Population Aging, Cohort Replacement, and the Evolution of Income Inequality in the United States*

Vesa-Matti Heikkuri[1] and Matthias Schief[2]

[1]Tampere University
[2]Economics Department, OECD

[Link to most recent version]

Abstract

This paper examines the impact of demographic change on household income inequality in the United States, both historically and prospectively. We emphasize the distinct roles of population aging and cohort replacement and develop a methodology to study their joint compositional effect. We document that cohorts born later in the 20th century embody higher levels of income inequality compared to earlier-born cohorts, and we argue that most of the increase in inequality over the past two decades can be accounted for by demographic change. Moreover, we predict that future demographic change will continue to put significant upward pressure on household income inequality in the United States.

Keywords: Demographic change, income inequality, compositional effects
JEL classification: C46, D31, J11

*Address correspondence to matthias.schief@oecd.org and vesa-matti.heikkuri@tuni.fi. We thank Joaquin Blaum, Pedro Dal Bó, James Feyrer, John Friedman, Oded Galor, Cecilia Garcia-Peñalosa, Jarkko Harju, Suvi Heikkuri, Stelios Michalopoulos, Sebastian Ottinger, Cosimo Petracchi, Tommaso Porzio, Devesh Rustagi, and David Weil for valuable comments and suggestions. We also thank seminar participants at the Growth Lab at Brown University, the Max Planck Institute for Demographic Research, the Stone Center on Socio-Economic Inequality, the PAA 2022 Annual Meeting, and the ECINEQ 2023 Annual Meeting. We acknowledge support from the James M. and Cathleen D. Stone Wealth and Income Inequality Project.

I Introduction

Income inequality in the United States is much higher today than half a century ago. A large body of empirical work has identified changes in the technological and institutional environment that have contributed to rising inequality, such as skill-biased technological change and de-unionization.[1] At the same time, the characteristics of the population have also changed substantially. Newer generations are smaller in size and live longer, leading to population aging. They are also generally better educated, tend to marry partners with similar education levels, and are selected differently for higher education and professional careers compared to earlier generations. Such demographic changes can affect inequality, even without technological or institutional changes.

In this paper, we study how demographic change affects household income inequality by altering the composition of the population in terms of age groups and birth cohorts. Our goal is to explain how demographic change has affected the evolution of income inequality in the past and to project how income inequality will evolve under future demographic change. While predicting future changes in the technological and institutional environment is difficult, demographics can be reliably projected several decades into the future.

We emphasize that demographic change affects inequality both through population aging and cohort replacement. Changes in the age structure of the population affect inequality because the distribution of income among older households differs systematically from that among younger households. Experience, on-the-job training, and subsequent retirement produce a hump-shaped age profile for average income. Moreover, heterogeneous returns to experience, persistent idiosyncratic shocks, and differential rates of wealth accumulation all imply that income inequality is higher among older households. Replacement of older cohorts by more recent birth cohorts affects income inequality because the distribution of income-relevant characteristics, such as the distribution of human capital (Altonji et al., 2012), the allocation of talent across professions (Hsieh et al., 2019), and the degree of positive assortative mating (Eika et al., 2019), differs between cohorts.

We study how demographic change affects household income inequality by considering the following question: How would income inequality evolve over time if the economic environment is held fixed in a given base year and only demographic change is allowed to occur? The goal of this thought experiment is to isolate the role of changing

[1] See, for example, Acemoglu and Autor (2011) and Firpo et al. (2009).

population characteristics while holding the returns to these characteristics fixed. In this thought experiment, households age within a static economic environment. New cohorts enter the economy with characteristics that are identical to those of the youngest cohort in the base year whose cohort-specific characteristics were shaped in the same economic environment. Older cohorts eventually exit the economy and the population age structure evolves as observed in the past or projected in the future.

A challenge to implementing our thought experiment is that cohorts likely differ in a large number of income-relevant characteristics, not all of which are readily observable. We therefore develop a parametric methodology that allows us to account for cohort differences in both observed and unobserved characteristics. Using household income data for the United States, we estimate life-cycle profiles and cohort differences in mean incomes and income Gini coefficients using an additive age-period-cohort model. We document important cohort differences in income distributions that are not accounted for by differences in age and educational attainment. The estimated cohort effects suggest that birth cohorts have become progressively more unequal in their income-relevant characteristics since the mid-20th century. We then use the estimated age and cohort effects to predict how the moments of subgroup income distributions evolve under demographic change when the economic environment is held fixed.

