraint parameter ϕ_t to 0.25 for individuals younger than 65 and to 0 for those 65 and older.[12] The coefficient on consumption in the consumption-housing composite, ζ

[12] As can be inferred from Table 1 in Fulford (2015), borrowing limits as a fraction of annual income vary from about 1/3 to 1/2. We choose a slightly lower ratio, 1/4, to account for possible selection in access to credit. We shut down borrowing in retirement as relatively few individuals age 65 or older carry credit card balances.

(see equation (8)), equals 0.855, the ratio of non-housing consumption expenditures to total consumption in the 2023 National Accounts.[13]

We model out-of-pocket medical spending ($mx_{i,t}$) as

$$mx_{i,t} = MX_{i,t} \times \big(1 - \Pr(i \text{ is insured at time } t) \cdot cir_{i,t}\big), \tag{11}$$

where $MX_{i,t}$ denotes total medical spending and $cir_{i,t}$ is the coinsurance rate. The process for $MX_{i,t}$ and the coinsurance rate $cir_{i,t}$—which equals 0.229 and 0.291 before and after age 65, respectively—are from Conesa et al. (2018). The probability that the individual is insured is the probability of having any insurance reported in Card et al. (2008, Table 1). Black and less-educated individuals are less likely to be insured, with the differences shrinking substantially when individuals qualify for Medicare at age 65. In addition to the coverage specified in equation (11), individuals may have some of their expenses covered through the consumption floor (equation (5)), which incorporates Medicaid.

Our estimates of survival rates come for Sasson (2016), who tabulates race-, sex- and education-specific mortality rates for calendar years 1990, 2000 and 2010. These data allow us to construct the life-cycle profile of survival rates for the cohort aged 25-34 in 1990; we use Social Security projections (Social Security Administration, 2018, Table 2.2) to extrapolate beyond 2010. Table 1 shows the life expectancies associated with our estimated survival rates. Comparing races, Table 1 reveals that the smallest racial differences in life expectancy are among those without a high school degree. The largest differences are for high school graduates, with White women on average living 3.7 years longer than Black women, and White men living 5.1 years longer than Black men. The shorter lifespans of Black individuals should, all else equal, weaken their life-cycle savings motives relative to those of Whites.

Table 1: Life Expectancy at Age 25

	Women		Men	
	Black	White	Black	White
Less than high school	53.0	53.0	45.8	47.0
High school graduate	54.5	58.2	47.7	52.8
Some college	55.9	58.8	52.5	54.9
Bachelor's degree+	58.7	62.1	56.1	59.5

Source: Sasson (2016), extrapolated using Social Security projections (Social Security Administration, 2018, Table 2.2).

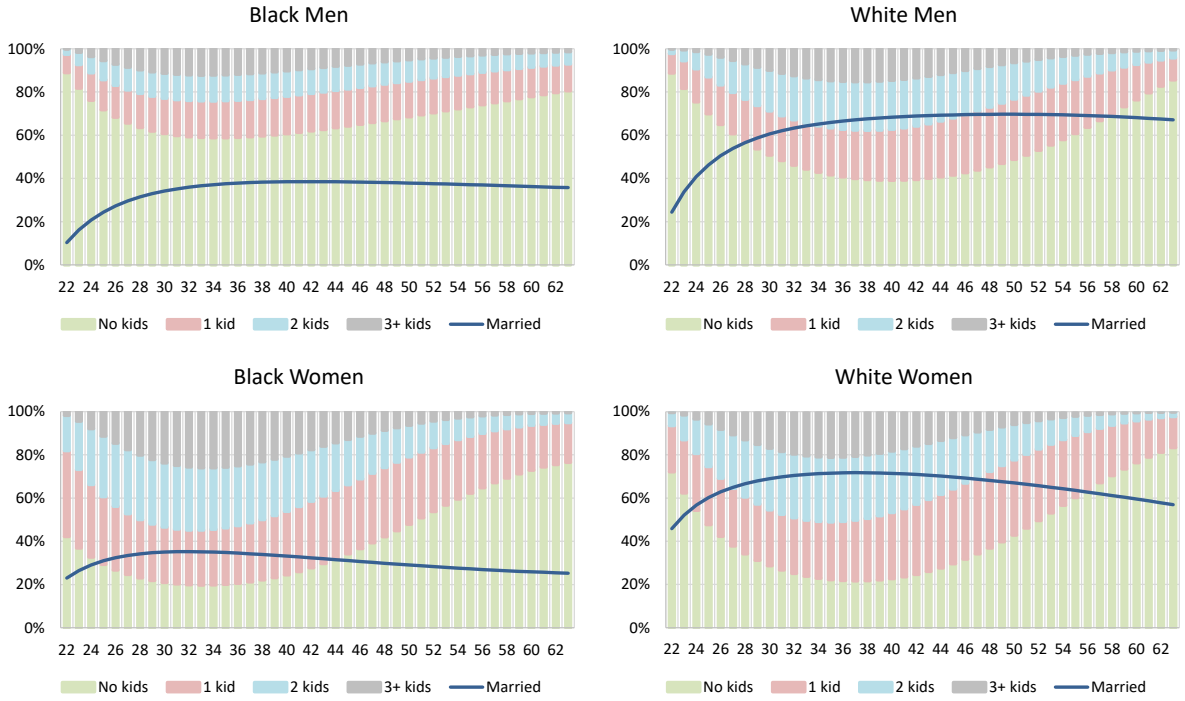
[13] See NIPA Table 2.4.5, lines 1 and 50.

5.2 Second stage: Forcing processes estimated from the data

5.2.1 Demographic transitions

The individual's marital status and number of co-residing children, $f_{i,t} \in \{0,1\} \times \{0,1,2,3+\}$, are modeled as an exogenous Markov Chain. We estimate separate sets of transition probabilities for each race-sex-education cell. We also account for non-reporting bias by allowing a missing value to be one of the observed outcomes. Appendix B provides a detailed description.

Figure 3: Marital Status and Number of Co-residing Children, High School Graduates



Note: Authors' calculations using data from the NLSY79. The solid line plots the fraction of households that are married, while the shaded areas give the distribution of co-residing children at each age.

Figure 3 shows life-cycle profiles for high school graduates. The solid line in each figure shows the incidence of marriage, while the shaded areas show the distribution of co-residing children at each age. Both the probability of marriage and the number of children peak when people are in their 30s, then fall as people divorce and children leave home. At any age, Black individuals are less likely to be married. Because children are more likely to live with their mothers, this results in Black men having fewer co-residing children than their White counterparts.

5.2.2 Employment, incarceration and earnings

Earnings, $e_{i,t}$, have eight possible outcomes: nonemployment, incarceration, and six positive earnings values (the first four correspond to mean earnings in the bottom four earnings quintiles; the fifth captures earnings in between the 80th and 94th percentiles of the earnings distribution, and the sixth captures earnings in the top 6%).

The distribution of observed earnings outcomes in turn depends on realizations of the persistent latent state, whose values are incarceration, persistent non-employment, and six positive earnings *potential* bins. The mapping from latent states to observed outcomes is not trivial. Incarceration and persistent non-employment map one-for-one into observed incarceration and non-employment, respectively. However, individuals in any positive latent earnings potential bin can realize (transitory) non-employment or any positive earnings outcome. Because the latent states follow a persistent Markov Chain, while the distribution of observed outcomes is transitory given the latent state, our earnings process is in many ways a generalization of the well-known "AR(1) plus White Noise" process.

Our earnings model builds on the one developed in Gordon et al. (2023). An important difference is that we now account for marriage: when a person is married, their earnings equal a fraction of the couples' total earnings. Earnings also depend on the number of children, and whether any of them are newborns. These linkages mean that the likelihood of an individual's observed earnings history depends on the likelihood of their observed family composition history.[14] To keep our estimation tractable, we assume that earnings outcomes do not affect family composition. This allows us to proceed in two steps, first estimating the earnings process, and then using the resulting coefficients when calculating the likelihood function for earnings. Appendix D describes the earnings process and its estimation in more detail.[15]

To fix ideas, Table 2 shows the latent state transition probabilities for single, childless, 25-year-old Black men with high school degrees. Three patterns bear noting. The first is that men with low earnings potential are more likely to transit to persistent nonemployment than those with high earnings potential. A man with no incarceration record

[14] Were the NLSY79 data available on an uninterrupted annual basis, we could condition our estimates of earnings transitions on realized family composition. However, there are gaps in the data, including when it becomes biennial, that require us to account for unobserved family transitions.

[15] For some groups—mostly women—the incidence of incarceration is so low that their incarceration-related parameters cannot be estimated with any precision. In these cases, we estimate a simplified model of employment and earnings, to which we add incarceration-related parameters estimated for other, similar groups. Our imputation approach, described in Appendix D.7, follows closely the one used in Gordon et al. (2023).

(IR = 0) in the bottom quintile of the earnings-potential distribution (Q1) has a roughly 5% chance of entering the non-employment state. The non-employment probability for an otherwise identical man in the top bin (Q6) is around 0.1%. Nonemployment in turn predicts lower future earnings: the first row of the matrix shows that men who exit the nonemployment state usually wind up in the bottom earnings potential bin.

Second, men who are nonemployed or have low earnings potential are much more likely to transit to jail. A 25-year-old Black man with no incarceration history (IR = 0) in the bottom quintile of the earnings-potential distribution (Q1) has a roughly 2.5% chance of becoming incarcerated at age 26. The incarceration probability for an otherwise identical man in the top bin (Q6) is 0.4%.

The third pattern is the prevalence of recidivism. A Black man currently in the bottom earnings potential quintile with an incarceration record has a 18% chance of being in jail the following year, an increase of almost 16 percentage points (pp) over that of a man with no record. Moreover, men who are currently incarcerated, should they leave jail, are most likely to exit to nonemployment or to the bottom earnings potential bin, where the odds of reincarceration are the highest. [16]

Table 3 shows the probabilities mapping from latent states to observed outcomes, once again using single, childless, 25-year-old Black men with high school degrees for illustration, but this time restricting ourselves to men with no incarceration record.[17] Notable here is the substantial likelihood that a worker with low earnings potential will be nonemployed. For example, a man in the bottom earnings potential quintile will be nonemployed 7% of the time.

A central contribution of our work is the careful estimation of the forcing processes that we feed into our consumption-savings model. To illustrate that our estimates are indeed reasonable, we close this section by highlighting life-cycle profiles drawn from the data and our estimation.

Looking first at nonemployment and incarceration, a general finding is that both rates are significantly lower for White individuals than for Black individuals with the same observable characteristics.[18] For example, at age 55, 34% of Black male high-school

[16] Of course, the effect of incarceration on employment depends on selection, specifically, the extent to which these individuals would, in the absence of incarceration, still have poor earnings prospects. Table 2 shows that negative selection into incarceration is indeed present. Gordon et al. (2023, Table 7) report that eliminating incarceration would reduce the Black-White difference in lifetime employment by 0.6 years (from 5.0 to 4.4 years) among male high school graduates, and by 1.6 years (from 6.5 to 4.9 years) among men without a high school degree.

[17] We do this for brevity, as the block for those with an incarceration record is qualitatively similar to the block for those without.

[18] Our model is silent on the root causes of differences in nonemployment. Potential causes include

Table 2: Latent State Transition Probabilities: 25-year-old Black Male High School Graduate, Single, No Kids

Current State ↓		Future State: Incarceration Record = 0 N	Q1	Q2	Q3	Q4	Q5	Q6	Jail	Incarceration Record = 1 N	Q1	Q2	Q3	Q4	Q5	Q6	Jail
IR = 0	N	0.76	0.21	0.00	0.00	0.00	0.00	0.00	0.02								
IR = 0	Q1	0.05	0.44	0.26	0.15	0.06	0.01	0.00	0.02								
IR = 0	Q2	0.02	0.21	0.54	0.21	0.01	0.00	0.00	0.01								
IR = 0	Q3	0.01	0.01	0.21	0.61	0.16	0.00	0.00	0.01								
IR = 0	Q4	0.00	0.00	0.02	0.21	0.62	0.14	0.00	0.00								
IR = 0	Q5	0.00	0.00	0.00	0.04	0.21	0.43	0.32	0.00								
IR = 0	Q6	0.00	0.00	0.00	0.01	0.05	0.15	0.78	0.00								
IR = 0	Jail									0.18	0.16	0.08	0.07	0.07	0.06	0.05	0.33
IR = 1	N									0.62	0.20	0.00	0.00	0.00	0.00	0.00	0.18
IR = 1	Q1									0.04	0.36	0.23	0.13	0.06	0.01	0.00	0.18
IR = 1	Q2									0.01	0.17	0.49	0.22	0.01	0.00	0.00	0.09
IR = 1	Q3									0.01	0.01	0.17	0.58	0.18	0.00	0.00	0.05
IR = 1	Q4									0.00	0.00	0.01	0.19	0.60	0.16	0.00	0.04
IR = 1	Q5									0.00	0.00	0.00	0.03	0.19	0.42	0.32	0.03
IR = 1	Q6									0.00	0.00	0.00	0.01	0.05	0.15	0.76	0.03
IR = 1	Jail									0.05	0.04	0.02	0.02	0.02	0.02	0.02	0.81

Source: Authors' calculations using data from NLSY79.
Note: Rows are indexed by the latent state at age 25, columns by the latent state at age 26. IR or "incarceration record" indicates previous incarceration. "Jail" denotes current incarceration. N denotes not employed but not currently incarcerated. Qi denotes earnings potential bin i.

graduates are nonemployed; the corresponding figure for White men is less than 17% (Figure 4).[19] Similarly, Figure 5 shows that over 20% of Black male high school graduates have been incarcerated at least once by age 55, compared to 3.3% of White high school graduates. Figures 4 and 5 also show that our estimates fit the underlying data quite well.

Finally, Figure 6 shows our estimates of mean earnings, found by simulating our earnings transitions and applying the bin-specific means estimated from the CPS. (Details on how we estimate the bin means can be found in Appendix E.) At any education level or sex, White individuals earn considerably more. This difference reflects both higher rates of employment and higher earnings when employed.

well-documented differences in health (see, e.g., Russo et al. 2024).

[19]Some of the differences among women arises because a married women is considered employed if the husband works (Appendix C), and White women are more likely to be married. As (Gordon et al., 2023, Table D.1) show, much (although not all) of the nonemployment differences among women disappears when nonemployment is measured at the individual level.

Table 3: Observation Probabilities: 25-year-old Black Male High School Graduate, Single, No Kids, No Incarceration Record

Latent State ↓		Observed Outcome: Incarceration Record = 0 N	Q1	Q2	Q3	Q4	Q5	Q6	Jail
IR = 0	N	1.00							
IR = 0	Q1	0.07	0.89	0.04	0.00	0.00	0.00	0.00	
IR = 0	Q2	0.03	0.16	0.63	0.16	0.01	0.00	0.00	
IR = 0	Q3	0.02	0.06	0.23	0.40	0.23	0.05	0.01	
IR = 0	Q4	0.01	0.00	0.03	0.22	0.51	0.19	0.03	
IR = 0	Q5	0.01	0.00	0.00	0.00	0.21	0.64	0.14	
IR = 0	Q6	0.01	0.00	0.00	0.00	0.01	0.41	0.57	
IR = 0	Jail								1.00

Source: Author's calculations using data from NLSY79.
Note: Rows are indexed by the latent state at age 25, columns by the observed state at the same age. IR or "Incarceration Record" indicates previous incarceration. "Jail" denotes current incarceration. N indicates not employed but not currently incarcerated. For rows, Qi denotes earnings potential bin i. For columns, Qj denotes observed earnings bin j.

5.2.3 Rates of Return

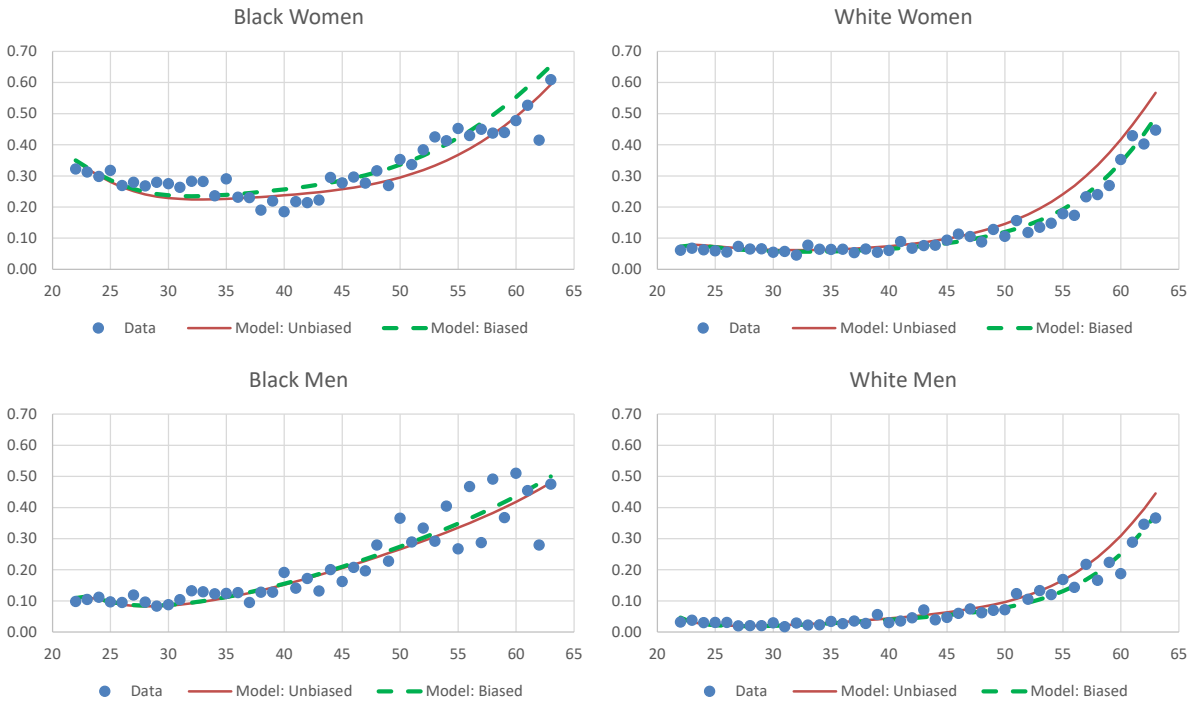
The final forcing process that we estimate from the data is the financial return on net worth $r_{i,t} = r(a_{it}, d_i, t, \rho_{i,t})$. We do this in three steps. First, we impute the portfolio share $s_{j,i,t}$ of each portfolio item j in a household's net worth using data from the SCF. Next we multiply each item's share by an item-specific return $r_{j,i,t}$. Then we add in an education-, sex- and wealth-based return adjustment $r^{d_i}(a_{i,t})$. Finally, we calculate the financial rate of return as $r_{i,t} = \sum_j s_{j,i,t} r_{j,i,t} + r^{d_i}(a_{i,t})$ Each of these steps is described in turn below, beginning with portfolio shares.

Table A1 in Appendix F provides a descriptive view of the SCF portfolio shares data. The first column lists the asset and debt categories that comprise net worth. Note that wealthier households are more heavily invested in high-yielding assets (such as stocks and private equity), less invested in low-yielding assets (such as cash and vehicles), and less indebted.[20] This pattern holds for both Black and White households. Notably, conditional on wealth, Black and White households have reasonably similar portfolios.

To formally estimate portfolio shares, we run a series of nonlinear regressions, using pooled data from all waves of the SCF between 1989 and 2019. In addition to a third-order spline in wealth, the regressions control for age, education, marital status and an

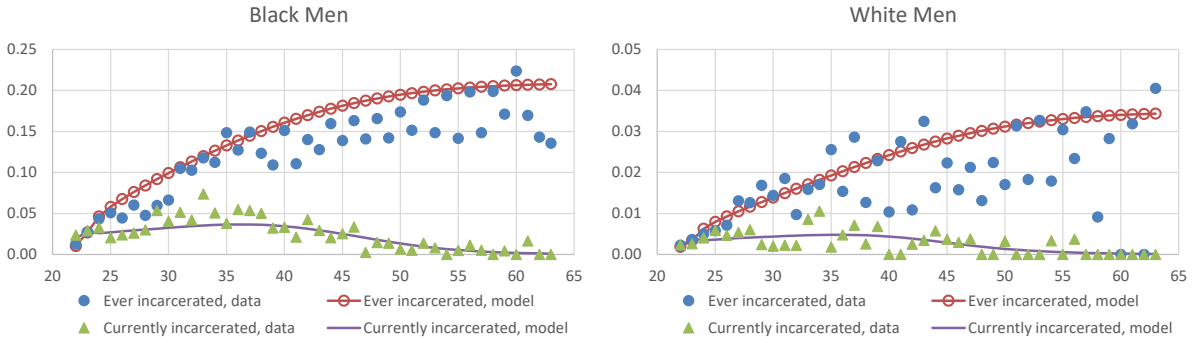
[20] This is a well-known relationship; see e.g. Campbell 2006

Figure 4: Nonemployment Rates, High School Graduates



Note: Authors' calculations using data from the NLSY79. Green lines are model predictions comparable to the nonemployment rates observed in the NLSY79. Red lines account for non-response bias.