We derive the aggregate Gini coefficient from subgroup moments using a novel aggregation methodology that follows the principle of maximum entropy. In particular, we model subgroup income distributions, using the parametric distribution that maximizes entropy for given mean and Gini coefficient. This distribution imposes the least amount of information in addition to knowing the mean and Gini coefficient. We show that our methodology is able to aggregate subgroup Gini coefficients into the aggregate Gini coefficient with only limited loss of information. We then use this aggregation methodology on the predicted subgroup moments and population shares to study how demographic change affects the aggregate Gini coefficient.

We find that demographic change plays an important role in the evolution of household income inequality in the United States – both in the past and in the future. Our thought experiment suggests that the compositional effects of demographic change can account for all the increase in income inequality over the past two decades. Moreover, we predict that demographic change will further increase inequality in the near future, with our estimates suggesting an increase in the income Gini coefficient of between one and six percentage points by the year 2040.

Our methodology allows us to isolate the effect of population aging and cohort replacement. We find that both population aging and cohort replacement have contributed substantially to the rise in household income inequality in the recent past. However, projected further aging of the US population in the near future will not affect household income inequality. Instead, the predicted increase in inequality will be driven almost exclusively by cohort replacement.

In our main analyses, we estimate cohort differences under a range of normalizations for linear trends in the age-period-cohort model. Acknowledging the important role of the choice of normalization for our results, we assess the plausibility of the estimated cohort differences using additional data for selected birth cohorts. Specifically, we use data on birthplace, education, spousal and parental education, race, as well as cognitive test scores from the National Longitudinal Surveys of Youth (NLSY) for the birth cohorts 1957-1961 and 1980-1984 to estimate the influence of changing cohort characteristics on income inequality. Adjusting the characteristics of the older NLSY79 cohorts to reflect those of the younger NLSY97 cohorts reveals important differences in inequality for both college-educated and non-college-educated households. Reassuringly, these differences align closely with the cohort differences implied by our preferred choice of normalization in the age-period-cohort model.

Our findings suggest a more important role for demographic change in explaining the evolution of household income inequality than is generally found in the literature (e.g., Kuhn et al. (2020)). Previous studies of compositional effects typically rely on using re-weighting methods, following DiNardo et al. (1996), and explicitly account for differences in only a limited number of characteristics. For comparison, we also implement our thought experiment using a re-weighting method that accounts for compositional changes in terms of the age structure and educational attainment of the US population. In this exercise, we find no effect of demographic change on household income inequality; however, we argue that the results from this re-weighting analysis are misleading. By assigning more weight to older households in a given cross-section to track population aging, we also assign more weight to earlier and more equal birth cohorts. As a result, re-weighting fails to capture the effect of cohort replacement and confounds the effect of population aging.

The rest of the paper is structured as follows. In section I.1, we discuss how our paper relates to the existing literature, and we introduce our data sources in section I.2. In section II, we develop a parametric method to implement the thought experiment. In section III, we present the main results. Section IV implements the thought experiment using a re-weighting analysis and discusses its shortcomings. Section V concludes.

I.1 Related literature

Recent papers on the compositional effects of demographic change on income and wealth distributions include Kuhn et al. (2020) and Auclert et al. (2021), who study the effects of population aging, and Eika et al. (2019), who study the impact of changing household characteristics.[2] Kuhn et al. (2020) assemble a new micro data set for household income and wealth in the United States going back to 1949 and study, among other things, the effect of demographic change on income and wealth inequality in the past. They find that population aging has moderately increased income inequality throughout their sample period. Auclert et al. (2021) use population projections to predict the compositional effect of demographic change on the future evolution of the wealth-to-output ratio in the United States and a number of other countries. They predict that population aging will have a significant impact on the wealth-to-output ratio in the United States over the coming decades. Eika et al. (2019) study the role of educational assortative mating on household income inequality. They find that educational assortative mating accounts for a non-negligible share of cross-sectional inequality but that the trend in sorting has hardly affected income inequality. They also find that the increases in college attendance and completion rate by women have slowed the increase in household income inequality.

There is also a literature on compositional effects of a changing population structure on wage inequality, where the focus has predominantly been on the skill composition of the population and the role of skill-biased technical change (Juhn et al., 1993; Lemieux, 2006; Autor et al., 2008; Altonji et al., 2012; Hoffmann et al., 2020). Lemieux (2006), for example, studies how changes in the composition of the US population in terms of experience and educational attainment affect residual wage inequality, using a re-weighting analysis. He finds that increases in within-group inequality are concentrated in the 1980s and that the increase in population-level wage inequality in the subsequent decade is driven by composition effects.

Cohort differences in income distributions play an important role in our paper. The role of cohort differences in the dynamics of inequality are also emphasized by Guvenen et al. (2022), who show that increased inequality in lifetime incomes is mostly resulting from more recent cohorts having higher initial income inequality while the increase in inequality over the life cycle remains similar across cohorts. Unlike Guvenen et al. (2022), we study the effect of cohort differences on the evolution of cross-sectional

[2] Older papers in this literature include Burtless (1999), Daly and Valletta (2006), Larrimore (2014), and Greenwood et al. (2014).

inequality and also discuss the implications of cohort differences for future inequality.

Altonji et al. (2012) study differences in income-relevant characteristics between the cohorts surveyed in the NLSY79 and NLSY97 surveys. They find that the skill distribution has widened between cohorts born around the year 1960 and the cohorts born in the early 1980s and use this observation to predict that wage inequality will increase substantially by 2025. We also use data from the NLSY79 and NLSY97 surveys and show that the more recent cohorts have characteristics that lead to higher inequality in household incomes compared to the earlier cohorts. Our paper differs from Altonji et al. (2012) in that we extend estimates of cohort differences in income distributions to cover all cohorts born between 1888 and 1994 using repeated cross-sections from Current Population Survey and Survey of Consumer Finances data and an age-period-cohort model. We also study the role of demographic change in the past as well as in the future.

Secular trends in cohort-specific characteristics giving rise to important cohort differences have been studied in a number of papers. Card and Lemieux (2001) attribute the rising college premium to a slowdown in educational attainment for cohorts born after 1950. Hendricks and Schoellman (2014) explain the same phenomenon by documenting growing test-score gaps between college-bound and non-college-bound students. More recently, Hsieh et al. (2019) argue that cohort-specific improvements in the allocation of talent have contributed significantly to US economic growth. Similarly, the literature on structural change has documented that a large share of labor reallocation can be accounted for by new cohorts entering growing industries (Lee and Wolpin, 2006; Hobijn et al., 2019; Porzio et al., 2022).

An additional source of cohort differences in income distributions that has recently received increased attention is scarring. The literature on this topic has documented long-lasting negative effects on earnings and employment for cohorts entering the labor market in a bad economy, and common findings are that these effects are heterogeneous and therefore affect inequality (Raaum and Røed, 2006; Kahn, 2010; Oreopoulos et al., 2012; Rothstein, 2019; Schwandt and Von Wachter, 2019).

To construct counterfactuals, we estimate how income distributions depend on age and birth cohort, using a standard age-period-cohort model. In this respect, our paper is also related to the literature devoted to studying life-cycle profiles of economic inequality. In particular, we build on Deaton and Paxson (1994a,b), who estimate age profiles for within-cohort income and consumption variance in the United States and propose a normalization for dealing with the linear dependence of age, period, and cohort effects.[3] Heathcote et al. (2005) point out the importance of the choice of normalization in estimating the age profile of income inequality. To deal with this issue, we follow Lagakos et al. (2018), who suggest exploring the results under a range of different normalizations.