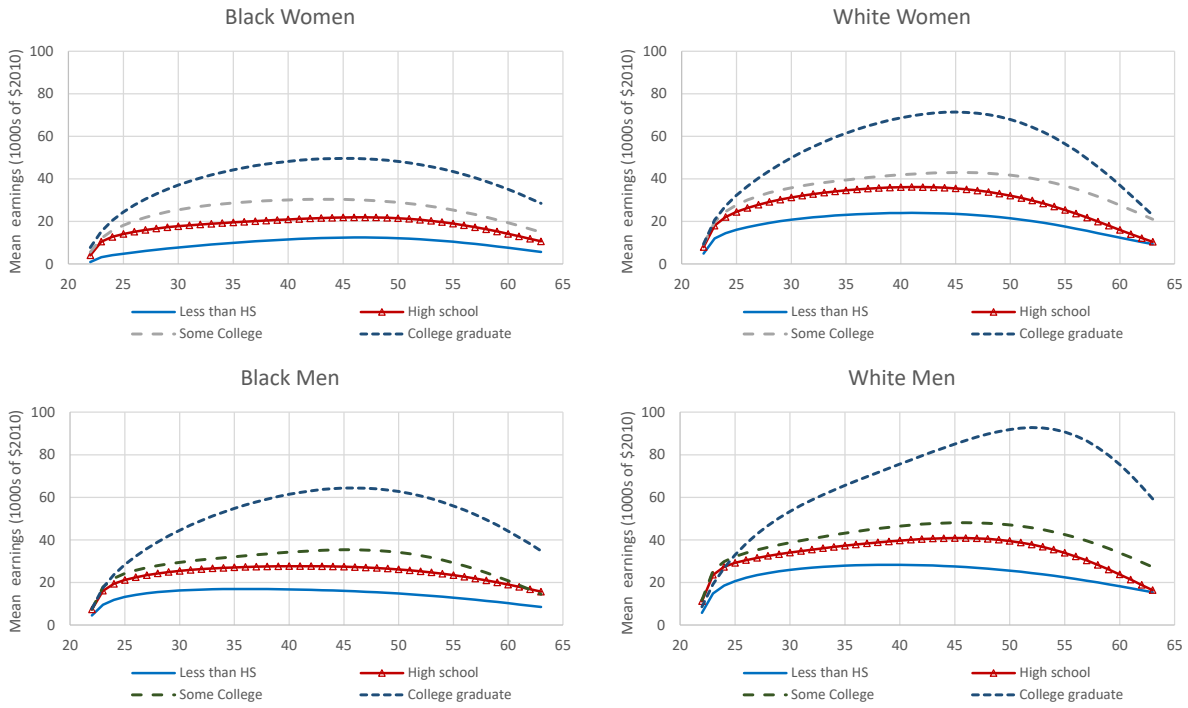
Figure 5: Incarceration Rates, Male High School Graduates



Note: Authors' calculations using data from the NLSY79.

indicator of whether the household head is White. Appendix F provides a full description of our specification and our estimation procedure. Table A2 of the same appendix shows the coefficient on the race (White) indicator for each portfolio item. In almost every case, White households hold more assets and less debt than Black households with the same covariates—the one exception is education debt. The table also shows, however, the

Figure 6: Mean Earnings by Race, Sex and Education



Note: Authors' calculations using data from the NLSY79 and the CPS.

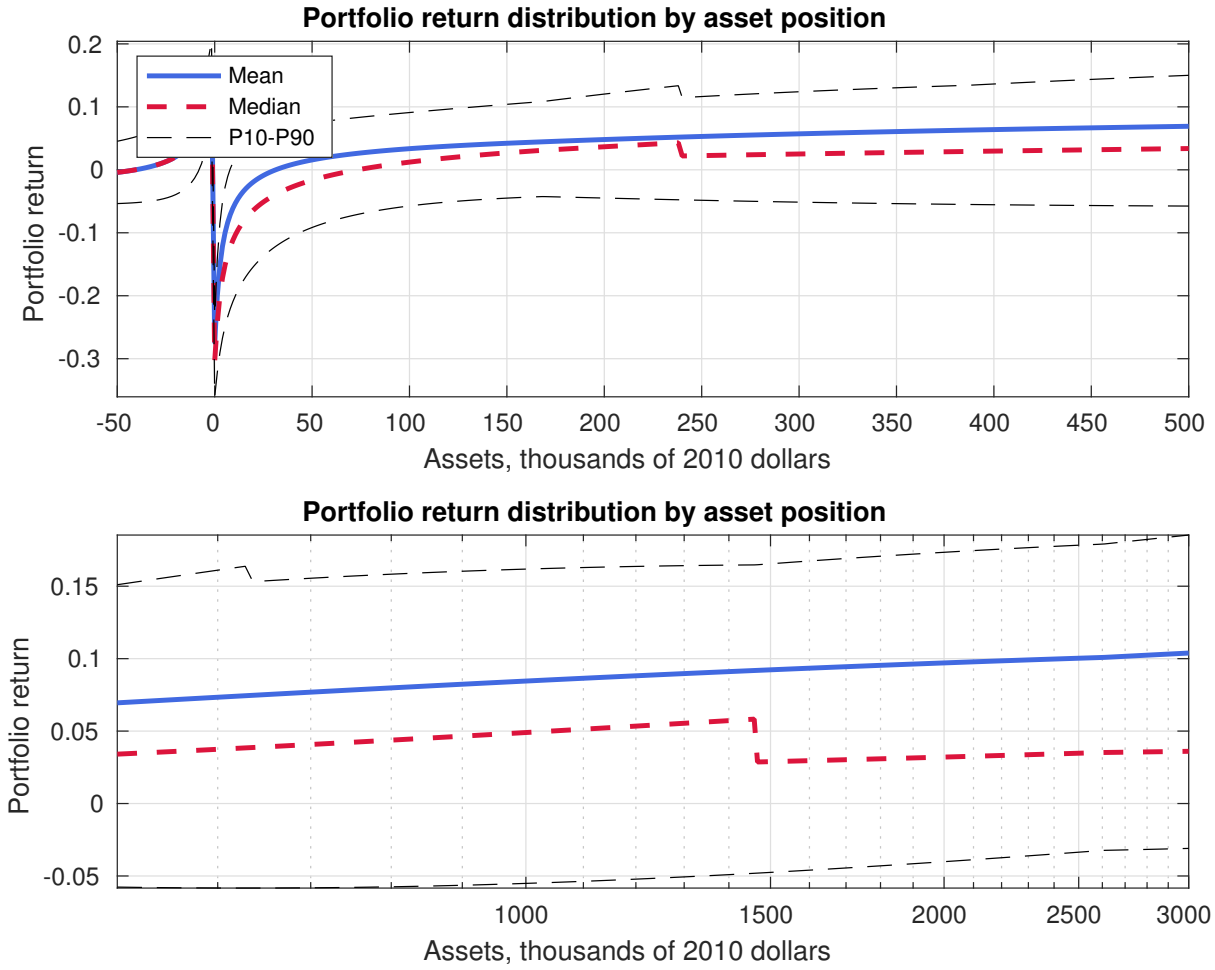
incremental gain in R^2 from adding the race indicator to the regression is uniformly small. Taken together, Tables A1 and A2 suggest that much of the variation in portfolio shares can be attributed to factors such as total wealth.

The second step is to obtain asset-specific rates of return. We assume—with the exception of housing—that these returns do not differ by race. Fagereng et al. (2020) report that, all else equal, men and highly-educated individuals earn higher overall rates of return; we use their estimates to construct $r^{d_i}(a_{i,t})$. For housing, we apply Kermani and Wong's (2021) estimate that the returns earned by Black homeowners on average are 2.32pp lower than those earned by White homeowners.

Finally, each individual's overall rate of return is found by multiplying their portfolio shares by the asset/debt-specific rates of return $\{r_{j,i,t}\}_j$ and adding the premium $r^{d_i}(a_{i,t})$. Return shocks are independent across time and individuals. We also impose a small scale premium in total returns, with the scale coefficient α_{sc} estimated along with the preference parameters. Appendix G describes the distributions and adjustments in detail.

Figure 7 shows the mean, median, 10th percentile and 90th percentile of the returns implied by our estimates. At low levels of net worth, returns are negative, as interest

Figure 7: Portfolio Returns as a Function of Net Worth



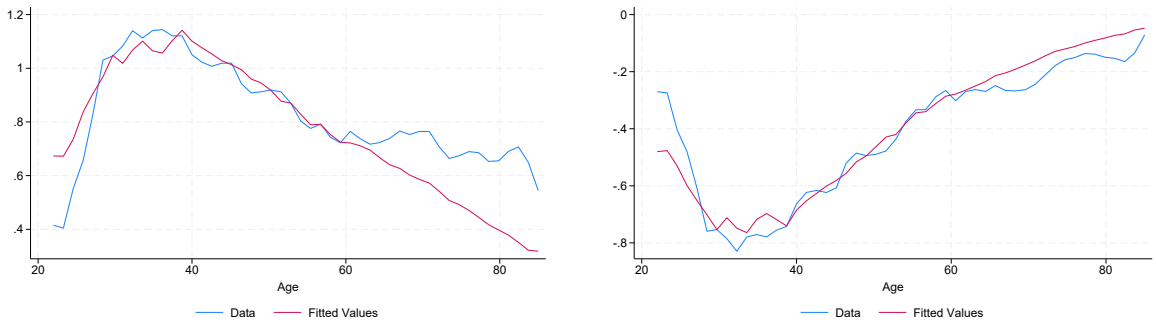
Note: Authors' calculations using the SCF and returns described in Appendices F and G.

rates on debt exceed returns on low-yielding assets. As wealth rises, so do mean returns. Households with more than $500,000 in wealth on average earn real returns of 7% per year or more. The yields on high-return assets tend to be riskier and more skewed, however, leading median returns to rise more slowly and at some point begin to decline in wealth. The richest individuals experience negative returns more than half of the time.

Our estimation also yields life-cycle profiles of household portfolios that are consistent with the data. As an example, Figure 8 shows the life cycle trajectories for housing shares, comparing fitted values from our estimates to the data. It shows that the portfolio shares for real estate and debt thus rise dramatically when individuals are in their late 20s, peak between ages 30 and 40, and fall significantly thereafter. This is consistent with individuals borrowing to buy homes when young, accumulating wealth as they pay down

their mortgages, and shifting their portfolios towards other assets as they retire.

Figure 8: Real Estate and Mortgage Shares by Age



(a) Real Estate (b) Mortgages

Note: Authors' calculations using the SCF. See Appendix F for details.

5.2.4 Inheritances and inter-vivos transfers received

The SCF queries respondents about the inheritances and gifts they have received in the past. Because individuals can receive multiple bequests over their lives, to calibrate our model we need a cumulative transfer measure. We thus rescale the transfers to the value they would have if the individual had received them at age 60, rather than the actual date of receipt, discounting at an annual rate of 3 percent. Summing the rescaled transfers yields lifetime totals, measured in age-60 dollars, at the time of the survey interview. The top panel of Appendix Table A4 shows the distribution of lifetime transfers found in the data. A well-known feature of the transfer distribution is that even as most individuals never receive any transfers whatsoever, some receive significant amounts. To fix ideas, among White individuals with bachelor's degrees, the median lifetime transfer is $6,000, but the mean exceeds $200,000. Black individuals receive much fewer and much smaller transfers. Among college graduates, the average lifetime transfer for Black individuals ($70,000) is roughly one third of that for White individuals.

To reduce the number of state variables in the model, we assume that intergenerational transfers are independent across age. The probability of receiving a transfer at any given age is assumed to be proportional to the probabilities found in the SCF. In contrast, we assume that when positive, transfers follow the same two-point distribution at every age. We estimate the parameters of the transfer process via simulation, targeting selected quantiles of total lifetime transfers, along with their mean. Appendix H describes our

calibration procedure in more detail, and Table A4 shows the quantiles implied by our estimated transfer process.

In our model initial wealth endowments play a role similar to that of intergenerational transfers and are likewise taken from the data. For each race-education-sex/marital status cell, we construct initial wealth distributions from the SCF: see Appendix I for details.

5.2.5 Taxes

Our model of taxes ($\tau_{i,t}$) builds off a specification dating back to Feldstein (1969) and recently utilized by Heathcote et al. (2017). Because a significant portion of asset income consists of capital gains, which are taxed at lower rates than other income and only when realized,[21] we modify this specification to account for the preferential tax treatment of asset income. This results in:

$$\tau_{i,t} = \max\{y^a_{i,t} + y^e_{i,t}, 0\} - \theta^0_{i,t}\left(\max\{y^a_{i,t} + y^e_{i,t}, 0\}\right)^{1-\theta^1_{i,t}}, \tag{12}$$

where $y^a_{i,t}$ and $y^e_{i,t}$ are the taxable components of asset and non-asset income, respectively.

Appendix J presents the parameters of our tax function and describes how we estimate them. We calculate tax liabilities (income and FICA) for estimation with SCF data, the NBER TAXSIM program (Feenberg and Coutts, 1993), and code substantively the same as that in Gale et al. (2022). Appendix Figure A4 shows the distribution of marginal tax rates for non-asset income (y^e) and wealth ($a_{i,t}$) associated with our tax specification.

5.3 Third stage: Matching the distribution of White wealth

5.3.1 Targets

In the final stage of our calibration, we find the parameter vector $\Theta = (\beta, \underline{h}, \kappa_b, \chi_b, N(\cdot), \underline{z}, \varsigma, \alpha_{\text{sc}})$ by targeting wealth means and quantiles for White individuals:

1. Mean wealth by age, household structure (single man, single woman, couple) and educational attainment.

2. 10th, 50th and 90th percentiles of wealth by age, household structure and education.

[21] In addition to capital gains, some forms of interest income receive favorable treatment. The double taxation of asset income through the corporate income tax is not relevant, as dividends and stock prices are both net of this tax.

3. The fraction of individuals with nonpositive wealth by by age, household structure and education.

The targets come from polynomial regressions on a synthetic panel of SCF data with cohort controls: our approach is similar to (if simpler than) the one for portfolio shares described in Appendix F. The aggregate profiles shown in Figure 1 are found using the same approach.

5.3.2 Simulation and estimation

We begin at age 22 with a set of individuals drawn from the data. These individuals differ by race, education and sex and are endowed with a household structure and earnings state drawn from our estimated initial distribution. We also endow them with wealth drawn from race-education-sex specific initial distributions constructed from the SCF: see Appendix I for details. We then give individuals draws of realized earnings, rates of return and medical expenses, determining their cash-on-hand, and simulate their lives forward. Each period, the model's decision rules, evaluated at an individual's current state, determine their consumption and saving. Saving decisions, along with draws from the estimated transition probabilities, then determine their updated state. We simulate individuals until they die, imposing an maximum age of 100, and calculate means and quantiles from the simulated data.

We set our parameters by minimizing

$$\sum_{qn,educ,f,age} \left(\frac{stat_{qn,educ,f,age} - \widehat{stat}_{qn,educ,f,age}}{\omega_{qn,educ,f,age}} \right),$$

where: $stat_{qn,educ,f,age}$ denotes the observed mean or quantile qn for an education-family composition-age subset of white individuals (family composition is one of {single man, single woman, married}); $\widehat{stat}_{qn,educ,f,age}$ denotes the simulated statistic; and ω is a weight. Our targets cover ages 22-80.

6 Wealth Differences: Model Fits and Decomposition

6.1 Parameter estimates

Appendix K presents our parameter estimates and discusses how they are identified in the data. Among the most important of these are the estimates of the discount factor

β. To help the model fit a diverse set of disaggregated targets, we allow β to differ by education and marital status. Our estimates (table A7) suggest that people with higher educational attainment tend to be more patient and that married people are more patient than singles. The discount factors for singles mostly lie between 0.8 and 0.85; the discount factors for couples are 12% (not pp) larger.

Another important parameter is α_{sc}, which controls the pure scale premium in returns. Our estimate of α_{sc} implies that, all else equal (including portfolio shares), the rate of return earned by a person with a net worth of \$2,000,000 is 1.755pp higher than the return earned by a person with a net worth of \$300,000.

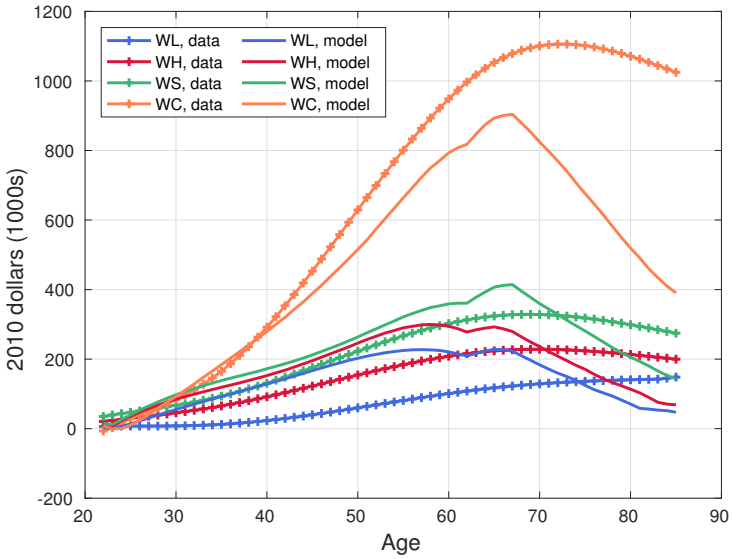
6.2 Model Fits

Figure 9 compares the model-predicted wealth profiles for White individuals to those found in the data. Although these profiles, aggregating across sex and marital status, are not our actual targets, they provide a sense of the model's fit of the data. Panel (a) of the figure shows that the model fits median wealth closely for most of the life cycle. The fit is even closer in panel (b), which shows results for the 90th percentile.

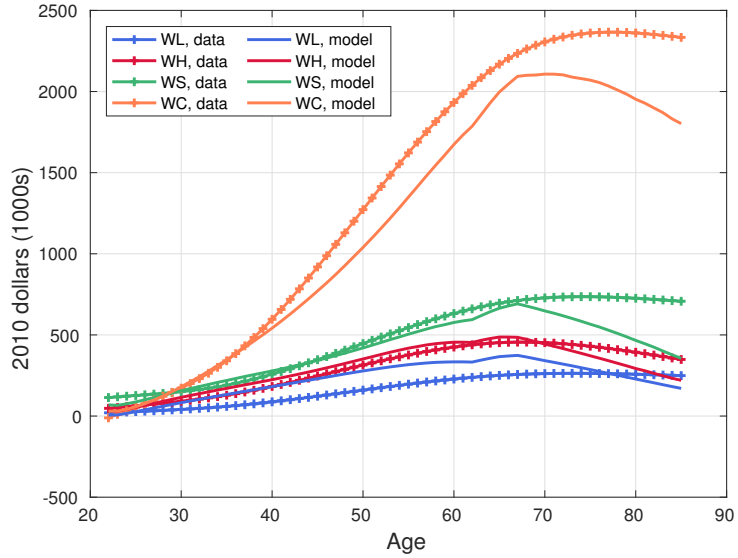
Figure 10 shows median wealth profiles for Black individuals. Given that we do not target Black individuals when calibrating the preference parameters, the model does an acceptable job. The main failing of the model is that it overstates the wealth holdings of those with college degrees. This suggests that our model, for all its richness, is still missing some mechanisms. It may be the case, for example, that affluent Black individuals are more likely to provide informal financial assistance to family and friends than their White counterparts (O'Brien, 2012). There is also evidence that Black workers utilize less of their employer-provided savings subsidies (Choukhmane et al., 2024).

Combining across years, we can pool the model-simulated histories into a cross-sectional distribution comparable to the pooled synthetic panel that we used to estimate the wealth profiles. Figure 11 compares the rank relationships implied by the two distributions. The model reproduces 81% (0.34/0.42) of the empirical rank Gini. The model likewise reproduces 86% (\$187,000/\$217,000; see Table 7 below) of the observed difference in mean wealth, 93% (\$63,600/\$68,500) of the difference in medians, and 118% (11.9pp/10.1pp) of the difference in the likelihood of non-positive wealth. These comparisons give us confidence in the model's predictions.