I.2 Data sources

In our main analyses, we use data on household income for the years 1968-2020 from the Current Population Survey (CPS). Our measure of household-level income is the total money income during the previous calendar year of all adult household members.[4] Total money income is the sum of wages and salaries, income from professional practice and self-employment, rental income, interest, dividends, and transfer payments, as well as business and farm income. We complement the CPS data with a longer series of harmonized repeated cross-sections based on archival data from historical waves of the Survey of Consumer Finances that was recently made available by Kuhn et al. (2020). This data set spans the time period 1949-2019 and reports household-level total income, which has the same definition as total money income in the CPS. We follow Kuhn et al. (2020) and refer to these data as the SCF+ data set.[5] The CPS and SCF+ data sets complement each other. While the CPS data set has a larger sample size, the SCF+ data cover two more decades.

About 0.16% of all households in the CPS data and 0.31% of all households in the SCF+ data report negative total incomes.[6] We censor the income distribution by recoding negative values as zeros.[7] To avoid problems associated with topcoding in the CPS and SCF+ data sets, we focus on income inequality among the bottom 99%.

For the CPS data, we use annual surveys between 1968 and 2020. For the SCF+ data, we use triennial waves constructed by Kuhn et al. (2020), which leaves us with data for every third year between 1950 and 2019, with the exception of the years 1974,

[3]Other papers that also use an age-period-cohort model to study life-cycle behavior include Attanasio (1998), Gourinchas and Parker (2002), Storesletten et al. (2004), Low et al. (2010), Huggett et al. (2011), Aguiar and Hurst (2013), Heathcote et al. (2014), Lagakos et al. (2018), and Jedwab et al. (2023).

[4]Specifically, we use the variable HHINCOME from the IPUMS CPS harmonized microdata (Ruggles et al., 2020).

[5]The data set in Kuhn et al. (2020) covers the time period 1949-2016. We added to this data set the 2019 Survey of Consumer Finances.

[6]After applying sampling weights, households with negative income make up 0.13% of the population in both the CPS and the SCF+ data sets.

[7]Negative income levels pose a challenge for the interpretation of differences in income inequality across different subgroups of the population because they can inflate the Gini coefficient even if the dispersion of income is low.

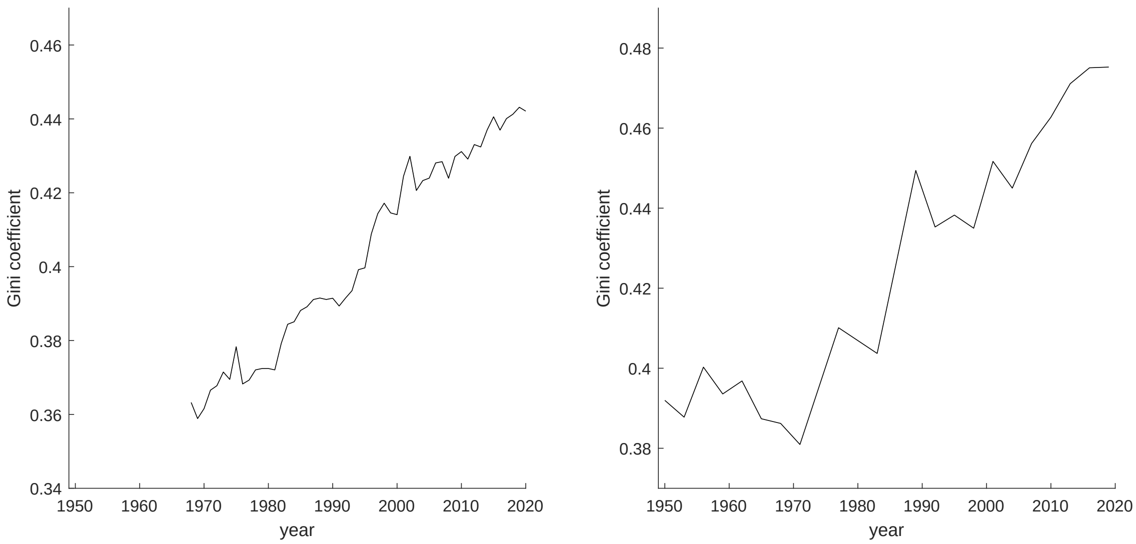

(a) CPS **(b)** SCF+

Figure 1: Evolution of household income inequality in the United States.