Figure 9: Wealth Profiles by Education, Data and Model, White Individuals



(a) Median



(b) Mean

6.3 Decompositions

Our first quantitative exercise is to allocate the wealth differences generated by the model among its endowments and forcing processes with the Shapley-Owen decomposition (Shorrocks, 2013). Letting $y = f(x_1, x_2, .., x_N)$ denote a nonlinear function, the Shapley-Owen decomposition measures the effects of removing each variable x_i, $i \in \{1, ..., N\}$, on y. It has two desirable properties: first, the sum of the N individual effects equal the

Figure 10: Median Wealth Profiles by Education, Data and Model, Black Individuals

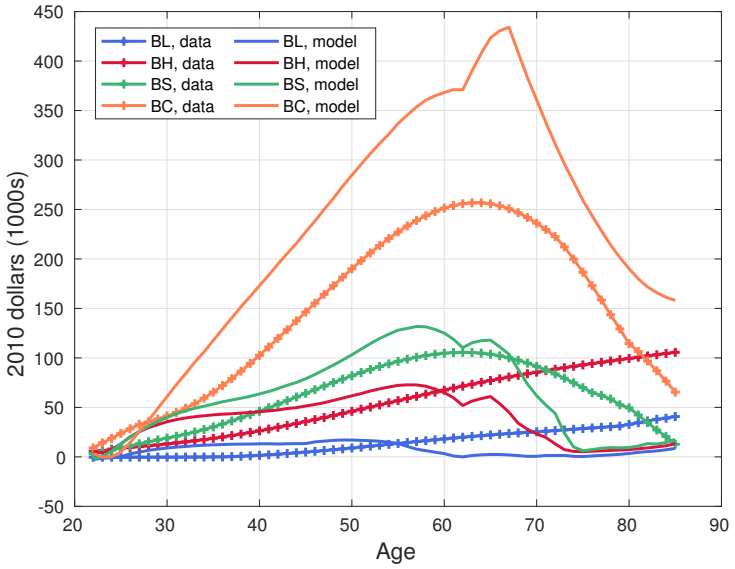
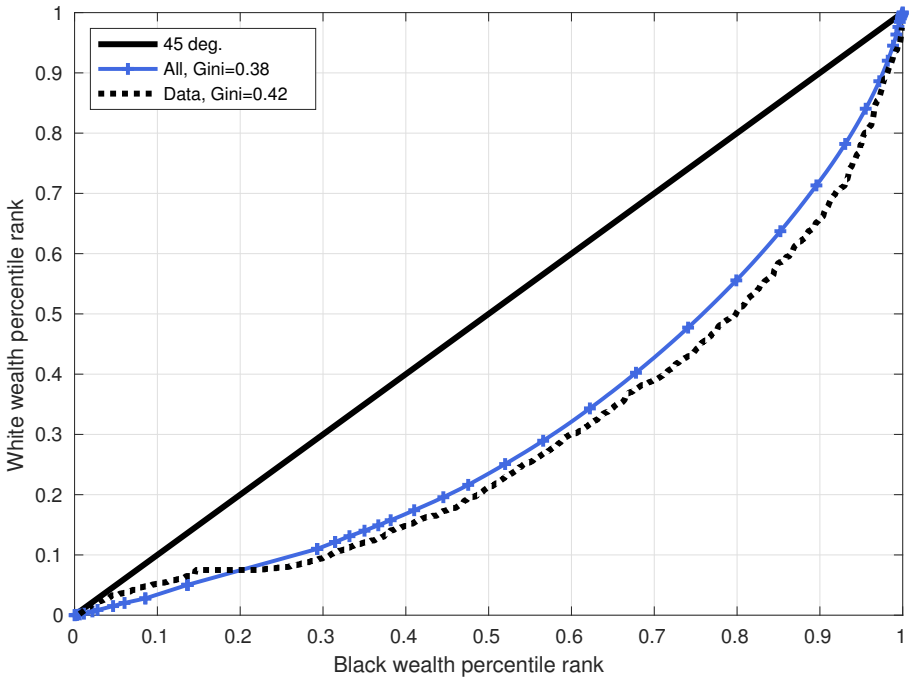


Figure 11: Rank Differences in Wealth: Data and Model



total difference y; and second, the decomposition does not depend on the order in which the variables are removed in the calculation. Audoly et al. (2024, Appendix E) provide a concise description. For our purposes, y is a model-predicted wealth difference and x_i is the difference in forcing process/endowment i, such as earnings or employment. If the set of forcing processes and endowments contained in $\{x_1, x_2, .., x_N\}$ is complete, setting all

values of x_i to zero by, say, giving Black individuals the endowments and forcing processes of White individuals, would make the Black and White wealth distributions identical. The Shapley-Owen algorithm thus provides a complete decomposition of the model-generated wealth differences.

Table 4 shows decomposition results for the mean and median differences, the difference in the fraction of households with nonpositive wealth, and the rank Gini coefficient. Each line of the table shows how the wealth differences change, as percentages of the total differences, after replacing a particular set of processes/endowments for Black individuals with those of White individuals.

Table 4: Shapley-Owen Decomposition of Wealth Differences (in Percent)

	Difference in means	Difference in medians	Fraction net worth ≤ 0	Rank Gini coefficient
Portfolios, taxes and medical expenses	5.3	12.1	6.3	9.3
Education	20.9	10.5	6.8	12.1
Other demographics	2.7	-0.6	-2.2	-0.6
Wealth transfers and initial wealth	11.7	7.3	8.8	8.3
Earnings when employed	19.0	18.5	13.4	18.6
Employment	40.3	52.2	66.8	52.2

Source: Author calculations described in text.

The largest shares are for employment, especially at the bottom of the distribution.[22] Giving Black individuals the nonemployment and incarceration rates of whites would reduce the nonpositive fraction difference by 62%. But even for means, which depend in large part on the right tail of the distribution, nonemployment and incarceration make up 30% of the difference.

The other major contributors are educational attainment and earnings when employed. These factors are most important for the mean difference, together accounting for 46% of the total. A key takeaway from Table 4, then, is that the bulk of the wealth differences can be attributed to earnings, employment or education. Moreover, the differences in earnings and employment considered in Table 4 are the differences that remain *after* controlling for education; as Table 5 shows, the contribution of education to the differences

[22]Given the complexity of the earnings process (see Appendix D), our calculation involves giving Black individuals the entire distribution of employment outcomes, including (positive) earnings bins, of White individuals. Fortunately, the bins are estimated separately for each race-education-sex-age group and based on a common set of percentile ranges (Appendix E).

is mediated in part through its effects onemployment and employment, especially at the mean. (Education also operates through preference parameters and demographics.) Our results are thus consistent with papers such as Ashman and Neumuller (2020) or Aliprantis et al. (2022), which show that models with non-homothetic saving behavior attribute large shares of the wealth differences to differences in income.

Table 5: Shapley-Owen Decomposition of Earnings Differences (in Percent)

	Difference in means	Difference in medians	Fraction nonemployed	Rank Gini coefficient
Portfolios, taxes and medical expenses	0.0	0.0	0.0	0.0
Education	22.4	13.4	13.8	18.9
Other demographics	-0.3	-0.8	-13.8	-0.7
Wealth transfers and initial wealth	0.0	0.0	0.0	0.0
Earnings when employed	37.8	33.0	0.0	35.3
Nonemployment & incarceration	40.1	54.4	100	46.4

Source: Author calculations described in text.
Note: Mean and median earnings differences are inclusive of nonworkers.

Another notable result is that explicit racial differences in returns, due to race-based differences in portfolios or race-specific returns on housing, make relatively small contributions, the largest being 15% for the median difference. (See the first row; these shares also include contributions from other items.) Within the model, most of the Black-White differences in portfolio shares and returns can be attributed to factors, such as total wealth or education, that affect both races similarly.

Intergenerational transfers and initial wealth endowments also play modest roles, accounting for between 9% and 12% of each difference. This is roughly consistent with Sabelhaus and Thompson's (2023) finding that inheritances and intervivos transfers explain about 14-16% of racial wealth differences. One reason that intergenerational transfers have limited contributions is that we assume individuals know the distribution of transfers conditional on their race and education, and adjust their saving accordingly. This contrasts with Ashman and Neumuller (2020), who assume that all transfers are unanticipated and find that they account for 26% of the median wealth difference. Given that most individuals' lifetime earnings are much larger than their initial wealth, it is not surprising that initial wealth endowments play a modest role as well.

Differences in mortality, marital status, and children ("other demographics") make the

smallest contribution. Because it is *changes* in marital composition that affect saving—consider the Euler equation—the relatively slow evolution of these states (see Figure 3) likely limits their role.

Table 6: Shapley-Owen Decomposition of Wealth Differences (in Percent)

	Difference in means	Difference in medians	Fraction net worth ≤ 0	Rank Gini coefficient
		Panel A. Age 35		
Portfolios, taxes and medical expenses	4.7	10.9	8.6	9.1
Education	14.3	8.0	5.4	9.6
Other demographics	-1.3	-2.8	-2.7	-2.5
Wealth transfers and initial wealth	19.5	6.7	3.7	7.8
Earnings when employed	19.4	21.9	14.7	22.1
Employment	43.3	55.2	70.3	53.9
		Panel B. Age 70		
Portfolios, taxes and medical expenses	6.4	14.2	2.4	10.5
Education	22.9	14.2	11.2	14.8
Other demographics	0.5	-1.4	-10.8	-1.9
Wealth transfers and initial wealth	10.4	7.1	7.0	8.5
Earnings when employed	19.5	17.7	20.3	18.8
Employment	40.3	48.2	70.0	49.3

Source: Author calculations described in text.

To examine the life-cycle dimension, Table 6 decomposes wealth differences by race for the age-35 (panel A) and age-70 (panel B) distributions of wealth. Relative to the all-ages cross-section (Table 4), at age 35 education plays a smaller role, while earnings and employment play larger ones. Also notable is that at age 35, intergenerational transfers and initial wealth account for 22% of the mean wealth difference, as opposed to the all-ages share of 12%; initial conditions matter more at earlier ages. It is nonetheless the case that employment, earnings and education account for most of the differences. The decomposition and age 70 is even more akin to the overall decomposition, except for the fraction with nonpositive wealth. Given that large wealth holdings are more likely at older ages, such a close correspondence is not surprising.

7 Model Mechanisms

The decompositions in the previous section showed that the lion's share of Black-White wealth differences are attributable to differences in employment, earnings and education. Employment differences play large roles across the wealth distribution, but especially at the bottom; earnings when employed and education are most important at the top. In order for labor income to play such a dominant role, saving within the model must be non-homothetic. We now examine several of the mechanisms that make this the case. Recall that even though these mechanisms expand differences between races, they are not explicitly race-based.

7.1 Individual mechanisms

In Table 7, we examine the effects of shutting down selected model mechanisms. These analyses differ fundamentally from the decomposition exercises, where we systematically eliminated Black-White differences in endowments/forcing processes by giving Black individuals the endowments/forcing processes of White individuals. In this section, we alter model mechanisms while retaining its race-specific endowments and forcing processes.

The first experiment in Table 7 is to have every individual earn a constant return of 3%.[23] This reduces the mean wealth difference by $111,000, or 60% of the model-predicted total. Notably, the relative change in the *proportional* difference, 27% (24.8pp/92pp), is less than half the relative change in the level difference. The divergent results arise because eliminating scale effects reduces the wealth holdings of both races, reducing both the numerator (the level difference) and the denominator (average wealth holdings) of the proportional difference, so that the level and proportional differences to move in distinct ways. All of this contrasts markedly with the decomposition exercises, where the wealth holdings of White individuals are held fixed while the endowments and predicted wealth of Black individuals are altered, causing the numerator and denominator of the proportional differences to always move in opposite directions—and causing the level and proportional differences to always move in the same direction. Given these mechanics, it is not surprising that the reduction in the rank Gini (0.07 out of 0.34) is much smaller than the reductions for the mean or median level differences. As Bayer and Charles (2018) stress, factors that stretch or compress the overall wealth distribution, such as scale effects in rates of return, can expand or reduce level differences without inducing large changes

[23] This change is not completely race-independent, as it eliminates explicit racial differences in portfolio shares and housing returns. As shown in the previous section, however, the latter play fairly modest roles.

in relative rank.

Table 7: Counterfactual Experiments: Model Mechanisms

	Means		Medians		NW ≤ 0	Rank
	($1000s)	(%/ppts)	($1000s)	(%/ppts)	(%/ppts)	Gini
Data difference	217	116	68.5	124	10.1	0.42
Baseline model difference	177	94.9	70.0	123	13.0	0.38
Change from baseline model difference:						
No portfolio/RoR heterogeneity	-95.9	-23.5	-29.2	-46.8	-2.76	-0.07
No bequest motive	-5.63	8.77	-4.01	20.3	2.29	0.03
No bequest motive or portfolio/RoR heterogeneity	-95.4	-12.9	-29.2	-30.3	-1.62	-0.04
Lower consumption floor	0.6	-5.02	-0.10	-23.7	-3.46	-0.01

Source: Author calculations described in text.

Note: Percentage differences for mean and median wealth calculated as $\frac{X^W - X^B}{0.5X^W + 0.5X^B}$.

The second row of Table 7 shows the effects of shutting down bequest motives. The mean and median differences close by $25,000 and $14,000, respectively, in level terms, and *increase* by 11pp and 28pp, respectively, in proportional terms. As with the return scale effects, eliminating bequest motives reduces wealth holdings across both races, in this case moving the level and proportional differences in opposite directions.

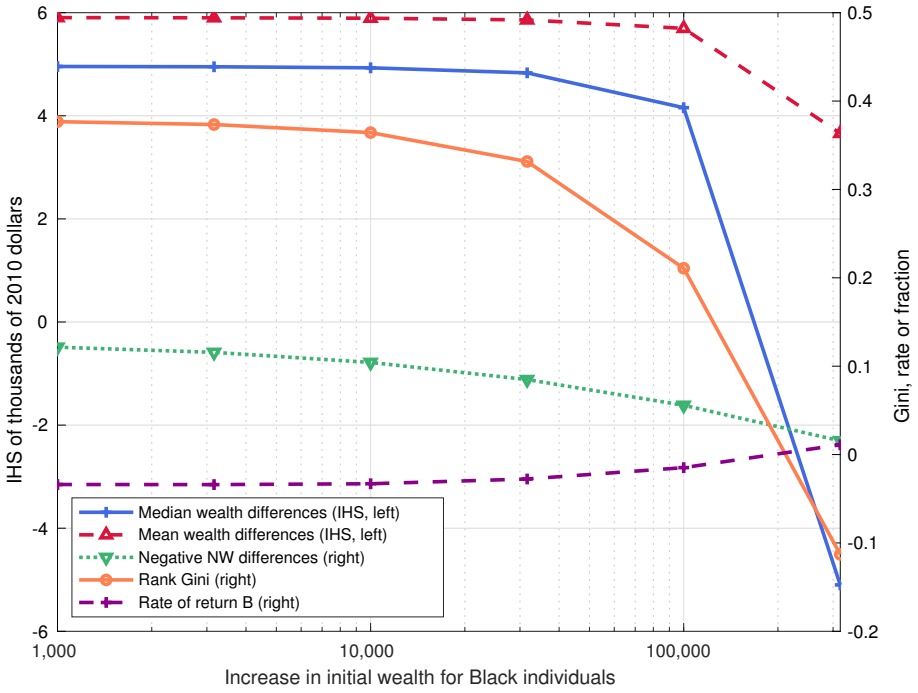
The last row of Table 7 examines the effect of reducing the consumption floor by decreasing the floor for singles, $\underline{z}$, from $2,900 to $2,000. Given that mean-tested transfers discourage the saving of poorer individuals (Hubbard et al., 1995), it is not surprising that reducing the consumption floor has the biggest effects at the bottom of the wealth distribution, decreasing the difference in nonpositive wealth by 10%. Reducing the consumption floor also reduces the proportional difference in medians by 10%, but by increasing the denominator—the level difference in the median (the numerator) actually increases.

7.2 Initial Wealth

Another way to quantify the model's mechanisms to examine how increases to the initial wealth of Black individuals affect wealth differences. Figure 12 shows how the differences shrink as Black wealth endowments increase. Because part of the increased endowment will be consumed, the transfers needed to close the mean and median wealth differences

are considerably larger than the differences themselves.[24] Eliminating the mean wealth difference of $187,000 requires giving Black individuals an incremental wealth endowment of around $300,000. Moreover, the first $100,000 of additional endowments reduces the mean difference by less than $40,000.[25] Because wealth is a luxury good, the marginal propensity of Black individuals to save out of their initial wealth increases markedly in wealth itself. One reason wealth is a luxury is the presence of scale effects in rates of return. As Figure 12 shows, the scale effects are substantial.

Figure 12: The Role of Black Wealth Endowments: Cross-sectional Differences



Note: The left-side vertical axis gives the Inverse Hyperbolic Sine (IHS) of the mean and median wealth differences.

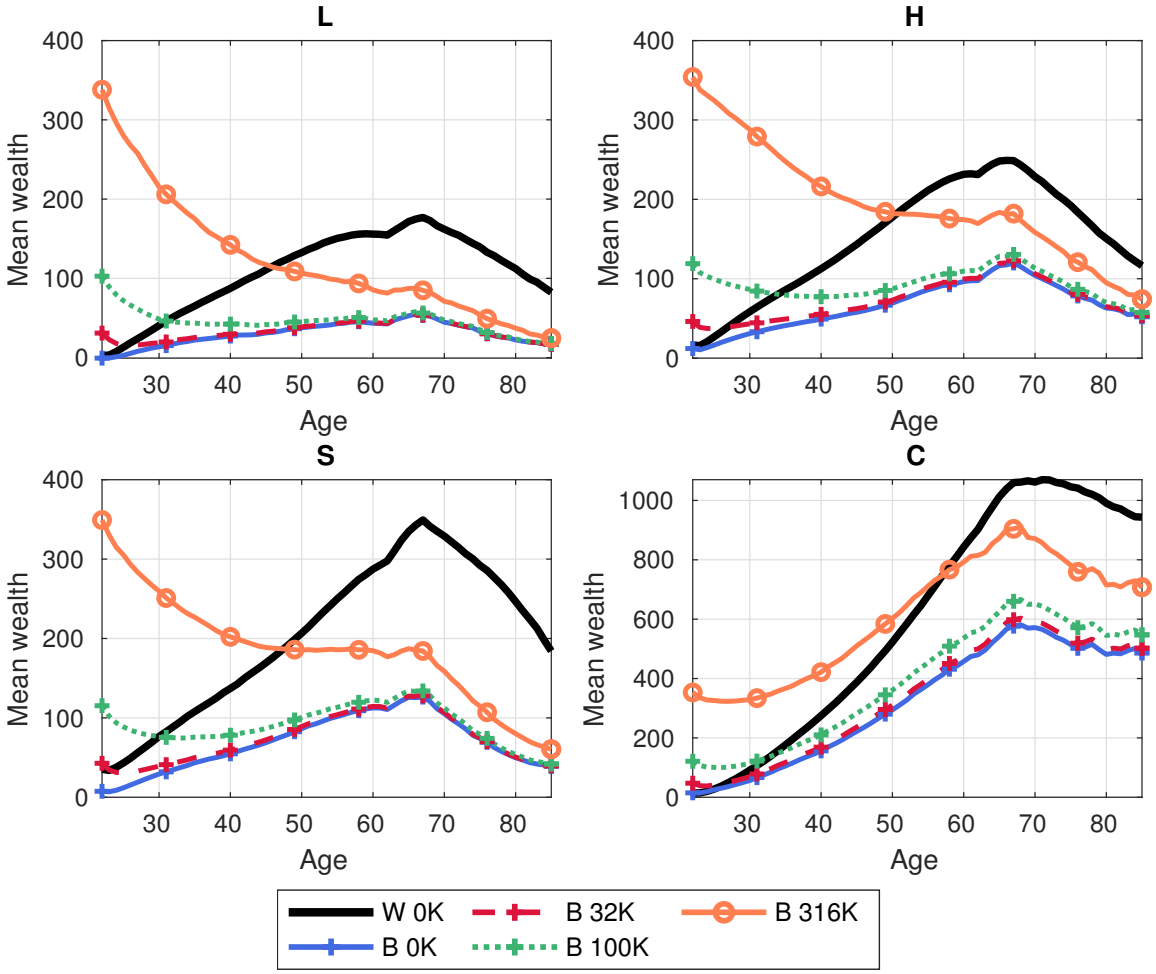
Recall that the cross-section in 12 is constructed from a single cohort. Figure 13 uses mean wealth profiles to show how Black individuals manage their incremental wealth endowments over their life cycles. Except for those with bachelor's degrees, most of the additional wealth is spent by age 60, and most of what remains at age 60 is spent dur-

[24] In our framework, saving is the only margin along which individuals adjust to these increments: labor supply and human capital are both exogenous. Moreover, we do not model the increments' funding. The exercises in this section and the next should thus be interpreted as model illustrations, rather than policy assessments.

[25] The left-side vertical axis of this graph gives the Inverse Hyperbolic Sine (IHS) of the median and median wealth differences, where $\text{IHS}(a) = \ln\left(a + \sqrt{a^2+1}\right)$ is a generalization of the natural log. See Appendix F.

ing retirement. Future cohorts will receive smaller wealth boosts than the one receiving the initial increment. In contrast, individuals with bachelor's degrees spend down only a small share of their additional wealth. This reflects the higher lifetime earnings of college graduates, which makes savings more of a luxury good. Nonetheless, in the aggregate our findings echo those of Aliprantis et al. (2022): they conclude, using an overlapping generations framework, that "[o]ne-time wealth transfers have only transitory effects unless they address the racial earnings gap."

Figure 13: The Role of Black Wealth Endowments over the Life Cycle



Note: "L", "H", "S" and "C" denote individuals without a high school degree, high school graduates, college attendance without a bachelor's degree, and bachelor's degree or higher, respectively. "B" and "W" denote Black and White individuals, respectively, while "xK" denotes the incremental wealth endowment in $1000s.

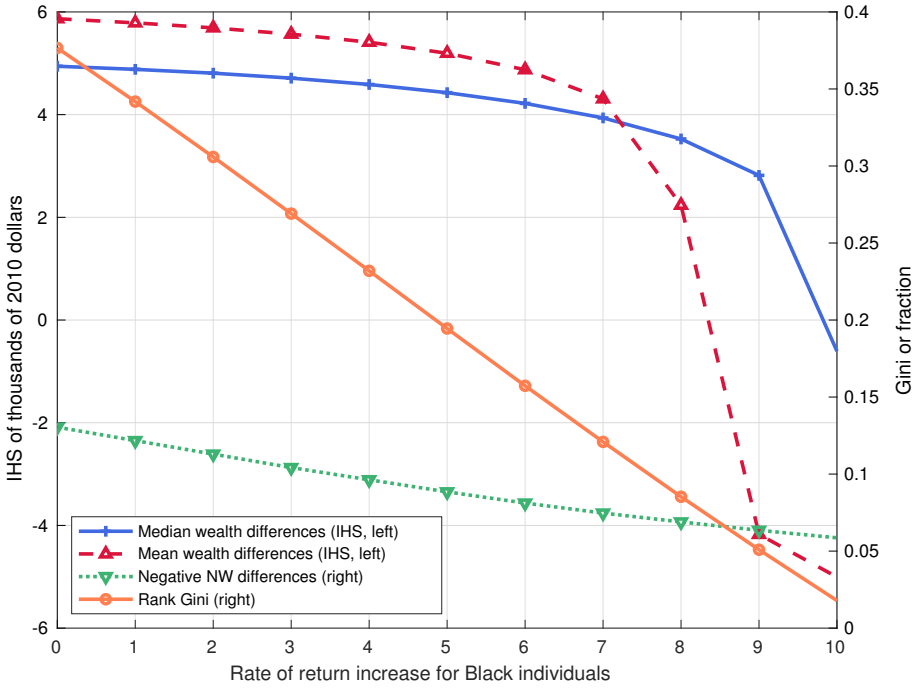
Given the small holdings of those at the bottom of the wealth distribution, endowment increments have bigger effects on the differences there. Increasing Black wealth endow-

ments by $100K reduces the level difference in medians by 60% and cuts the rank Gini in half.

7.3 Return subsidies

Figure 14 shows the effects of increasing the rate of return Black individuals earn on their wealth. Consistent with the preceding results, higher returns have the biggest effect at the top of the distribution. While a 6pp wealth increment almost completely eliminates the mean difference, the other differences remain substantial, even at higher increments. Taken together, Figures 12 and 14 reinforce the main messages of Tables 4 and 7. Earnings matter the most at the bottom of the wealth distribution, while rates of return matter more at the top.

Figure 14: The Role of Rates of Return: Cross-sectional Differences



Note: The left-side vertical axis gives the Inverse Hyperbolic Sine (IHS) of the mean and median wealth differences.

8 Conclusion

In this paper, we measure the contribution of a variety of forcing processes and endowments to observed wealth differences between Black and White households from the same

birth cohort. We are particularly focused on whether these factors differ over the life cycle and across the wealth distribution within this cohort.

Our approach involves, first, writing down a detailed life-cycle consumption savings model with an array of forces that includes employment and earnings dynamics, household composition dynamics, intergenerational transfers, mortality, among others. We calibrate the model to to match the life-cycle wealth profiles of White individuals disaggregated by education, sex, and marital status. Inputs to our calibration include exogenous forcing processes carefully estimated from data on White individuals and households. When we replace these exogenous processes with those estimated from data on Black individuals and households, our model largely reproduces the racial wealth heterogeneity observed in the data. We are thus able to allocate most of Black-White differences in wealth to endowments and forces present in our model. We employ Shapley-Owen decompositions for this purpose.

Our decomposition results show that different factors dominate across the wealth distribution. Differences in employment dominate at the bottom of the distribution, while differences in education and the earnings of the employed play larger roles at the top. Factors such as intergenerational transfers and demographics play relatively small roles throughout.

We also find that the absolute sizes of wealth differences depend on model mechanisms that are independent of race. The most important of these is the tendency of wealthier individuals to earn higher rates of return, regardless of race, which amplifies differences at the top of the wealth distribution. These scale differences—which are largely independent of race—are necessary for our model to fully account for the differences in Black and White wealth observed in our data.

By focusing on a single cohort, we have been able to estimate the forcing processes and financial mechanisms behind our results with a fair bit of precision. A fruitful area for future research is to continue examining these objects, bringing new data to bear, especially among different cohorts.

References

Aliprantis, D., D. R. Carroll, and E. R. Young (2022): "The Dynamics of the Racial Wealth Gap," Federal Reserve Bank of Cleveland Working Paper.

Altonji, J. G. and R. M. Blank (1999): "Race and Gender in the Labor mMrket," in *Handbook of Labor Economics, vol. 3*, ed. by O. Ashenfelter and D. Card, Elsevier, 3143–3259.

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

Audoly, R., R. McGee, S. Ocampo-Diaz, and G. Paz-Pardo (2024): "The Life-cycle Dynamics of Wealth Mobility," Staff Report 1097, Federal Reserve Bank of New York.

Bartolucci, F., A. Farcomeni, and F. Pennoni (2010): "An Overview of Latent Markov Models for Longitudinal Categorical Data," *arXiv preprint arXiv:1003.2804*.

Bayer, P. and K. K. Charles (2018): "Divergent Paths: A New Perspective on Earnings Differences Between Black and White Men Since 1940," *The Quarterly Journal of Economics*, 133, 1459–1501.

Benhabib, J. and A. Bisin (2018): "Skewed Wealth Distributions: Theory and Empirics," *Journal of Economic Literature*, 56, 1261–1291.

Benhabib, J., A. Bisin, and M. Luo (2019): "Wealth Distribution and Social Mobility in the US: A Quantitative Approach," *American Economic Review*, 109, 1623–1647.

Board of Governors (2023): "2022 Survey of Consumer Finances," Tech. rep., Board of Governors of the Federal Reserve System.

Boerma, J. and L. Karabarbounis (2023): "Reparations and Persistent Racial Wealth Gaps," *NBER Macroeconomics Annual*, 37, 171–221.

Bozio, A., C. Emmerson, C. O'Dea, and G. Tetlow (2017): "Do the Rich Save More? Evidence from Linked Survey and Administrative Data," *Oxford Economic Papers*, 69, 1101–1119.

Brandsaas, E. E. (2024): "Illiquid Homeownership and the Bank of Mom and Dad," Mimeo.

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

Campbell, J. Y. (2006): "Household Finance," *The Journal of Finance*, 61, 1553–1604.

Card, D., C. Dobkin, and N. Maestas (2008): "The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare," *American Economic Review*, 98, 2242–2258.

Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010): "Quantile and probability curves without crossing," *Econometrica*, 78, 1093–1125.

Choukhmane, T., J. Colmenares, C. O'Dea, J. L. Rothbaum, and L. D. Schmidt (2024): "Who Benefits from Retirement Saving Incentives in the US? Evidence on Gaps in Retirement Wealth Accumulation by Race and Parental Income," Working Paper 32843, National Bureau of Economic Research.

Conesa, J. C., D. Costa, P. Kamali, T. J. Kehoe, V. M. Nygard, G. Raveendranathan, and A. Saxena (2018): "Macroeconomic Effects of Medicare," *The Journal of the Economics of Ageing*, 11, 27–40.

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

De Nardi, M. and G. Fella (2017): "Saving and Wealth Inequality," *Review of Economic Dynamics*, 26, 280–300.

De Nardi, M., E. French, J. B. Jones, and R. McGee (2023): "Why Do Couples and Singles Save during Retirement? Household Differences and their Aggregate Implications," Working Paper w28828, National Bureau of Economic Research.

Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–1248.

Derenoncourt, E., C. H. Kim, M. Kuhn, and M. Schularick (2022): "Wealth of Two Nations: The US Racial Wealth Gap, 1860-2020," Working Paper 30101, National Bureau of Economic Research.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004): "Do the Rich Save More?" *Journal of Political Economy*, 112, 397–444.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020): "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 88, 115–170.

Farmer, L. E. (2021): "The Discretization Filter: A Simple Way to Estimate Nonlinear State space Models," *Quantitative Economics*, 12, 41–76.

Feenberg, D. and E. Coutts (1993): "An Introduction to the TAXSIM Model," *Journal of Policy Analysis and Management*, 12, 189–194.

Feldstein, M. S. (1969): "The Effects of Taxation on Risk Taking," *Journal of Political Economy*, 77, 755–764.

Fulford, S. L. (2015): "How Important is Variability in Consumer Credit Limits?" *Journal of Monetary Economics*, 72, 42–63.

Gale, W. G., S. Joshi, C. Pulliam, and J. Sabelhaus (2022): "Simulating Income Tax Liabilities in the Survey of Consumer Finances," Tech. rep., Tax Policy Center.

Gordon, G. and A. Hedlund (2019): "Accounting for the Rise in student Tuition," in *Education, Skills, and Technical Change: Implications for Future US GDP Growth*, ed. by C. R. Hulten and V. A. Ramey, University of Chicago Press, 357–394.

Gordon, G., J. B. Jones, U. Neelakantan, and K. Athreya (2023): "Incarceration, Employment, and Earnings: Dynamics and Differences," *Review of Economic Dynamics*, 51, 677–697.

Hamilton, J. D. (1994): *Time Series Analysis*, Princeton University Press.

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 132, 1693–1754.

Houle, J. N. (2014): "Disparities in Debt: Parents' Socioeconomic Resources and Young Adult Student Loan Debt," *Sociology of Education*, 87, 53–69.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995): "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103, 360–399.

Huggett, M. and G. Ventura (2000): "Understanding Why High Income Households Save More than Low Income Households," *Journal of Monetary Economics*, 45, 361–397.

Jones, J. B. and U. Neelakantan (2022a): "How Big Is the Inheritance Gap Between Black and White Families?" *Richmond Fed Economic Brief*, 22.

——— (2022b): "A More Comprehensive Measure of the Black-White Wealth Gap," *Richmond Fed Economic Brief*, 22.

——— (2023): "Portfolios Across the U.S. Wealth Distribution," *Richmond Fed Economic Brief*, 23.

Kermani, A. and F. Wong (2021): "Racial disparities in housing returns," Working Paper 29306, National Bureau of Economic Research.

Krusell, P. and A. A. Smith, Jr (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

Kumaraswamy, P. (1980): "A Generalized Probability Density Function for Double-Bounded Random Processes," *Journal of Hydrology*, 46, 79–88.

Lang, K. and J.-Y. K. Lehmann (2012): "Racial Discrimination in the Labor Market: Theory and Empirics," *Journal of Economic Literature*, 50, 959–1006.

Lockwood, L. M. (2018): "Incidental Bequests and the Choice to Self-insure Late-life Risks," *American Economic Review*, 108, 2513–2550.

Nakajima, M. and I. A. Telyukova (2017): "Reverse Mortgage Loans: A Quantitative Analysis," *The Journal of Finance*, 72, 911–950.

O'Brien, R. L. (2012): "Depleting Capital? Race, Wealth and Informal Financial Assistance," *Social Forces*, 91, 375–396.

Quadrini, V. (1999): "The Importance of Entrepreneurship for Wealth Concentration and Mobility," *Review of income and Wealth*, 45, 1–19.

Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyerand, J. Pacas, and M. Sobek (2020): "IPUMS USA: Version 10.0 [Dataset]," Https://doi.org/10.18128/D010.V10.0.

Russo, N., R. McGee, M. D. Nardi, M. Borella, and R. Abram (2024): "Health Inequality and Economic Disparities by Race, Ethnicity, and Gender," in *Society for Economic Dynamics, 2024 Annual Meeting*.

Sabelhaus, J. and J. P. Thompson (2023): "The Limited Role of Intergenerational Transfers for Understanding Racial Wealth Disparities," Current Policy Perspectives 95748, Federal Reserve Bank of Boston.

Sabelhaus, J. and A. H. Volz (2024): "Wealth Inequality and Retirement Preparedness: A Cross-Cohort Perspective," in *Real-World Shocks and Retirement System Resiliency*, ed. by O. S. Mitchell, J. Sabelhaus, and S. Utkus, Oxford University Press.

Sakong, J. (2022): "Cyclical Transactions and Wealth Inequality," Working Paper 2022-05, Federal Reserve Bank of Chicago.

Sasson, I. (2016): "Diverging Trends in Cause-specific Mortality and Life Years Lost by Educational Attainment: Evidence from United States Vital Statistics Data, 1990-2010," *PloS ONE*, 11, e0163412.

Scott, S. L. (2002): "Bayesian Methods for Hidden Markov Models: Recursive Computing in the 21st Century," *Journal of the American Statistical Association*, 97, 337–351.

Shorrocks, A. (2013): "Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value," *The Journal of Economic Inequality*, 11, 99–126.

Social Security Administration (2018): "The Long-Range Demographic Assumptions for the 2018 Trustees Report," Technical appendix.

Survey Research Center (1992): "A Panel Study of Income Dynamics: Procedures and Tape Codes, 1989 Interviewing Year; Wave XXII: A Supplement," Tech. rep., Institute for Social Research, University of Michigan.

Wolff, E. N. (2018): "The Decline of African-American and Hispanic Wealth since the Great Recession," Working Paper 25198, National Bureau of Economic Research.

——— (2022): "Heterogenous Rates of Return on Homes and Other Real Estate: Do the Rich Do Better? Do Blacks Do Worse?" *NBER Working Paper*, 30543.

Appendices for Online Publication

A Constructing the Net Worth Profiles

Our wealth data come from the 1989 - 2022 waves of the Survey of Consumer Finances (SCF), a repeated cross-section updated every three years by the Federal Reserve Board (Board of Governors, 2023). The SCF includes detailed information on the balance sheets, income and demographics of U.S. families. The SCF oversamples the wealthy; we use the sample weights that accompany the data.

Our measure of wealth is net worth: the value of a family's financial and nonfinancial assets net of all its debt. At each SCF wave we drop the (weighted) top 0.5% of observations. Our measure of wealth excludes defined benefit pension wealth; our model explicitly accounts for Social Security benefits, and private defined benefit pensions are rare for our cohort.

The smoothed profiles come from a set of regressions on individuals born between 1933 and 1988 with cohort controls. Wealth is expressed on an individual basis, with the wealth holdings of couples split evenly. In addition to means, medians and nonpositive-wealth fractions, we target the 10th, 25th, 75th and 90th percentiles of wealth over the life cycle. Because we use a fairly sparse cubic spline to fit these curves, they sometimes cross, especially at very young and very old ages. To make the curves monotonic in rank, we employ a variant of the strategy developed by Chernozhukov et al. (2010). In particular, for each grouping of interest, we estimate life cycle profiles for the 1st to 99th percentiles; we will refer to these as the "raw" predicted percentiles. At each age t, we then rank the values predicted by the raw percentiles. The qth-ranked value among this set is the qth "rearranged" percentile for that age. The rearranged profiles also give us the fraction of individuals with wealth less than or equal to zero; this is simply the fraction of raw percentile values with nonpositive values.

Over the bulk of the life cycle, the rearranged percentiles correspond very closely to the raw percentiles used to construct them—the two sets of profiles differ the most among people in their 20s. Using the average value of the raw percentiles to calculate means is more problematic, however, as they do not account for the very upper and lower tails of the distribution. We therefore use the raw means for ages 32 and above. For ages 22-31, we use a weighted average of the raw and percentile-based means. The weight on the percentile-based mean is $\frac{32-t}{10}$, which decreases linearly from 1 to 0 over ages 22 to 32.

B Family Composition Dynamics

Let $f_{i,t} \in \mathbb{F} = \{F_0, F_1, ..., F_{I^f-1}\}$ denote individual i's underlying, latent, family composition, and let $g_{i,t} \in \mathbb{G} = \{G_0, G_1, ..., G_{J^f-1}\}$, denote the outcome observed in the NLSY. There are two key assumptions. The first is that $f_{i,t}$ is condtionally Markov, with the $I^f \times I^f$ transition matrix, $\mathbf{A}^f_x$:

$$\mathbf{A}^f_{j,k\,|\,x} = \Pr(f_{i,t+1} = F_k \,|\, f_{i,t} = F_j,\, x_{i,t}) = \Pr(f_{i,t+1} = F_k \,|\, \mathcal{H}_t), \tag{A1}$$

where $x_{i,t}$ is a vector of explanatory variables and $\mathcal{H}_t$ denotes the time-t information set. (In a slight abuse of notation, we use f to denote family transition parameters.) The second is that the distribution of the observed outcome $g_{i,t}$ depends only on the contemporaneous realization of $f_{i,t}$. We can then place the probabilities that map $f_{i,t}$ to $g_{i,t}$ in the $I^f \times J^f$ matrix $\mathbf{B}^f_z$:

$$\mathbf{B}^f_{j,k\,|\,z} = \Pr(g_{i,t} = G_k \,|\, f_{i,t} = F_j,\, z_{i,t}) = \Pr(g_{i,t} = G_k \,|\, \mathcal{H}_{t-1}, f_{i,t}), \tag{A2}$$

where $z_{i,t}$ is a second vector of explanatory variables. The final element of our statistical model is the $1 \times I^f$ row vector $\mu^f_{1\,|\,x}$, which gives the unconditional distribution of the latent variable $f_{i,1}$ as a function of the vector $x_{i,1}$.

The remainder of this appendix describes how $\mathbf{A}^f_x$, $\mathbf{B}^f_z$ and $\mu^f_{1\,|\,x}$ are populated. We estimate separate sets of parameters for each race-sex-education combination; to reduce clutter, we will suppress this dependence in the exposition below.

B.1 Specialization

At this point, we will cease indexing individuals but will continue to index time. Let age_t denote age.

Let $m_t \in \{0, 1\}$ denote the individual's marital status, with 0 denoting single and 1 denoting married. Let $nc_t \in \{0, 1, 2, 3\}$ denote the number of co-residing children. (Empirically, we code families with more than 3 children as having 3.) It follows that $\mathbb{F} = \{0, 1\} \times \{0, 1, 2, 3\}$ and $I^f = 8$.

Let O_t indicate whether the individual is observed in the current survey wave. The set of observed outcomes, $\mathbb{G}$, consists of the latent states ($\mathbb{F}$) and missing ($O_t = 0$), so that $J^f = 9$.

B.2 State Transition Matrix

We model the non-zero elements of $\mathbf{A}_x^f$ in two steps. First, we find the probability that the individual will be married next period, using a univariate logit:

$$\Pr\left(m_{t+1}=1 \,|\, f_t=j,\, x_t\right) = \frac{\exp\left(\mathbf{x}_1(age_t, f_t)\boldsymbol{\zeta}_m\right)}{1+\exp\left(\mathbf{x}_1(age_t, f_t)\boldsymbol{\zeta}_m\right)}, \quad j \in \mathbb{F}, \tag{A3}$$
$$\mathbf{x}_1(age_t, f_t) := \left[\begin{array}{cccccccc} 1 & \frac{age_t}{10} & \frac{age_t^2}{1000} & m_t & m_t\frac{age_t}{10} & nc_t & \frac{nc_t^2}{10} \end{array}\right].$$

In an abuse of notation, we use $\mathbf{x}_1(age_t, f_t)$ to denote the explanatory variables in this regression.