1980, and 1986. We correspondingly aggregate the data into three-year age groups and birth cohorts. We assign households to their respective birth cohort based on the age of the household head. Furthermore, we restrict our attention to households in which the household head is between 26 and 79 years old in the CPS data and between 26 and 80 in the SCF+ data.[8] Figure 1 shows the evolution of household income inequality in the CPS and SCF+ data after applying our sample restrictions.

We also use data from the NLSY79 and NLSY97 surveys to study cohort differences in household income distributions for selected cohorts. We restrict the sample to birth cohorts born in 1957-1961, which we refer as the older cohort, and 1980-1984, which we refer as the younger cohort. We use the 1996 wave from NLSY79 when the older

[8]The definition of the household head in the SCF+ data is the male partner in mixed-sex couples, the older partner in same-sex couples, and the single core individual in households without a core couple. In the CPS data, the definition of household head changed in the year 1980. While it was similar to the SCF+ definition in the years prior to 1980, the CPS has since discontinued the use of the term "household head" and has replaced it with "householder." A householder is the person in whose name the housing unit is owned or rented, or in the case of a married couple jointly owning or renting the house, it is either spouse. In our main results, we assign households to their respective age groups based on the age of the individual who is considered the household head/householder in the respective data set. As a robustness check, we re-define household heads in the CPS to be always the male partner in mixed-sex couples and the older person in same-sex couples, to be consistent with the definition in the SCF+ data. The results hardly change under this alternative definition. Finally, we drop household heads aged 80 in the CPS data from the analysis because individuals that are older than 80 are coded as 80 for some waves of the survey.

cohort was aged 35-39, and the 2019 wave from NLSY97 when the younger cohort was similarly aged 35-39. We drop observations for which family income is not observed and focus on inequality among the bottom 95% to avoid issues with topcoding.

To track the evolution of the age composition of the US population, we rely on population projections from the US Census Bureau. The median age in the US population has increased from 27 in 1970 to 38 in 2019, and is predicted to increase to 42 by the year 2060.[9] The leftmost panel in figure 2 shows the age distribution in the US population in the years 1970 and 2010, as well as projections for the year 2050. Over this time period, older people progressively make up a higher fraction of the US population.

The middle and the right panels show the corresponding age distribution among household heads in the CPS and the SCF+ data sets. As we want to study the role of demographic change not only in the past but also in the future, we need to translate the predicted changes in the age structure of the US population into corresponding changes in the age structure of household heads in the survey data. We do this by using the constant headship rate method. That is, we compute the probability that an individual of a given age in the latest survey wave is recorded as the household head, and we assume that these probabilities remain fixed in the future.

II Methodology

In this section, we develop a parametric approach to implement our thought experiment of holding the economic environment fixed while allowing demographic change to take place. The methodology proceeds in three steps. First, we estimate the life-cycle profiles of average income and income inequality for different birth cohorts and education groups. Second, we use these estimates to model how subpopulation income distributions evolve as cohorts age. In line with our thought experiment of fixing the economic environment, we assume that cohorts entering the economy after any given base year have the same characteristics as the youngest cohort in that base year and thus share the same life-cycle profiles of average income and income inequality. Third, we use predicted population shares to construct a population-level income distribution and to study the evolution of its Gini coefficient. In the remainder of this section, we discuss each of these steps in detail.

[9]United States Census Bureau (2017).

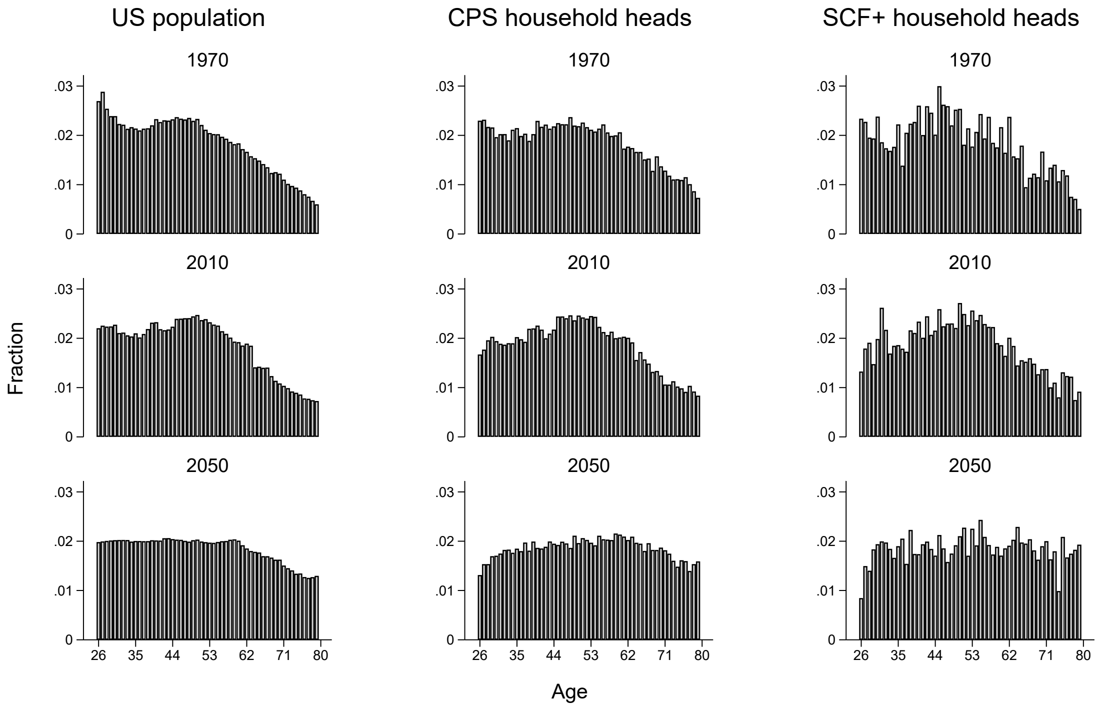

Figure 2: The evolution of the age composition of the US population and the implied age distributions of household heads in the CPS and the SCF+ data sets.

II.1 Estimating age, period, and cohort profiles

We assume that two key moments of the income distribution, the logarithms of the mean and the Gini coefficient, are described by an additively separable age-period-cohort model. This allows us to use repeated cross-sections in the CPS and SCF+ data sets to estimate how the distribution of income differs across birth cohorts and how it evolves within cohorts as they age. We motivate the age-period-cohort model by showing that a simple income process leads to additively separable age, period, and cohort profiles in the logarithms of mean income and the income Gini coefficient.

II.1.1 A simple income process

In each period, the income distribution of households with given age and education receives an income shock that has two components: a level component and an inequality component. The level component increases or decreases all incomes by a given factor, while leaving inequality between the households unchanged. The inequality component stretches or compresses the income distribution, while leaving the average income unchanged. In particular,

$$y_{i,t} = \big(1+\beta(e,a,t)\big)\Big(y_{i,t-1} + \gamma(e,a,t)\big(y_{i,t-1} - \mathsf{E}[y_{i,t-1}|e,a]\big)\Big)$$

where $y_{i,t}$ denotes household i's income in period t, and $\beta \in [-1, \infty)$ and $\gamma \in [-1, \infty)$ denote the level and inequality component, respectively; e and a are the level of education and the age of the household head; and $\mathsf{E}[y_{i,t-1}|e,a]$ is the mean income in the household's demographic subgroup. A positive β means that the shock increases average income, while a positive γ means that the shock increases inequality. Negative values of β and γ achieve the opposite.

We assume that the level and the inequality components of the income shocks are separable in age and period, so that

$$\begin{aligned} 1+\beta(e,a,t) &= \big(1+\beta_a(e,a)\big)\big(1+\beta_t(e,t)\big), \\ 1+\gamma(e,a,t) &= \big(1+\gamma_a(e,a)\big)\big(1+\gamma_t(e,t)\big). \end{aligned}$$

Finally, we allow the initial income distribution for each cohort to be arbitrary, reflecting differences due to characteristics such as assortative mating, sorting into occupation, or the quality of education.

With this income process, we can model how changes in the technological and institutional environment affect the distribution of income in the economy. For example, skill-biased technological progress that disproportionately increases the incomes of high earners corresponds to an income shock where β_t and γ_t are both greater than zero. A recession on the other hand corresponds to an income shock where β_t is negative. These income shocks also allow us to model life-cycle dynamics. For example, accumulation of labor market experience that complements skills is captured by positive β_a and γ_a.