The second step is to find the distribution of the number of children conditional on the individual's updated marital status. This is done using the Kumaraswamy distribution $K\big(p;\, \alpha, \beta\big)$, where $p \in [0,1]$, and $\alpha, \beta > 0$ are parameters to be estimated (Kumaraswamy, 1980). Like the Beta distribution, the Kumaraswamy distribution is a parsimonious but flexible function defined over the $[0,1]$ interval; however, it is much simpler to evaluate:

$$K(p;\, \alpha, \beta) = \Pr(y \leq p;\, \alpha, \beta) = 1-(1-p^{\alpha})^{\beta}. \tag{A4}$$

To apply the Kumaraswamy distribution, we break the unit interval into 4 equally-sized bins, each corresponding to a particular family size. We introduce state persistence by letting α and β depend on marital status and lagged family size:

$$\begin{aligned}
\Pr\left(nc_{t+1}=k \,|\, f_t=j, x_t, m_{t+1}\right) \;=\; & K\left(\frac{k+1}{4};\, \alpha(age_t, f_t, m_{t+1}), \beta(age_t, f_t, m_{t+1})\right) \\
& - K\left(\frac{k}{4};\, \alpha(age_t, f_t, m_{t+1}), \beta(age_t, f_t, m_{t+1})\right), \quad j \in \mathbb{F}, \quad k \in \{0,1,2,3\}, \qquad \text{(A5)} \\
\alpha(age_t, f_t, m_{t+1}) =\; & \exp\left(\mathbf{x}_2(age_t, f_t, m_{t+1})\boldsymbol{\zeta}_{\alpha}\right), \\
\beta(age_t, f_t, m_{t+1}) =\; & \exp\left(\mathbf{x}_2(age_t, f_t, m_{t+1})\boldsymbol{\zeta}_{\beta}\right), \\
\mathbf{x}_2(age_t, f_t, m_{t+1}) := & \left[\begin{array}{cccccccccc} 1 & \frac{age_t}{10} & \frac{age_t^2}{1000} & m_t & nc_t & \frac{nc_t^2}{10} & nc_t\frac{age_t}{100} & nc_t^2\frac{age_t}{1000} & \Delta m_{t+1} \end{array}\right].
\end{aligned}$$

$K\big(0;\, \alpha, \beta\big) = 0$ and $K\big(1;\, \alpha, \beta\big) = 1$ by definition.

Convoluting the two sets of probabilities yields $\mathbf{A}_x^f$.

B.3 Observation Probabilities

We rule out mis-reports, which gives the 8×9 observation matrix $\mathbf{B}_z^f$ a simple structure. It consists of an 8×1 column giving the probability that the individual is not observed, to which is appended an 8×8 diagonal matrix giving the probability that the individual is observed and accurately reports their underlying latent state. To populate $\mathbf{B}_z^f$, we assume that the probability that the individual is missing follows a logit specification:

$$\mathbf{B}_{j,0\,|\,z}^f = 1 - \frac{\exp\big(\mathbf{z}(age_t, f_t, O_{t-1})\gamma\big)}{1 + \exp\big(\mathbf{z}(age_t, f_t, O_{t-1})\gamma\big)}, \quad j \in \mathbb{F}, \tag{A6}$$

where

$$\mathbf{z}(age_t, f_t, O_{t-1}) := \left[\begin{array}{ccccccc} 1 & \frac{age_t}{10} & \frac{age_t^2}{1000} & m_t & \mathcal{I}_t^K & O_{t-1} \end{array}\right].$$

We assume that the observation probability depends on the presence of children but not their number, using $\mathcal{I}_t^K$ to indicate whether any children are present.

B.4 Initial Probabilities

We construct the initial distribution of latent states, $\mu_{1\,|\,x}^f$, in much the same way we found their transition probabilities. First, we find the probability that the individual is married. Then, conditional on marriage, we find the distribution across initial family sizes using the Kumaraswamy distribution $K\big(p;\, \alpha_{0m}, \beta_{0m}\big), m \in \{0, 1\}$; the calculations parallel those in equation (A5). The product of these two distributions is the initial distribution.

B.5 Likelihood

We estimate the parameters of the model using maximum likelihood, using the same approach as Gordon et al. (2023). To declutter the notation, we suppress the f superscripts and x (or z) subscripts on the matrices $\mathbf{A}$ and $\mathbf{B}$.

Suppose that an individual has the sequence of observed outcomes $\{g_t\}_{t=1}^T$. The likelihood of this sequence can be found via forward recursion (Bartolucci et al. 2010; Scott 2002; see also Hamilton 1994, chapter 22, and Farmer 2021):

1. Begin with the $1 \times I$ vector of initial latent state probabilities, μ_1. Our model starts at age 22. Observations from younger ages are dropped.[1]

[1] Some individual histories have differences such that the first time the individual is observed in our sample is after age 22. In those cases we use the model to update μ_t until that later age is reached.

2. Letting j_1 index the realization of g_1, calculate the $1 \times I$ vector $\eta_1 = \mu_1 \odot \left(\iota_{j_1} \mathbf{B}_1'\right)$, where $\odot$ denotes the Hadamard product, element by element multiplication. Here ι_j is the $1 \times J^f$ row vector with 1 at position j and zeros elsewhere; the product $\iota_{j_1} \mathbf{B}_1'$ returns (transposed) column j_1 of $\mathbf{B}_1$. Element i of η_1 thus gives the joint time-1 probability of latent state i and observed outcome g_{j_1}.

3. For $t = 1, 2, ..., T-1$, calculate $\eta_{t+1} = \frac{1}{\eta_t \mathbf{1}} \left(\eta_t \mathbf{A}_t\right) \odot \left(\iota_{j_{t+1}} \mathbf{B}_{t+1}'\right)$.[2] The $1 \times I$ vector η_{t+1} gives the joint probability distribution of the latent state f_{t+1} and the outcome observed at time $t+1$ ($g_{j_{t+1}}$), conditional on all outcomes observed through time t. The sum $\eta_t \mathbf{1}$ is the probability of the outcome observed at time t, conditional on all prior observed outcomes. The ratio $\frac{\eta_t}{\eta_t \mathbf{1}}$ thus gives the distribution of the latent states at time t, conditional on all outcomes observed through time t, and $\mathbf{A}_t$ updates the distribution to time $t+1$. The right-hand term of the Hadamard product accounts for g_{t+1}.

4. Calculate the cumulative probability, $\Pr(g_1, g_2, ..., g_T) = \prod_{t=1}^{T} (\eta_t \mathbf{1})$.

As Scott (2002) observes, if T is large, the product in item (4) may be very small, creating the risk of underflow errors. We thus calculate the log-likelihood as the sum of the logged probabilities, $\sum_t \ln(\eta_t \mathbf{1})$, rather than the log of the product, $\ln\left(\prod_t (\eta_t \mathbf{1})\right)$.

B.6 Condensed model

As described in Appendix D, the transition and observation probabilities for earnings depend on the individual's marital status (m_t) and whether they have any co-residing children ($\mathcal{I}_t^K$). Because family composition does not depend on earnings, we can estimate it transition probabilities independently of earnings and take them as given when estimating the process for earnings. It is useful, however, to condense the vector $f_{i,t} \in \mathbb{F}$ to $\tilde{f}_{i,t} \in \widetilde{\mathbb{F}} = \{0,1\} \times \{0,1\}$, by condensing the outcomes for children to $\{nc = 0, nc > 0\}$. This will halve the size of the composite state space used when estimating the earnings process.

The transition probabilities for the condensed process are found in two steps.

1. Find the sequence $\{\mu_t^f\}_{t=1}^T$, which gives the distribution of f_t at each date t. The sequence starts with μ_1^f and is updated using $\mu_{t+1}^f = \mu_t^f \mathbf{A}_t^f$.

[2] When the NLSY79 moves to two-year frequency, we multiply successive transition matrices.

2. Find the 4×4 transition matrix for $\tilde{f}_{i,t}$, $\mathbf{A}_t^{\tilde{f}}$, as

$$\mathbf{A}_{t,\tilde{j},\tilde{k}}^{\tilde{f}} = \frac{1}{\sum_{j \in \tilde{j}} \mu_{t,j}^{f}} \left[\sum_{j \in \tilde{j}} \mu_{t,j}^{f} \sum_{k \in \tilde{k}} \mathbf{A}_{t,j,k}^{f} \right], \quad \tilde{j}, \tilde{k} \in \widetilde{\mathbb{F}}, \quad t = 1, 2, ..., T, \tag{A7}$$

Suppose one constructs the condensed initial distribution $\mu_{1,j}^{\tilde{f}} = \sum_{j \in \tilde{j}} \mu_{1,j}^{f}$ and then applies $\mu_{t+1}^{\tilde{f}} = \mu_t^{\tilde{f}} \mathbf{A}_t^{\tilde{f}}$. The resulting sequence, $\{\mu_t^{\tilde{f}}\}_{t=1}^{T}$, will equal the sequence of distributions found by condensing $\{\mu_t^{f}\}_{t=1}^{T}$ year-by-year.

C Earnings and Employment Measures

In this appendix, we describe how we measure employment and earnings in the data. We begin by describing how we measure the employment status and earnings of singles in, successively, the NLSY79 and the CPS. Our variable construction is identical to that in Gordon et al. (2023), and the description is a verbatim copy of Appendix C in that paper. We conclude by describing how we extend our variable definitions to couples.

C.1 NLSY79

We measure earnings as the sum of wage income, salary income and the labor portion of farm and business income, with the latter found using the approach found in the Panel Study of Income Dynamics (PSID) (Survey Research Center, 1992). We will refer to individuals with no farm or business income as "workers." The NLSY79 reports both total hours and total weeks of work. We include military hours in the total: These are fairly small, as we do not use the NLSY79's military subsample.

In general, employed individuals have positive (80 or more) hours of work, positive (2 or more) weeks of work, and positive ($250 or more in 1980 dollars) earnings. When the three measures contradict, we define employment as follows.

1. If a person has positive earnings **and either** positive hours **or** positive weeks, she is employed.

2. A person with no hours **and** no weeks of work is nonemployed, regardless of earnings.

3. A **worker** with no earnings is nonemployed, regardless of hours or weeks.

4. A **nonworker** with positive hours **and** positive weeks is employed, regardless of earnings.

5. If a **nonworker** has no earnings **and either** no hours **or** no weeks, she is nonemployed.

The NLSY79 includes a variable that gives the respondent's residence status. Individuals who report "jail/prison" as their response are classified as being incarcerated.

C.2 CPS

Consistent with our approach in the NLSY, our CPS earnings measure includes not just wage and salary income, but also the labor component of farm and business income, again applying the PSID (Survey Research Center, 1992) methodology.

The CPS reports both the number of weeks worked in the prior year and the usual number of hours worked each week. We consider individuals who worked for less than two weeks to be nonemployed, along with those who worked less than 80 hours over the entire year (i.e., the product of usual hours worked each week and the number of weeks worked was less than 80). Workers (no business or farm income) with less than $250 (in 1980 dollars) of earnings are also considered to be nonemployed.

C.3 Couples

Married individuals are assigned fixed shares of total household earnings: if total household earnings are $e^*_{i,t}$, the wife and husband receive $\varsigma e^*_{i,t}$ and $(1-\varsigma)e^*_{i,t}$ respectively. A married individual is considered nonemployed only if the individual and their spouse are both not employed. In contrast, individuals residing in jail are always classified as incarcerated, even if they have an employed spouse. Individuals who are not nonemployed and not incarcerated are considered employed.

D The Stochastic Process for Earnings, Employment and Incarceration

Our model of earnings builds on the one developed in Gordon et al. (2023), and this appendix borrows heavily from the appendix of that paper. This appendix also mimics parts of Appendix B above and recycles some of its notation.

As discussed above and in the main text, the distribution of earnings outcomes depends on whether the individual is married or has co-residing children. This means that the log-likelihood for earnings will include family transition probabilities, especially when accounting for transtions between non-consecutive years. It is easiest to proceed, however, by first developing a model of earnings transitions with family composition taken as given. This is consistent with our assumption that causality runs from family composition to earnings, but not in the opposite direction.

Let $\ell_{i,t} \in \mathbb{L} = \{L_0, L_1, ..., L_{I^e-1}\}$ denote individual i's underlying, latent labor market state at date t, and let $e_{i,t} \in \Xi = \{E_0, E_1, ..., E_{J^e-1}\}$ denote the earnings outcome observed by the researcher. Our model relies on two key assumptions. The first is that $\ell_{i,t}$ is conditionally Markov, with the $I^e \times I^e$ transition matrix, $\mathbf{A}^e_x$:

$$\mathbf{A}^e_{j,k\,|\,x} = \Pr(\ell_{i,t+1} = L_k \,|\, \ell_{i,t} = L_j,\ \tilde{f}_{i,t+1},\ x_{i,t}), \tag{A8}$$

where $x_{i,t}$ is a vector of exogenous variables, and $\tilde{f}_{i,t+1} \in \widetilde{\mathbb{F}}$ denotes the condensed family composition variable defined in Appendix B.6. (In a modest abuse of notation, we will use e to denote earnings transition parameters.) Note that we allow earnings outcomes to depend on contemporaneous as well as lagged family outcomes. The second is that the distribution of the observed outcome $e_{i,t}$ depends on only the contemporaneous realization of $\ell_{i,t}$. We can place the probabilities that map $\ell_{i,t}$ to $e_{i,t}$ in the $I^e \times J^e$ matrix $\mathbf{B}^e_z$:

$$\mathbf{B}^e_{j,k\,|\,z} = \Pr(e_{i,t} = E_k \,|\, \ell_{i,t} = L_j,\ \tilde{f}_{i,t+1},\ z_{i,t}). \tag{A9}$$

Our statistical model is completed by the $1 \times I^e$ row vector $\mu^e_{1\,|\,x,f}$, which gives the distribution of the initial latent state $\ell_{i,1}$ as a function of $\tilde{f}_{i,1}$ and the vector $x_{i,1}$.

We turn now to describing how $\mathbf{A}^e_x$, $\mathbf{B}^e_z$ and μ^e_1 are populated. We estimate separate sets of transition and observation probabilities for each race-sex-education combination. To reduce clutter, we will suppress the dependence on race, sex, and education in the exposition below.

D.1 Specialization

At this point, we will cease indexing individuals but will continue to index time. Let age_t denote age.

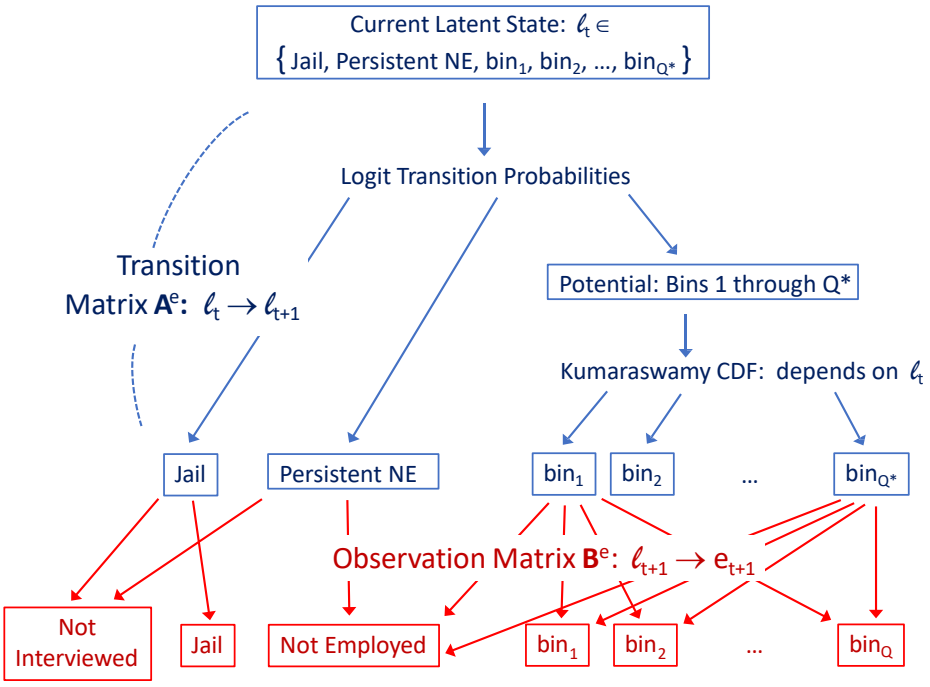
The set of latent states, $\mathbb{L}$, consists of incarceration, long-term nonemployment, and Q^* earnings *potential* bins. This set of latent values is then interacted with a $\{0, 1\}$

incarceration record flag, so that $\mathbb{L}$ contains $I^e = 2(Q^* + 2)$ elements.

Let O_t indicate whether the individual is observed in the current survey wave. The set of observed outcomes, Ξ, consists of incarceration, nonemployment (short- or long-term), Q positive earnings bins, and not interviewed/missing ($O_t = 0$). The nonmissing outcomes are also interacted with the incarceration record flag, so that Ξ contains $J^e = 2Q + 5$ elements.

Let p_q, $q = 0, 1, ..., Q$ denote the probability cutoffs for the earnings bins. We assume that $Q = 6$ and $p_q \in \{0.0, 0.2, 0.4, 0.6, 0.8, 0.94, 1\}$. We also assume that the bins for latent earnings potential are the same as those for observed earnings, so that $Q^* = Q = 6$. It is of course not necessary that Ξ and $\mathbb{L}$ align so closely, but it simplifies the analysis. We will impose this equality below.

Figure A1: Earnings process transitions



D.2 Latent state transitions

We define a person as having an incarceration record if he has been incarcerated in any *previous* period. Once a person is incarcerated, he will have an incarceration record in all subsequent periods. Figure A1 illustrates the remainder of the process for populating the matrices $\mathbf{A}^e$ and $\mathbf{B}^e$. As the top half of this figure shows, we find the elements of the

transition matrix $\mathbf{A}_x^e$ in two steps. First, we use a multinomial logit regression to determine the one-period-ahead probabilities of incarceration (IC), long-term nonemployment (NE), or potential employment ($\mathcal{Q}^*$):

$$\Pr\left(\ell_{t+1} \in k \,|\, \ell_t = j,\, \tilde{f}_{t+1},\, x_t\right) = \lambda_{j,k} \Big/ \sum_{m \in \{NE, \mathcal{Q}^*, IC\}} \lambda_{j,m}, \tag{A10}$$

$$j \in \mathbb{L}, \quad k \in \{NE, \mathcal{Q}^*, IC\},$$

$$\lambda_{j,NE} \equiv 1, \quad \forall j,$$

$$\lambda_{j,m} = \exp\left(\mathbf{x}_1(age_t, \ell_t, \tilde{f}_{t+1})\boldsymbol{\zeta}_m\right), \quad m \in \{\mathcal{Q}^*, IC\},$$

where $\{\boldsymbol{\zeta}_m\}_{m \in \{\mathcal{Q}^*, IC\}}$ are coefficient vectors for future states. Nonemployment is the benchmark state. In an abuse of notation, $\mathbf{x}_1(age_t, \ell_t, \tilde{f}_{t+1})$ denotes the explanatory variables in the logit regression. The elements of this vector include a polynomial in current age (age_t), indicators for the current state, and interactions:

$$\mathbf{x}_1(age_t, \ell_t, \tilde{f}_{t+1}) =$$
$$\left[\begin{array}{cccccccccccc} 1 & age_t & \frac{age_t^2}{100} & \mathcal{I}_t^{NE} & \mathcal{I}_t^{NE} age_t & \tilde{p}_j & \tilde{p}_j age_t & \tilde{p}_j^2 & \tilde{p}_j^2 age_t & \mathcal{I}_t^{IC} & \mathcal{I}_t^{IR} & m_{t+1} \quad \mathcal{I}_{t+1}^K \end{array}\right],$$

where: age_t denotes the individual's age at calendar year t; $\mathcal{I}_t^{NE}$ and $\mathcal{I}_t^{IC}$ are 0-1 indicators for long-term nonemployment or incarceration, respectively; $\mathcal{I}_t^{\mathcal{Q}^*} = 1 - \mathcal{I}_t^{NE} - \mathcal{I}_t^{IC}$ indicates positive earnings potential; $\mathcal{I}_t^{IR}$ is the 0-1 indicator for an incarceration record (previous incarceration).

The variable $\tilde{p}_j$ gives the individual's (approximate) earnings rank. In particular, when state j corresponds to earnings potential bin q_j^*, $\tilde{p}_j = [p_{q_j^*} + p_{q_j^*-1}]/2$. In the case at hand, $\tilde{p}_j \in \{0.1, 0.3, 0.5, 0.7, 0.87, 0.97\}$. When j indicates incarceration or persistent nonemployment, $\tilde{p}_j$ is set to 0. Because we treat $\tilde{p}_j$ as continuous rather than categorical, the number of variables in the logistic regression can be invariant to the number of bins. Moreover, using the midpoint value $\tilde{p}_j$ to characterize the current earnings potential bin means that as the number of bins grows large, our discretized distribution converges to a continuous one. It is natural to view our approach semiparametric approximations that will become more complicated as the sample size grows.

We also include two measures of time-$t+1$ household composition: $m_{t+1} \in \{0, 1\}$ gives the individual's marital status, with 1 denoting married; and the 0-1 indicator $\mathcal{I}_{t+1}^K$ equals 1 if any children under 18 reside in the individual's household. At this point in the exposition, we are simply taking these variables as given.

Second, we estimate the distribution of next period's earnings potential, conditional on being employed, across the bins. To do this, we assume that the conditional distribution of ranks follows the Kumaraswamy (1980) distribution,

$$K(p;\, \alpha, \beta) = 1 - (1 - p^{\alpha})^{\beta},$$

where $p \in [0, 1]$ and the parameters α and β are both strictly positive. (See the discussion of equation (A4).) In particular, if bin q^* covers quantiles p_{q^*-1} to p_{q^*},

$$\begin{aligned}
\Pr\big(\ell_{t+1} = \text{bin } q^* \,|\, \ell_t = j,\, \tilde{f}_{t+1},\, \tilde{f}_t,\, x_t\big) = {} & \Pr\big(\ell_{t+1} \in \mathcal{Q}^* \,|\, \ell_t = j,\, \tilde{f}_{t+1},\, \tilde{f}_t,\, x_t\big) \\
& \times \Big[K\Big(p_{q^*};\, \alpha\big(\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t)\big), \beta\big(\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t)\big)\Big) \\
& \quad - K\Big(p_{q^*-1};\, \alpha\big(\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1})\big), \beta\big(\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t)\big)\Big) \Big], \\
& j \in \mathbb{L}, \quad q^* \in \{1, 2, ..., Q\}, \\
\alpha\big(\mathbf{x_2}(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t)\big) = {} & \exp\big(\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t)\boldsymbol{\zeta}_\alpha\big), \\
\beta\big(\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t)\big) = {} & \exp\big(\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t)\boldsymbol{\zeta}_\beta\big).
\end{aligned} \tag{A11}$$

$K\big(0;\, \alpha, \beta\big) = 0$ and $K\big(1;\, \alpha, \beta\big) = 1$ by definition.

The Kumaraswamy parameters α and β are functions of the individual's age and current state, through

$$\begin{aligned}
&\mathbf{x}_2(age_t, \ell_t, \tilde{f}_{t+1}, \tilde{f}_t) = \\
&\quad \Big[\; 1 \quad age_t \quad \tfrac{age_t^2}{100} \quad \mathcal{I}_t^{NE} \quad \mathcal{I}_t^{NE} age_t \quad \tilde{p}_j \quad \tilde{p}_j age_t \quad \tilde{p}_j^2 \quad \tilde{p}_j^2 age_t \quad \mathcal{I}_t^{IC} \quad \mathcal{I}_t^{IR} \quad \Delta m_{t+1} \quad \mathcal{I}_{t+1}^{K} \;\Big].
\end{aligned}$$

We include the change in marital status, Δm_{t+1}, because entering or leaving a marriage changes the population used to calculate the individual's relative earnings ranking, potentially changing the ranking as well. $\boldsymbol{\zeta}_\alpha$ and $\boldsymbol{\zeta}_\beta$ are the associated coefficient vectors.

Recall that once a person is incarcerated, he will have an incarceration record for the rest of his life. As a result, $\mathbf{A}^e$ will be approximately block diagonal (see Table 2). The first $Q + 1$ rows of this matrix denote cases where at time t the individual has no incarceration record ($\mathcal{I}_t^{IR} = 0$) and is not currently incarcerated; his current state is long-term nonemployment (state 0) or one of the earnings potential bins. Such a person will not have an incarceration record at time $t+1$; even if he becomes incarcerated at $t+1$, he will not have a prior conviction at that point. The first $Q+1$ rows thus have (potentially)

nonzero values in the first $Q+2$ columns and zeros for the remainder. (Given that we start our indexing at 0, this corresponds to rows 0 through Q and columns 0 through $Q+1$.) The final $Q+2$ rows are for people who have an incarceration record at time t, $\mathcal{I}_t^{IR}=1$. These rows will have zeros in the first $Q+2$ columns and (potentially) nonzero values for the remainder. This leaves one row to consider, namely the one for a person who at time t has no incarceration record—he has not been incarcerated in the past—but is currently in jail. This person will have an incarceration record at time $t+1$, so row $Q+1$ is configured like the rows for those with an incarceration record at time t. The transition probabilities for this person will differ, however, from that of a person who at time t is both in jail and in possession of an incarceration record. To sum, the matrix $\mathbf{A}^e$ has a $(Q+1)\times(Q+2)$ block in the upper left corner, a $(Q+3)\times(Q+2)$ block in the lower right corner, and zeros everywhere else.

D.3 Observation probabilities

Recall that the set of observed outcomes, $\boldsymbol{\Xi}$, consists of the set of latent states, $\mathbb{L}$, and missing ($O_t=0$). To populate $\mathbf{B}^e$, we first find the probability that the individual is not interviewed at time t, using a logit specification:

$$\begin{aligned}\mathbf{B}^e_{j,0\,|\,z} &= \Pr\left(O_t=0\,|\,\ell_t=j,\,z_t,\,\tilde{f}_t,\,O_{t-1}\right)\\ &= 1-\frac{\exp\left(\mathbf{z}_1(age_t,\ell_t,\tilde{f}_t,O_{t-1})\boldsymbol{\gamma}_O\right)}{1+\exp\left(\mathbf{z}_1(age_t,\ell_t,\tilde{f}_t,O_{t-1})\boldsymbol{\gamma}_O\right)},\quad j\in\mathbb{L},\end{aligned}\tag{A12}$$

where

$$\mathbf{z}_1(age_t,\ell_t,\tilde{f}_t,O_{t-1})=\left[\begin{array}{cccccccc}1 & age_t & \frac{age_t^2}{100} & \mathcal{I}_t^{NE} & \mathcal{I}_t^{IR} & m_t & \mathcal{I}_t^{K} & O_{t-1}\end{array}\right].$$

The interview probability depends on the presence of earnings but not their rank.

Next, we find the distribution of states for individuals who are interviewed. We assume that the individual's incarceration record is measured accurately, so that the latent incarceration state maps 1-for-1 into observed incarceration. Likewise, we assume that the latent nonemployment state maps 1-for-1 into observed nonemployment, consistent with our view that the state represents long-term unemployment. We assume further that

incarceration records are reported accurately. With these assumptions, we get

$$\mathbf{B}^e_{j,j+1\,|\,z} = \frac{\exp\big(\mathbf{z}_1(a_t, \ell_t, \tilde{f}_t, O_{t-1})\boldsymbol{\gamma}_O\big)}{1+\exp\big(\mathbf{z}_1(age_t, \ell_t, \tilde{f}_t, O_{t-1})\boldsymbol{\gamma}_O\big)}, \quad j \in \{0, Q+1, Q+2, 2Q+3\}, \tag{A13}$$

$$\mathbf{B}^e_{j,k\,|\,z} = 0, \quad j \in \{0, Q+1, Q+2, 2Q+3\},\ k \notin \{0, j+1\}. \tag{A14}$$

To populate the remaining elements of $\mathbf{B}^e$, we assume that when the latent state ℓ_t is the earnings potential bin q^* the individual's observed outcome may be any observed earnings bin, $q \in \mathcal{Q}$, or nonemployment. Nonemployment realized in these circumstances is purely transitory, and has no effect on the individual's latent earnings prospects. Transitory incarceration shocks are ruled out, and the corresponding elements of $\mathbf{B}^e$ are set to zero. The non-diagonal elements of $\mathbf{B}^e$ can reflect either transitory earnings shocks or measurement error. For the exercise conducted here, we take the former interpretation.

Our approach for finding these probabilities is similar to the one employed for the latent states. First, we find the probability that the individual will be nonemployed or working, using a logit:

$$\begin{aligned}\Pr\big(e_t \in k \,|\, \ell_t = j,\, O_t = 1,\, z_t,\, \tilde{f}_t\big) = \lambda_{j,k} \Big/ \sum_{h\in\{NE,\mathcal{Q}\}} \lambda_{j,h}, \qquad & \text{(A15)}\\ j \in \{\text{bin } 1, \text{bin } 2, ..., \text{bin } Q\} \times \{0,1\},& \\ k \in \{NE, \mathcal{Q}\},& \\ \lambda_{j,NE} \equiv 1, \quad \forall j,& \\ \lambda_{j,\mathcal{Q}} = \exp\big(\mathbf{z}_2(age_t, \ell_t, \tilde{f}_t)\boldsymbol{\gamma}_{\mathcal{Q}}\big),&\end{aligned}$$

with the conditioning vector $\mathbf{z}_2$ given by

$$\mathbf{z}_2(age_t, \ell_t, \tilde{f}_t) = \Big[\ 1 \quad age_t \quad \tfrac{age_t^2}{100} \quad \tilde{p}_j \quad \tilde{p}_j age_t \quad \tilde{p}_t^2 \quad \tilde{p}_t^2 age_t \quad \mathcal{I}_t^{IR} \quad \mathcal{I}_t^{IR} age_t \quad m_t \quad \mathcal{I}_t^K\ \Big].$$

Note that $E_0 = NE$ is again the benchmark state. We continue to assume that incarceration records are reported accurately.

Next, the probabilities for the individual earnings bins are found using a univariate logit distribution:

$$\begin{aligned}\Pr\big(e_t = \text{bin } q \,|\, \ell_t = j,\, O_t = 1,\, z_t,\, \tilde{f}_t\big) = \Pr\big(e_t \in \mathcal{Q} \,|\, \ell_t = j,\, O_t = 1,\, z_t,\, \tilde{f}_t\big) \qquad & \text{(A16)}\\ \times \big[L^\star\big(p_q, \tilde{p}_j;\, \sigma_{\mathbf{z}_2}\big) - L^\star\big(p_{q-1}, \tilde{p}_j;\, \sigma_{\mathbf{z}_2}\big)\big],&\end{aligned}$$

$$
\begin{aligned}
j &\in \{\text{bin } 1, \text{bin } 2, ..., \text{bin } Q\} \times \{0, 1\}, \\
q &\in \{1, 2, ..., Q\}, \\
L^{\star}\big(p, \tilde{p};\, \sigma\big) &= \frac{\exp(\sigma(p - \tilde{p}))}{1 + \exp(\sigma(p - \tilde{p}))} \qquad \text{(A17)} \\
&\quad \Big/ \left[\frac{\exp(\sigma(1 - \tilde{p}))}{1 + \exp(\sigma(1 - \tilde{p}))} - \frac{\exp(-\sigma\tilde{p})}{1 + \exp(-\sigma\tilde{p})}\right], \\
\sigma_{\mathbf{z}_2} &= \big|\mathbf{z}_2(age_t, \ell_t, \tilde{f}_t)\boldsymbol{\gamma}_E\big|. \qquad \text{(A18)}
\end{aligned}
$$

Our decision to center the distribution $L^{\star}\big(p, \tilde{p}_j;\, \sigma\big)$ around the latent earnings rank $\tilde{p}_j$ is an identifying assumption. Given that the logistic density is symmetric around zero, our assumption implies that if the earnings bins are also symmetric, the most likely observed earnings level is the current latent state. The denominator in equation (A17) ensures that $\sum_q \Pr\big(\text{bin } q \,|\, \ell_j,\, O = 1,\, z,\, \tilde{f}\big) = \Pr\big(\mathcal{Q} \,|\, \ell_j,\, O = 1,\, z,\, \tilde{f}\big),\, \forall j$.

Multiplying the probabilities in equations (A15) and (A16) by the observation probability given in (A13) completes the process.

D.4 Initial probabilities

We construct the initial distribution of latent states, $\mu^e_{1\,|\,x}$, in much the same way we found their transition probabilities. First, we find the probability that the individual is incarcerated, nonemployed (long-term), or in one of the positive earnings potential bins, using a multinomial logit specification. Conditional on being positive, we find the distribution of earnings potential across its bins using the Kumaraswamy distribution $K\big(p;\, \alpha_0, \beta_0\big)$; the calculations parallel those in equation (A11). The parameters α_0 and β_0 are scalars to be estimated. The final step is to estimate the probability that the individual has an incarceration record, conditional on the other latent states. Here we use a logistic distribution, allowing the probability of an incarceration record to depend on $\mathcal{I}_0^{NE}$, $\mathcal{I}_0^{IC}$, and q_0. The product of these two probabilities gives us our initial distribution.

D.5 Combining earnings and family transitions

If we had a balanced panel, with no annual differences in individual histories, we could estimate the parameters for the earnings transition probabilities taking the observed family composition outcomes as given. From 1994 onwards, however, the NLSY79 is biennial rather than annual, and even before then many individuals have gaps in their interview history. In this case the family composition outcomes needed to calculate the earnings

transitions are unobserved, and the likelihood function must integrate over these unobserved outcomes using their transition probabilities.

To do this, we work with the combined state vector $\urcorner_{i,t} = (\tilde{f}_{i,t}, \ell_{i,t}) \in \widetilde{\mathbb{F}} \times \mathbb{L}$. The size of this vector is $I^{\urcorner} = 4I^e$. We construct the $I^{\urcorner} \times I^{\urcorner}$ transition matrix for $\urcorner_{i,t}$, $\mathbf{A}_t^{\urcorner}$, as follows. Let $\mathbf{A}_t^e(j_t, k_{t+1})$ denote the $I^e \times I^e$ state transition matrix for time-t earnings for an individual whose (condensed) family composition takes on value $\tilde{f}_{j_t}$ at time t and value $\tilde{f}_{k_{t+1}}$ at time $t+1$, and let $\mathbf{A}_{t,j_t,k_{t+1}}^{\tilde{f}}$ give the probability that a person with family composition $\tilde{f}_{j_t}$ transitions to $\tilde{f}_{k_{t+1}}$. We then have

$$\mathbf{A}_t^{\urcorner} = \begin{bmatrix} \mathbf{A}_{t,1,1}^{\tilde{f}} \mathbf{A}_t^e(1,1) & ... & \mathbf{A}_{t,1,4}^{\tilde{f}} \mathbf{A}_t^e(1,4) \\ \vdots & \vdots & \vdots \\ \mathbf{A}_{t,4,1}^{\tilde{f}} \mathbf{A}_t^e(4,1) & ... & \mathbf{A}_{t,4,4}^{\tilde{f}} \mathbf{A}_t^e(4,4) \end{bmatrix}.$$

We can likewise construct the vector of observed outcomes, $\lrcorner_{i,t} = (\tilde{f}_{i,t}, e_{i,t})$. One complication is that when earnings are unreported, so is family status. This means that all $(\tilde{f}_{i,t}, O_t = 0)$ combinations are collapsed into one outcome $(O_t = 0)$. The size of $\lrcorner_{i,t}$, $J^{\urcorner}$, is thus $4(J^e - 1) + 1 = I^{\urcorner} + 1$. To construct the $I^{\urcorner} \times J^{\urcorner}$ observation matrix $\mathbf{B}^{\urcorner}$, let $\mathbf{B}_t^e(j_t)$ denote the $I^e \times J^e$ state observation matrix for time-t earnings when (condensed) family composition takes on value $\tilde{f}_{j_t}$ at time t. Partition this matrix as $\mathbf{B}_t^e(j_t) = \begin{bmatrix} \widehat{\mathbf{B}}_t^e(j_t) & \widetilde{\mathbf{B}}_t^e(j_t) \end{bmatrix}$, where the $I^e \times 1$ vector $\widehat{\mathbf{B}}_t^e(j_t)$ gives the probability that the individual does not respond to the survey, and the $I^e \times I^e$ matrix $\widetilde{\mathbf{B}}_t^e(j_t)$ gives the outcome probabilities for when the individual does respond. Recalling that family outcomes are measured accurately, we get

$$\mathbf{B}_t^{\urcorner} = \begin{bmatrix} \widehat{\mathbf{B}}_t^e(1) & \widetilde{\mathbf{B}}_t^e(1) & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \widehat{\mathbf{B}}_t^e(2) & \mathbf{0} & \widetilde{\mathbf{B}}_t^e(2) & \mathbf{0} & \mathbf{0} \\ \widehat{\mathbf{B}}_t^e(3) & \mathbf{0} & \mathbf{0} & \widetilde{\mathbf{B}}_t^e(3) & \mathbf{0} \\ \widehat{\mathbf{B}}_t^e(4) & \mathbf{0} & \mathbf{0} & \mathbf{0} & \widetilde{\mathbf{B}}_t^e(4) \end{bmatrix},$$

where $\mathbf{0}$ denotes an $I^e \times I^e$ matrix of zeros. Because the probability of being interviewed depends on its lagged value, there are two values of $\mathbf{B}_t^{\urcorner}$, for $O_{t-1} = 0$ and $O_{t-1} = 1$ respectively.

The initial distribution of the composite latent state, the $1 \times 4I^e$ vector $\mu_{1|x}^{\urcorner}$, is

$$\mu_{1|x}^{\urcorner} = \mu_{1|x}^{\tilde{f}} \otimes \mu_{1|x}^e,$$

with $\otimes$ denoting the Kronecker product.

D.6 Likelihood

Our procedure for estimating the parameters of the model is the one described in Appendix B.5, with $\mathbf{A}_t^f$, $\mathbf{B}_t^f$ and $\mu_{1\,|\,x}^f$ replaced by $\mathbf{A}_t^{\urcorner}$, $\mathbf{B}_t^{\urcorner}$ and $\mu_{1\,|\,x}^{\urcorner}$, respectively. In doing this we take as given the values of $\mathbf{A}_t^{\tilde{f}}$, $\mathbf{B}_t^{\tilde{f}}$ and $\mu_{1\,|\,x}^{\tilde{f}}$ estimated in the first stage (Appendix B) and estimate only the parameters related to earnings.

When unconstrained, our estimates imply that nonemployment significantly increases the probability of non-response for White women with a high school degree and White men with at least a high school degree. This leads to a significant increase in the model-predicted nonemployment rate. We bound this effect, by restricting the coeffection on $\mathcal{I}_t^{NE}$ in equation (A12), so that the non-reporting bias in the nonemployment rate is at most one-quarter of the observed nonemployment rate. By way of example, this means that if the nonemployment rate among respondents is 20%, the model-predicted nonemployment rate for all individuals is at most 25%.

D.7 Incarceration imputations

(This appendix has been copied verbatim from (Gordon et al., 2023, Appendix F).) For a number of groups—White men with a college degree, White women with at least a high school diploma, and Black women with either a high school or a college degree—the incidence of incarceration is so low that their incarceration-related parameters cannot be estimated accurately. In these cases, we use the data to estimate a simplified model that omits incarceration. To this set of parameters, we add incarceration-related parameters estimated for other, similar groups, namely White men with some college education, or White women without a high school diploma, or Black women with some college experience. These coefficients are identified by an entry of "NA" in the standard error slot.

When making these imputations, we adjust the intercepts for the incarceration probabilities to match the probabilities observed in the NLSY for the group. The logic of our adjustments is the following. Consider a simple static logistic model, where

$$
\begin{aligned}
u &= \frac{1}{1+E+IC}, \\
e &= \frac{E}{1+E+IC}, \\
ic &= \frac{IC}{1+E+IC},
\end{aligned}
$$

give the probabilities of being nonemployed, employed, or incarcerated, respectively. Sup-

pose the incarceration constant for group i is known to be $\ln(IC_i)$, and we want to build off this constant to impute $\ln(IC_j)$ for group j. If we know the probabilities ic_i, ic_j, u_i, u_j, we have

$$ic_i = \exp(\ln(IC_i)) \cdot u_i,$$
$$ic_j = \exp(\ln(IC_j)) \cdot u_j,$$

which can be rearranged to yield

$$\exp\left(\frac{\ln(IC_j)}{\ln(IC_i)}\right) = \frac{ic_j}{ic_i} \cdot \frac{u_i}{u_j},$$
$$\Rightarrow \ln(IC_j) = \ln(IC_i) + \big(\ln(ic_j) - \ln(ic_i)\big) - \big(\ln(u_j) - \ln(u_i)\big). \tag{A19}$$

Letting $\ln(IC)$ be the intercept for the incarceration-related expression, the latter two terms of (A19) comprise our adjustment. In practice, we calculate ic and u by averaging across all sample periods, and we estimate ic as the average probability of an incarceration record, which is somewhat less noisy than incarceration itself.

A further complication is that for White men and Black women with a college degree, and White women with a high school degree, the fraction of individuals with an incarceration record is very low, 0.03% or less. In these cases, we have a second layer of imputation, for the fraction ic itself. Black female college graduates are assigned the incarceration record fraction observed for Black female high school graduates; White female high school graduates are assigned the fraction for White female college graduates; and White male college graduates are given the fraction for White men with some college, scaled downward using the fractions observed for Black men. All of these imputations are admittedly *ad hoc*, but the groups to which they are applied have very low rates of incarceration, implying that any imputation errors will have small quantitative effects.

E Earnings Cutoffs and Bin Means

Although the NLSY79 is our principal data source, to calculate the cutoffs that delineate earnings quantiles, we use the larger sample available in the CPS (downloaded from IPUMS; Ruggles et al. 2020). CPS data are available from 1962 to 2023; however, we limit our sample to 1976 onward because data on hours worked, which we need for our measure of employment, are not available prior to then.

Since the CPS is not a panel, we employ a synthetic cohort approach. We include individuals born between 1941 and 1980 (i.e., the NLSY79 +/- two cohorts) and use cohort dummies to account for any cohort effects within this group. This consists of around 4.2 million observations. We restrict the sample to White and Black individuals, who together make up about 94% of the sample. We also exclude those for whom educational attainment is not reported and limit observations to those aged between 22 and 66. We choose 66 as the upper limit since that is the normal Social Security retirement age for the NLSY79 cohort. After applying these restrictions, around 3 million individuals remain. We define earnings broadly to include not just wage and salary income, but also the labor portions of farm and business income. Since we have separate categories for the nonemployed and incarcerated, we limit ourselves to those who were employed in the previous year; Appendix C describes how we measure employment in more detail.[3] Those who remain (about 2.4 million) form our sample of employed individuals. We weight the data to ensure that the sample is representative.

We estimate earnings bin cutoffs (quantiles) separately for each race-sex-marital status-education group. In particular, in each group, for each quantile q, we run the following quantile regression:

$$y_{i,t} = \beta_{q,0} + \beta_{q,1} age_{i,t} + \beta_{q,2} age_{i,t}^2 + \beta_{q,3} age_{i,t}^3 + \sum_{m=2}^{5} \gamma_{q,m} cohort_{m,i} + \epsilon_{q,i,t}. \tag{A20}$$

Here, $y_{i,t}$ denotes earnings, $age_{i,t}$ is age, and $cohort_{m,i}$ is a dummy variable for one of the five eight-year cohorts in our sample.

Applying the cutoffs to the CPS data, we calculate within-bin mean earnings at each age for each group. We then fit a cubic polynomial in age with cohort dummies through these age-specific means. We use this to construct life-cycle profiles of within-bin mean earnings for our cohort of interest.

F Portfolio Share Regressions

Before we introduce our regression specification, we provide some descriptive results in Table A1, which shows average portfolio shares by race and net worth bin, expressed as fractions of net worth.[4]. Note that debt has a negative sign and its share will be negative

[3] The CPS elicits income information for the year prior to the survey year.

[4] Tables A1 and A2, along with their discussion, parallel those found in Jones and Neelakantan (2023)

for individuals with positive net worth. Our decision to use net worth in the denominator is why the table excludes shares for households below the 25th net worth percentile. Small values of net worth often reflect large offsetting asset and debt holdings that generate large absolute shares. Moreover, when net worth is close to zero, small changes in either assets or debt can generate large proportional changes in net worth and thus portfolio shares. In such cases, the shares convey relatively little information.

Table A1 shows that wealthier households are more heavily invested in high-yielding assets (such as stocks and private equity), less invested in low-yielding assets (such as cash and vehicles), and less indebted. It also reveals that, conditional on wealth, Black and White households have reasonably similar portfolios. The largest differences appear between the 75th and 99th percentiles of the aggregate distribution, where White households tend to be more heavily invested in equity; this is also the bean where mean wealth holdings differ most substantially.

Table A1: Portfolio Share Percentages by Race and Net Worth

NW percentile bins	25th-50th		50th-75th		75th-99th		Top 1%	
Race	W	B	W	B	W	B	W	B
Cash	9	7	8	6	5	5	2	1
Vehicle Equity	38	35	15	14	7	10	2	2
Real estate	145	151	83	89	50	73	22	25
Business	3	2	5	10	13	12	41	44
Stocks	15	17	21	15	35	19	31	28
Other assets	5	2	9	7	11	8	16	7
Mortgages	-77	-78	-24	-24	-10	-15	-2	-1
Credit card	-4	-3	-1	-1	0	0	0	0
Education loans	-7	-9	-1	-4	0	-1	0	0
Other debt	-2	-3	0	-1	0	0	-1	-2
Mean NW ($2010)	68.6K	63.4K	223K	219K	1.21M	0.73M	16.6M	14.4M

Source: Authors' calculations using the 2019 and 2022 waves of the SCF.
Note: 'Percentile bins are for both races combined but calculated separately for each wave. "W" and "B" denote White and Black, respectively. "Other assets" consist of bonds and other safe assets. "NW" denotes net worth. "K" denotes thousands and "M" denotes millions.

Let $a_{i,t}$ denote household i's total net worth, $a_{j,i,t}$ its holdings of asset/debt item j, and $s_{j,i,t} = a_{j,i,t}/a_{i,t}$ item j's share of total net worth. In these calculations debt has a negative sign, so that its share will be negative for individuals with positive net worth and positive for individuals with negative net worth. Because $s_{j,i,t}$ is undefined when net worth is exactly zero, these observations will be excluded from our analysis. We also drop

the bottom and top one percent of the share distribution for each item: small values of net worth can reflect large offsetting asset and debt holdings that result in large absolute shares.

From the remainder of this appendix we drop the i and t subscripts. Consistent with our "separate books" assumption, net worth is measured on a per capita basis. Net worth is further transformed by applying the Inverse Hyperbolic Sine (IHS) function:

$$\widehat{a} \coloneqq \text{IHS}(a) = \ln\Big(a + \sqrt{a^2+1}\Big). \tag{A21}$$

The IHS transformation provides a generalization of the natural log. It can be shown that for large positive values of a, $\widehat{a} \approx \ln(2a)$; for large negative values, $\widehat{a} \approx -\ln(-2a)$; and in the interval $[-1,1]$, $\widehat{a} \approx a$. We will work with a third-order spline in $\widehat{a}$. Define

$$\begin{aligned}
\widehat{a}_1 &= \min\{\widehat{a}, 10\},\\
\widehat{a}_2 &= \max\{\widehat{a}_1 - 5, 0\},\\
\widehat{a}_3 &= \max\{\widehat{a} - 10, 0\},\\
\widehat{a}_{11} &= \widehat{a}\mathcal{I}^{\widehat{a}<0},\\
\widehat{a}_{12} &= (\widehat{a}+5)\mathcal{I}^{\widehat{a}<5},
\end{aligned}$$

where $\mathcal{I}^{\mathcal{A}}$ denotes the 0-1 indicator for event $\mathcal{A}$. It bears remembering that the the IHS function is not homothetic: in all our applications, we measure wealth in thousands of 2010 dollars.

Using this notation, the net worth share of asset j is

$$\begin{aligned}
\hat{s}_j = \big(\mathcal{I}^{\widehat{a}>0} - \mathcal{I}^{\widehat{a}<0}\big)\exp\Big(&\gamma_{j,0} + \gamma_{j,1}age + \gamma_{j,2}age^2 + \gamma_{j,3}\widehat{a}_1 + \gamma_{j,4}\widehat{a}_1^2 + \gamma_{j,5}\widehat{a}_2^3\\
&+ \gamma_{j,6}\widehat{a}_3^3 + \gamma_{j,7}\mathcal{I}^{\widehat{a}<0} + \gamma_{j,8}\widehat{a}_{11} + \gamma_{j,9}\widehat{a}_{11}^2 + \gamma_{j,9}\widehat{a}_{12}^3 + \gamma_{j,10}\widehat{a}_1 age\\
&+ \gamma_{j,11}\mathcal{I}^{ed=1} + \gamma_{j,12}\mathcal{I}^{ed=2} + \gamma_{j,14}\mathcal{I}^{ed=3} + \gamma_{j,14}\mathcal{I}^{ms=1} + \gamma_{j,15}\mathcal{I}^{ms=2}\\
&+ \gamma_{j,16}cohort + \gamma_{j,17}\mathcal{I}^{race=White}\Big),
\end{aligned} \tag{A22}$$

where: $ed = 1,2,3$ denotes, respectively, individuals without a high school degree, with a high school degree and no college attendence, and college attendence without bachelor's degree; $ms = 1,2$ denotes, respectively, single men and single women; and *cohort* denotes the rank order of the individual's 8-year birth cohort, proceeding from -3 (for individuals born in the years 1933 through 1940) to 3 (1980-1987).

The share equations for debts take the same form as those for assets, except that the first two terms of equation (A22) are reversed. This flips the signs of the shares to be consistent with our negative signing of debt.

Table A2: Race Coefficients in Portfolio Share Regressions

	Race (White) Indicator		Regression R^2	
Portfolio Item	Coefficient	Standard Error	No Race Indicators	Race Indicator
Vehicles	0.263	0.017	0.5848	0.5913
Real Estate	0.227	0.020	0.9999	0.9999
Cash	0.041	0.015	1.0000	1.0000
Stocks	0.159	0.018	0.9999	0.9999
Business	0.047	0.042	0.1832	0.1833
Bonds	0.359	0.087	0.0439	0.0446
Other Safe Assets	0.250	0.038	0.0362	0.0368
Education Loans	-0.072	0.019	0.4932	0.4936
Credit Cards	0.417	0.028	0.3632	0.3715
Mortgages	0.319	0.028	0.2427	0.2468
Vehicle Debt	0.137	0.027	0.3135	0.3145
Other Debts	0.418	0.034	0.9999	0.9999

Source: Authors' calculations using the 1989-2022 waves of the SCF.
Note: The first two columns show the coefficient and standard error, respectively, on race (White indicator) in a multivariate regression for portfolio shares. The third and forth columns show the R^2 statistics for regressions with and without the race indicator.

Table A2 shows the coefficient on the race (White) indicator for each portfolio item. In almost every case, White households hold more assets and less debt than Black households with the same covariates—the one exception is education debt. The table also shows the R^2 statistic for regressions with and without the race indicator. Comparing R^2 values suggests that race explains (statistically) only a small amount of the variation in portfolio shares.

Because our estimates are nonlinear approximations, the predicted shares need not add up to 1. Let s_{asst} denote the sum of the asset shares and s_{debt} denote the sum of the debt shares. While $s_{asst} + s_{debt}$ always equals 1, our signing conventions imply when net worth is positive, $s_{asst} \geq 1$ and $s_{debt} \leq 0$; when net worth is negative, we have $s_{asst} \leq 0$ and $s_{debt} \geq 1$. We therefore adjust the predicted shares as follows. When net worth is positive we check that the sum of the predicted asset shares ($\hat{s}_{asst}$) is at least 1, rescale it to 1 if it is not, and then set $\hat{s}_{debt} = 1 - \hat{s}_{asst}$. When net worth is negative we check that

the sum of the predicted debt shares ($\hat{s}_{debt}$) is at least 1, rescale it to 1 if it is not, and then set $\hat{s}_{asst} = 1 - \hat{s}_{debt}$.

G Return Distributions

With one exception, we assume that real asset returns are normally distributed. We calculate means and standard deviations for each asset class from observed returns as follows.

The nominal return of stocks inclusive of dividends is obtained from the Wilshire 5000 total market cap index (FRED series WILL5000INDFC) on a year-over-year basis. For autos, we assume that the implicit rent from vehicle ownership net of depreciation offsets the vehicle's expected capital loss, so that the nominal return is zero. The secondary market rate on 3-month Treasury bills (FRED series TB3MS) provides our estimate of the nominal bond return. The nominal return on cash is zero.

The nominal return on mortgage debt is taken to be the average 30-year fixed rate mortgage (FRED series MORTGAGE30US). The nominal return on credit card debt is the commercial bank interest on credit card plans (FRED series TERMCBCCALLNS). The nominal return on auto debt is the rate on consumer installment loans for 60-month auto loans (FRED series RIFLPBCIANM60NM).

With the exception of housing, we assume that Black and White individuals face the same asset-specific returns. For housing, we apply Kermani and Wong's (2021) estimate that the returns earned by Black homeowners on average are 2.32pp lower than those earned by Whites. The benchmark return on housing earned by White individuals consists of capital gains and implicit rents, less depreciation. For capital gains, we use the year-over-year growth of the all-transactions house price index constructed by the FHFA (FRED series USSTHPI). For the implicit rent net of depreciation, we follow Fagereng et al. (2020) who assume this to be a constant 2.88%.

To compute real returns for the above series, we subtract from each nominal return the year-over-year growth rate of the GDP deflator (FRED series GDPDEF). All our FRED data are of February 2, 2022 vintage.

The real student loan interest rate is annual and comes from Table 7 of Gordon and Hedlund's (2019) web appendix. The methodology used to construct the series is described in Section B.3 of the same document.

Table A3 reports the mean and standard deviation of these returns within each asset class.The credit card and auto loan series are quarterly and begin in 1994Q4 and 2006Q3,

respectively, and stop in 2021Q4. The student loan series is annual and runs from 1987 to 2010. The other real return series are available quarterly from 1976Q1 through 2021Q4.

Table A3: Real Asset Returns by Asset Class

	Mean	Standard deviation	Source
Housing	0.047	0.040	FRED-USSTHPI, Fagereng et al. (2020)
Autos	-0.031	0.021	nominal return = 0
Stocks	0.103	0.158	FRED-WILL5000INDFC
Bonds	0.012	0.024	FRED-TB3MS
Cash	-0.031	0.021	nominal return = 0
Business equity	0.102	0.245	Fagereng et al. (2020)
Mortgages	0.047	0.025	FRED-MORTGAGE30US
Credit cards	0.118	0.016	FRED-TERMCBCCALLNS
Auto loans	0.035	0.014	FRED-RIFLPBCIANM60NM
Student loans	0.045	0.016	Gordon & Hedlund (2019)

Note: Listed sources are for the nominal component of the returns. Real returns found by subtracting price growth as measured by the GDP deflator (FRED series GDPDEF). All returns are assumed to be normally-distributed, except for business equity, which follows a Pareto distribution. Housing returns are for White individuals.

The one exception to our assumption that returns are normally-distributed is business equity. To calibrate this return, we appeal to (Fagereng et al., 2020, Table 3), who use Norwegian tax data to estimate the mean and standard deviation of private equity returns, along with their 10th, 50th and 90th percentiles. Assuming the *gross* return follows a Pareto distribution, that is, $F(R) = 1 - (R_m/R)^\alpha$, we employ a GMM scheme to find the values of R_m and α that best match these statistics. The criterion function we use gives each percentile a weight of 5 and the mean and standard deviation a weight of 1. We compute the Pareto moments from 100,000 Monte Carlo draws. Our procedure returns $\alpha = 5.6145$ and and $R_m = 0.9058$. The resulting gross returns have a mean, standard deviation, 10th percentile, 50th percentile and 90th percentile of 1.102, 0.245, 0.923, 1.025, and 1.365, respectively.

Recall that the total realized return on wealth for person i at time t is $\sum_j s_{j,i,t} r_{j,i,t} + r^{d_i}(a_{i,t})$. We have just described how we calibrate the asset-specific returns $\{r_{j,i,t}\}_j$. Appendix F above describes how we impute the shares $\{s_{j,i,t}\}_j$. This leaves the adjustment $r^{d_i}(a_{i,t})$.

Fagereng et al. (2020) find that rates of return vary by sex and education even after controlling for portfolio composition. We impose their estimated return differentials as

follows. (Fagereng et al., 2020, Table 4, specification 2) report that each year of education increases the annual rate of return by 0.1492pp and that the return earned by men exceeds that of women by 0.0495pp. Treating our four education categories (LHS, HS, SC, and CG) as corresponding to 10, 12, 14, and 16 years of education, respectively, we adjust returns so that each education level earns 0.2984pp than the one immediately below. We also impose that men earn 0.0495pp more than women. We normalize these adjustments (by substracting a constant) so that their mean effect, when averaged over our initial distribution of education and sex, is zero. The return differential varies from roughly -0.5pp for women without a high school degree to 0.4pp for men with Bachelor's degrees.

Fagereng et al. (2020) also find scale effects within broad asset classes. We capture this affect by adding the scale term $\alpha_{\text{sc}} \times \max\{0, \widehat{a} - \text{IHS}(295)\}$. We assume that this "pure" scale effect, which is distinct from portfolio effects, operates only at the upper tail of the wealth distribution; the cutoff 295,000 equals the mean wealth of White individuals in our SCF sample. The parameter α_{sc} is calibrated along with the model's preference parameters.

H Distribution of Intergenerational Transfers

The SCF asks respondents whether they (or their spouses, if married) have received any significant inheritances or gifts, and if so, inquires further about dates, givers and amounts. The questions are retrospective, and the respondent is asked about each individual transfer. These data allow us to estimate the frequency and size of transfers received over the life cycle. In keeping with our general practice, for married individuals we include transfers received by either spouse and split the total in half.

We begin by estimating the probability an individual receives a positive transfer in a given year, using a logistic specification. For White individuals we estimate full sets of parameters for each education group, with additional controls for age and marital status. For Black individuals, who comprise a smaller portion of observations, we apply a common age term, and allow education and marital status to enter as shift terms. Because the odds of receiving a transfer in a given calendar year are fairly small, we collect data from every SCF wave for people born between 1941 and 1980 (i.e., the NLSY79 +/- two cohorts) and include cohort dummies to account for any cohort effects.

Because individuals can receive multiple bequests and gifts over their lives, we use a cumulative measure of these transfers when estimating them. To this end, we rescale the transfers in our data to the value they would have if the individual had received

them at age 60, rather than the actual date of receipt, using an annual discount rate of 3 percent. Summing the rescaled transfers yields lifetime totals, measured in age-60 dollars, at the time of the survey interview. The distribution of lifetime transfers as of age 70 provides our estimation targets—this is roughly the age in the data at which the lifetime totals peak. Panel A of Table A4 shows the mean and quantiles that we target. For every education group, Black individuals receive fewer and smaller transfers than their White counterparts (Jones and Neelakantan, 2022a). (Because of the small number of positive transfers, we treat Black individuals without a bachelor's degree as a single group.)

Table A4: Lifetime Receipts of Intergenerational Transfers

				Panel A. Data				
Quantiles								
0.25	0	0	0	0	0	0	0	0
0.50	0	0	0	6	0	0	0	0
0.60	0	0	0	36	0	0	0	0
0.70	0	0	27	81	0	0	0	0
0.75	0	6	44	113	0	0	0	0
0.80	15	29	79	167	0	0	0	0
0.85	49	64	107	233	0	0	0	29
0.90	111	115	201	390	0	0	0	129
0.95	181	262	424	758	22	22	22	374
0.975	235	427	557	1,423	89	89	89	546
Mean	30	52	130	202	6	6	6	70
				Panel B. Model				
Quantiles								
0.25	0	0	0	0	0	0	0	0
0.50	0	0	0	0	0	0	0	0
0.60	0	0	0	47	0	0	0	0
0.70	0	0	36	101	0	0	0	0
0.75	0	0	58	140	0	0	0	0
0.80	24	37	82	193	0	0	0	0
0.85	56	76	124	245	0	0	0	55
0.90	107	133	188	511	0	0	0	147
0.95	152	237	943	1,530	26	26	26	308
0.975	305	662	1,668	2,446	98	98	98	779
Mean	31	55	135	226	6	6	6	75

Notes: Data quantities are lifetime totals as of age 70, discounted to age 60 at an annual rate of 3%, and reported in thousands of 2010 dollars. Model quantities are simulations of the model described in this appendix, using parameters estimated from the SCF. Simulations run through age 85.

To reduce the number of state variables in the model, we assume that intergenerational transfers are independent across age. The probability of receiving a transfer at any age is proportional to the probabilities observed in the SCF data, estimated as described above.

In contrast, we assume that, when positive, transfers follow the same Pareto distribution at every age:

$$F(b \,|\, b > 0) = 1 - \left(\frac{b_0}{b}\right)^{\alpha}, \; b_0 > 0, \alpha > 1, \tag{A23}$$

with the parameters b_0 and α varying by race and education level. We estimate b_0 and α, along with the probability scaling factor p_0, via simulation, targeting the quantiles of total lifetime transfers, along with their mean. We discretize the estimated Pareto distribution into a two-point distribution as follows. Let b_H denote the high realization, which we assume occurs with probability $q_c = 0.2$, let b_L denote the low realization. Letting $\mathbb{E}(b)$ denote the mean bequest (when positive), we have

$$\begin{aligned}
\mathbb{E}(b) &= b_0 \frac{\alpha}{\alpha - 1}, \\
b_H &= q_c^{\frac{1}{\alpha}} \, \mathbb{E}(b), \\
b_L &= \left(1 - q_c^{1-\frac{1}{\alpha}}\right) \frac{\mathbb{E}(b)}{1 - q_c}.
\end{aligned}$$

We estimate (p_0, b_0, α) separately for each race-education group. Panel B of Table A4 shows the distribution of transfers implied by our estimated (and discretized) model. The model matches the mean and quantiles up to the 90th percentile quite well. It overstates some of the higher quantiles, due to the coarse (two-point) discretization.

I Distribution of Initial Wealth

We will assume that $\widehat{a}_{i,0} = \text{IHS}(a_{i,0})$, the Inverse Hypberbolic Sine of initial wealth, follows a normal distribution: $\widehat{a}_{i,0} \sim \mathcal{N}(\mu_{j(i)}, \sigma^2_{j(i)})$, with the index $j(i)$ identifying the cell to which the individual belongs at time 0.

To calibrate μ and σ, note that wealth itself can be expressed as

$$a_{i,0} = \text{IHS}^{-1}(\widehat{a}_{i,0}) = \frac{1}{2}\big[\exp(\widehat{a}_{i,0}) - \exp(-\widehat{a}_{i,0})\big]. \tag{A24}$$

(This can be confirmed by inserting equation (A21) into equation (A24).) Suppose individual i belongs to cell j_0. It then follows from the properties of lognormal distributions

that the expected value of $a_{i,0}$, $a^M_{j_0}$, equals:

$$\begin{aligned} a^M_{j_0} &= \frac{1}{2}\Big[\exp\big(\mu_{j_0} + \tfrac{1}{2}\sigma^2_{j_0}\big) - \exp\big(-\mu_{j_0} + \tfrac{1}{2}\sigma^2_{j_0}\big)\Big] \\ &= \text{IHS}^{-1}(\mu_{j_0})\exp\big(\tfrac{1}{2}\sigma^2_{j_0}\big). \end{aligned} \tag{A25}$$

Because IHS$(\cdot)$ is monotonic, the qth quantile of $a_{i,0}$, $a^q_{j_0}$, equals

$$a^q_{j_0} = \text{IHS}^{-1}(\widehat{a}^q_{j_0}). \tag{A26}$$

$$\widehat{a}^q_{j_0} = \mu_{j(i)} + \sigma_{j_0}\Phi^{-1}(q), \tag{A27}$$

where $\Phi^{-1}(q)$ returns quantile q of the standard normal distribution.

Given estimates of $a^M_{j_0}$ and $a^{50}_{j_0}$, equations (A26) and (A27) allow us to calibrate μ_{j_0} and $\sigma^2_{j_0}$ via the method of moments. For each cell j_0 we target the mean and the 10th, 25th, 50th, 75th and 90th percentiles of age-22 wealth.

J Taxes

J.1 Model

We model taxes $(\tau_{i,t})$ using a well-known specification (dating back to Feldstein 1969 and recently utilized by Heathcote et al. 2017) that has been modified to account for the preferential tax treatment of asset income:

$$\tau_{i,t} = \max\{y^a_{i,t} + y^e_{i,t}, 0\} - \theta^0_{i,t}\big(\max\{y^a_{i,t} + y^e_{i,t}, 0\}\big)^{1-\theta^1_{i,t}}, \tag{A28}$$

$$y^e_{i,t} = e_{i,t} + ss_{i,t}, \tag{A29}$$

$$y^a_{i,t} = \text{IHS}^{-1}\big(\xi^0_{i,t} + \xi^1_{i,t}\widehat{a}_{i,t} + \xi^2_{i,t}\max\{\widehat{a}_{i,t} - 3, 0\}^2\big), \quad \widehat{a}_{i,t} := \text{IHS}\big((1 + r_{i,t})a_{i,t}\big), \tag{A30}$$

where: $y^e_{i,t}$ and $y^a_{i,t}$ are the taxable components of asset and non-asset income, respectively; $\theta^0_{i,t}$, $\theta^1_{i,t}$, $\xi^0_{i,t}$, $\xi^1_{i,t}$ and $\xi^2_{i,t}$ are parameters; IHS$(\cdot)$ and IHS$^{-1}(\cdot)$ are the Inverse Hyperbolic Sine function and its inverse, as defined in equations (A21) and (A24); and the remaining variables retain the definitions assigned in Section 3 of the main text. Equation (A28) returns the total tax liability. Equations (A29) and (A30) determine the portions of non-asset (mostly earnings) and asset income, respectively, subject to taxation.

While equation (A29) is a simple identity not allowing for any exemptions, equation (A30) can be parameterized so that a portion of asset income is exempt ($y^a < ra$). The

function for y^a should reflect, among other things, portfolio composition (e.g., the fraction of tax-exempt municipal bonds) and the fraction of capital gains that are realized.

J.2 Estimation

Using TAXSIM (Feenberg and Coutts, 1993) and code substantively the same as that used by Gale et al. (2022), we compute federal income and payroll tax liabilities for SCF households.[5] We then increase this amount proportionately for each household by the ratio of aggregate state and local taxes to aggregate federal taxes (FRED series A074RC1Q027SBEA and W071RC1Q027SBEA), as the state to which the respondent pays taxes is observed only in the confidential geocoded SCF. This adjustment varies by year, but on average it increases tax liabilities by 28%. Expressing estimated tax liabilities on a real, per-adult equivalent basis gives us a direct measure of taxes, $\tilde{\tau}_{i,t}$.

As part of Gale et al.'s (2022) procedure, estimates of dividend income, interest income, short-term and long-term capital gains, non-tax-exempt pension income, and income from other properties are calculated. We sum these items into "taxable asset income", $\tilde{y}^a_{i,t}$, which we also express in real per-adult terms. Figure A2 presents a bin scatter of the raw data in IHS space. Taxable asset income is mostly linear in wealth at high levels of wealth, but flattens out at low wealth levels.

Rewriting (A30) and adding the residual $\varepsilon^a_{i,t}$, we have:

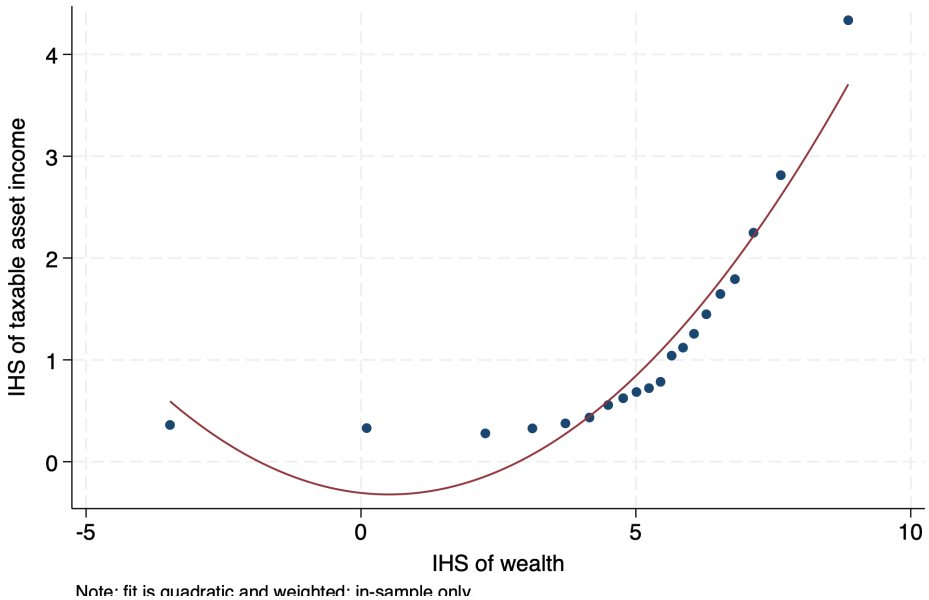
$$\begin{aligned} \text{IHS}(\tilde{y}^a_{i,t}) &= \text{IHS}(y^a_{i,t}) + \varepsilon^a_{i,t}, \quad \varepsilon^a_{i,t} \sim N(0, \sigma_a^2) \\ &= \xi^0_{i,t} + \xi^1_{i,t}\widehat{a}_{i,t} + \xi^2_{i,t} \max\{\widehat{a}_{i,t} - 3, 0\}^2 + \varepsilon^a_{i,t}, \end{aligned} \tag{A31}$$

with $\varepsilon^a_{i,t}$ independent of the other explanatory variables in equation (A31). As we document below, the parameters $\xi^0_{i,t}$ and $\xi^1_{i,t}$ are themselves functions of demographics and other explanatory variables. The balance sheet variables in the SCF give us a measure of $\widehat{a}_{i,t}$. (Because the income measures in the SCF are backward-looking, while the balance sheet variables date to the time of the interview, for this exercise we can treat our SCF-based measure of net worth as an end-of-period amount inclusive of any returns earned over that period).

Table A5 reports the estimated mapping from wealth to taxable asset income. The parameters for ξ^1 shown in specifications (1) and (2) imply a relationship between wealth and taxable income that is positive but weaker for Black households. This racial difference

[5]The latest version of TAXSIM can be accessed through https://www.nber.org/research/data/taxsim. We use v35.

Figure A2: Raw Data for Taxable Asset Income



Note: Authors' calculations using SCF data, the NBER TAXSIM program and Gale et al. (2022). Fitted line is quadratic and weighted, in-sample only.

might in part reflect sample construction. We remove observations with negative pre-tax income, which are more common at lower levels of wealth (see Figure A2), we may be truncating the relationship between wealth and taxable income at lower wealth levels, which are more common for Black individuals. We will use specification (3) in the main model.

We construct measures of pre-tax earnings, Social Security benefits, and inheritances/inter-vivos transfers and sum them into the income measure $\tilde{y}^e_{i,t}$. (We do not deduct medical expenses from this measure because: the SCF data lack this information; and medical expenses are tax-deductible only in excess of a substantial threshold.) The sum of $\tilde{y}^a_{i,t}$ and $\tilde{y}^e_{i,t}$ is our measure of taxable income. Figure A3 plots a bin scatter of the log of after-tax income $(\tilde{y}^a_{i,t} + \tilde{y}^e_{i,t} - \tilde{\tau}_{i,t})$ against the log of taxable income. The two variables have a near-linear relationship.

Rewriting (A28) and adding the residual $\varepsilon_{i,t}$, we have

$$\ln\left(\tilde{y}^a_{i,t} + \tilde{y}^e_{i,t} - \tilde{\tau}_{i,t}\right) = \ln(\lambda_{i,t}) + (1-\theta_{i,t})\ln(\tilde{y}^a_{i,t} + \tilde{y}^e_{i,t}) + \varepsilon_{i,t}, \quad \varepsilon_{i,t} \sim N(0,\sigma^2), \qquad \text{(A32)}$$

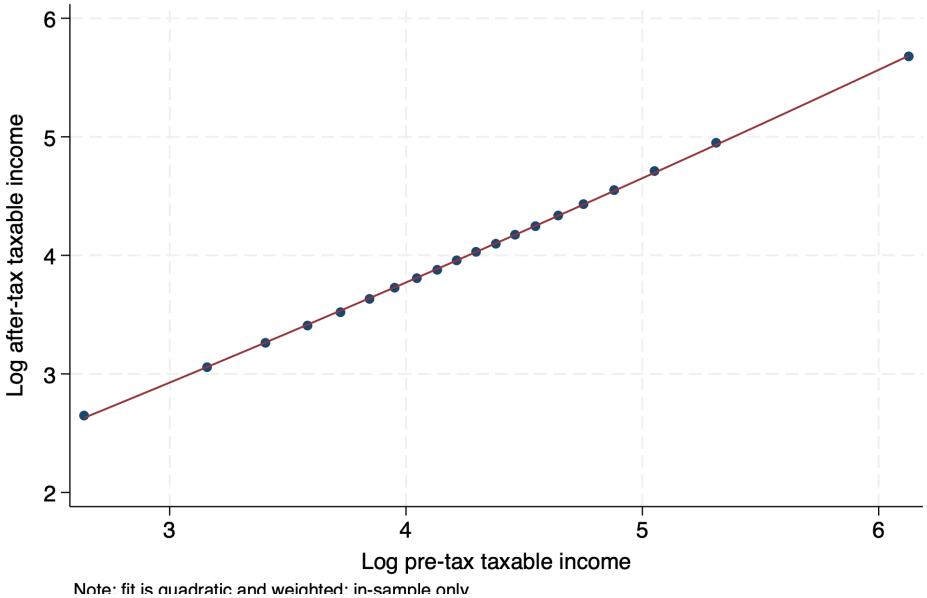
with $\varepsilon_{i,t}$ independent of the other explanatory variables in equation (A32). Parallel to equation (A31), $\lambda_{i,t}$ and $\theta_{i,t}$ are themselves functions of other explanatory variables.

Table A5: Parameter Estimates, Taxable Asset Income

	ξ^0 coefficients			ξ^1 coefficients			ξ^2 coefficients		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
Constant	0.319	0.314	-1.696	-0.020	-0.021	-0.162	0.118	0.119	-0.153
	(50.90)	(43.10)	(-5.80)	(-11.03)	(-10.53)	(-1.90)	(181.56)	(175.77)	(-2.72)
Black		0.030	-0.039		0.025	0.036		-0.023	-0.025
		(2.16)	(-2.76)		(5.08)	(7.47)		(-6.80)	(-7.96)
Married			-0.092			-0.000			-0.001
			(-7.87)			(-0.04)			(-0.62)
Education			0.017			-0.001			0.003
			(6.61)			(-1.14)			(11.34)
Number of kids			-0.000			-0.009			0.002
			(-0.08)			(-6.07)			(3.33)
Age/10 × 1[Age<67]			1.464			0.214			0.117
			(6.10)			(3.19)			(3.13)
Age^2/100 × 1[Age<67]			-0.408			-0.073			-0.018
			(-6.55)			(-4.33)			(-2.28)
Age^3/1000 × 1[Age<67]			0.040			0.007			0.001
			(7.71)			(5.01)			(1.41)
Age 67 or higher			3.468			0.227			0.189
			(11.49)			(2.61)			(3.36)
R^2	0.316	0.316	0.412	0.316	0.316	0.412	0.316	0.316	0.412
N	168313	168313	168308	168313	168313	168308	168313	168313	168308

Note: t-statistics in parentheses; the estimation sample is restricted to $\tilde{y}^a+\tilde{y}^e \geq 5000$ and $\tilde{y}^a+\tilde{y}^e-\tau \geq 5000$ in 2010 dollars (per-adult-equivalent).

Figure A3: Raw Data for Pre- and Post-Tax Income



Note: Authors' calculations using SCF data, the NBER TAXSIM program and Gale et al. (2022). Fitted line is quadratic and weighted, in-sample only.

Table A6: Parameter Estimates, Tax Function

	$\ln(\lambda)$ coefficients			θ coefficients		
	(1)	(2)	(3)	(1)	(2)	(3)
Constant	0.271	0.244	0.141	0.123	0.117	0.229
	(82.82)	(67.77)	(0.98)	(162.81)	(142.88)	(6.56)
Black		0.181	0.110		0.036	0.020
		(21.70)	(13.80)		(18.14)	(10.33)
Married			-0.089			-0.013
			(-16.13)			(-10.25)
Education			-0.020			-0.004
			(-17.83)			(-16.00)
Number of kids			0.107			0.020
			(36.03)			(28.60)
Age/10 $\times$ 1[Age<67]			0.429			-0.011
			(4.07)			(-0.43)
Age^2/100 $\times$ 1[Age<67]			-0.139			-0.003
			(-5.72)			(-0.56)
Age^3/1000 $\times$ 1[Age<67]			0.012			0.000
			(6.84)			(0.98)
Age 67 or higher			0.333			-0.103
			(2.29)			(-2.90)
R^2	0.944	0.944	0.953	0.944	0.944	0.953
N	168313	168313	168308	168313	168313	168308

Note: t-statistics in parentheses; the estimation sample is restricted to $\tilde{y}^a + \tilde{y}^e \geq 5000$ and $\tilde{y}^a + \tilde{y}^e - \tilde{\tau} \geq 5000$ in 2022 dollars (per-adult-equivalent).

Table A6 presents coefficient estimates for the tax function. Specifications (1) and (3) both imply progressivity ($\theta > 0$), but it is stronger in the richer specification. By way of comparison, Heathcote et al. (2017) estimate, using PSID data, $\theta = 0.181$ with an $R^2 = 0.91$. The results for our preferred specification, version (3), are close to this. The coefficient for a 45-year-old married White person with 1 child and 14 years of education is 0.172, while the R^2 for our regression is 0.96. Black households end up with a more progressive tax function in specification (2), but that is partly due to covariates, as in the richer specification (3) the coefficient on Black households is cut in half.

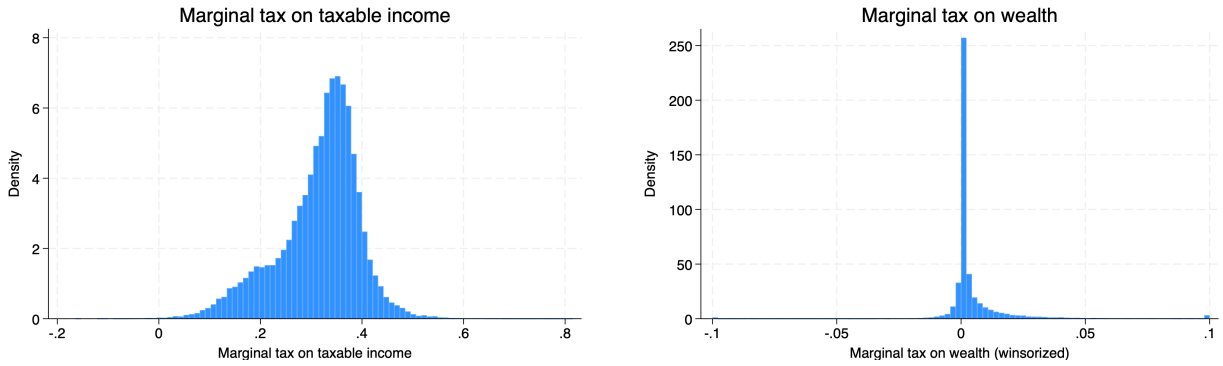
J.3 Marginal tax rates

Using the definition of IHS(a), the marginal tax rates for non-asset income and assets are given by, respectively:

$$mtr^e_{i,t} = 1 - \theta^0_{i,t}(1-\theta^1_{i,t})\big(\max\{y^a_{i,t} + y^e_{i,t}, 0\}\big)^{-\theta^1_{i,t}},$$
$$mtr^a_{i,t} = mtr^e_{i,t} \cdot \left(\frac{(y^a_{i,t})^2+1}{a^2_{i,t}+1}\right)^{1/2} \cdot (\xi^1_{i,t} + 2\,\xi^2_{i,t}\,\max\{\text{IHS}(a_{i,t}) - 3, 0\}).$$

Panel (a) of Figure A4 shows marginal tax rates for non-asset income. Noting that our tax rates include FICA and state and local taxes, the distribution is plausible. Panel (b) shows marginal tax rates for wealth. The tax rate on wealth is typically well below 1%, implying that the marginal tax rate on asset *income* is typically well below $0.01/r$.

Figure A4: Marginal Tax Rates on Non-asset Income and Wealth



(a) Non-asset income (b) Wealth

Note: Authors' calculations using SCF data, the NBER TAXSIM program and Gale et al. (2022). The marginal tax on wealth has been winsorized at ±10%. See Appendix J for details.

K Preference parameter estimates

Table A7: Preference Parameter Estimates

	Education-specific parameters				
Parameter	LTHS	HS	SC	CG	All
β: Discount factor					
β_L: Low value	0.812	0.832	0.610	0.849	
β_H: High value	0.812	0.957	0.838	1.018	
$\Pr(\beta = \beta_L)$ (in %)	99.4%	99.3%	8.2%	95.0%	
Couples premium (proportional)					1.12
Equivalence scale: $n = 1 + (\alpha_1 - 1)m_{i,t} + \alpha_2(\exp(\alpha_3 \cdot \#\text{kids}) - 1)$					
α_1: Marriage coefficient					0.98
α_2: Number of kids, scale					2.79
α_3: Number of kids, elasticity					0.196
Bequest parameters: $\upsilon_b(a) = \frac{\kappa_b^\sigma}{1-\sigma}(\kappa_b\chi_b + a)^{1-\sigma}$					
κ_b: Intensity parameter, singles					12.28
χ_b: Shift parameter, singles (in 1000s)					1.80
κ_b: Intensity parameter, couples					1.05
χ_b: Shift parameter, couples (in 1000s)					20.21
$\underline{h}$: Shift parameter, housing utility (in 1000s)					1,987
δ: Debt discharge rate, age < 65					0.000
δ: Debt discharge rate, age 65+					0.936
$\underline{z}$: Consumption floor for singles (in 1000s)					2.69
ς: Wife's share of household income					0.575
α_{sc}: Rate of return scale coefficient (multiplied by $\max\{0, \widehat{a} - \text{IHS}(295\text{K})\}$					0.00925

A key parameter is the wife's share of household earnings, ς, which we estimate to be 0.575—wives receive more than half of household earnings. This may reflect the tendency of widows to receive a significant portion of household wealth when their husbands die De Nardi et al. (2023). The model reproduces this feature by allowing wives to accumulate more wealth while married. Setting $\varsigma = 0.5$ has relatively little effect on our findings.

Another important parameter is α_{sc}, which controls the scale premium in returns, Our estimate of α_{sc} implies that, all else equal (including portfolio shares), the rate of return earned by a person with a net worth of \$2,000,000 is 1.755pp higher than the return

earned by a person with a net worth of \$300,000.

We capture this affect by adding the scale term $\alpha_{\text{sc}} \times \max\{0, \widehat{a} - \text{IHS}(295,000)\}$, where \$295,00 is the mean wealth holdings of White individuals.