We obtain an expression for the income of household i in terms of initial income and the complete history of income shocks:

$$\begin{aligned} y^e_{i,a,t} = &\prod_{k=1}^{a}\Big(1+\beta_a(e,k)\Big)\prod_{k=t-a+1}^{t}\Big(1+\beta_t(e,k)\Big) \\ &\left[y^e_{i,0,t-a} + \left(\prod_{k=1}^{a}\Big(1+\gamma_a(e,k)\Big)\prod_{k=t-a+1}^{t}\Big(1+\gamma_t(e,k)\Big) - 1\right)\Big(y^e_{i,0,t-a} - \mathsf{E}[y^e_{0,t-a}]\Big)\right]. \end{aligned}$$

Computing the mean and Gini coefficient within a given demographic subgroup reveals that this income process generates additively separable age, period, and cohort profiles for the logarithms of mean income and the income Gini coefficient:

$$\ln\left(\mathsf{E}[y^e_{i,a,t,c}]\right) = \underbrace{\sum_{k=1}^{a}\ln\left(1+\beta_a(e,k)\right)}_{\text{age effect}} + \underbrace{\sum_{k=1-a^{\max}+1}^{t}\ln\left(1+\beta_t(e,k)\right)}_{\text{period effect}} + \underbrace{\ln\mu_0^{e,c} - \sum_{k=1}^{c}\ln\left(1+\beta_t(e,k)\right)}_{\text{cohort effect}}$$

and

$$\ln\left(G(y^e_{i,a,t,c})\right) = \underbrace{\sum_{k=1}^{a}\ln\left(1+\gamma_a(e,k)\right)}_{\text{age effect}} + \underbrace{\sum_{k=1-a^{\max}+1}^{t}\ln\left(1+\gamma_t(e,k)\right)}_{\text{period effect}} + \underbrace{\ln G_0^{e,c} - \sum_{k=1}^{c}\ln\left(1+\gamma_t(e,k)\right)}_{\text{cohort effect}},$$

where $\mu_0^{e,c}$ and $G_0^{e,c}$ are the initial mean and the Gini coefficient of birth cohort $c = t - a$ with education e before receiving any income shocks; and $a^{\max}$ is the maximum age.[10] Hence, this income process is equivalent to a standard age-period-cohort model for the logarithms of mean income and the Gini coefficient.

II.1.2 Age-period-cohort model

We partition our main sample into year-by-age-by-education subsamples and compute the mean and Gini coefficient in each subsample. As a result, we obtain two balanced panels in age and survey waves—one for households with a college-educated household head and another for households without a college-educated household head. We model the income moments as being generated by additive age, period, and cohort effects.

The model can be written as

$$M_{apc} = \alpha_a + \pi_p + \kappa_c + \varepsilon_{apc}, \tag{1}$$

where M_{apc} is the observed moment at age a, in period p, and in cohort c. The age, period, and cohort effects are captured by α_a, π_p, and κ_c, respectively. A mean zero

[10] The additive separability for the logarithm of the Gini coefficient follows from the fact that the inequality shock multiplies the Gini coefficient by $1+\gamma$, see proposition 6 in Heikkuri and Schief (2022). Note also that additive separability does not strictly require that the income of each household evolves according to this income process as long as the percentiles of the income distributions follow this process. Put differently, only the distribution of incomes matters—not the positions of individual households in it.

error term, ε_{apc}, captures both sampling variance and unmodeled noise.[11] We specify this model separately for college-educated and non-college-educated households. Note that in this model we are not imposing any functional form on the age, period, and cohort profiles.

Unfortunately, equation (1) is not identified and cannot be estimated from the data. An obvious problem is that we need to normalize at least one each of the age, period, and cohort effects. A more fundamental identification problem, however, arises because of the linear dependency between age, period, and cohort. The nature of the identification problem can be seen more clearly when we decompose age, period, and cohort effects into two parts: (1) linear trends in age, period, and cohort, and (2) fixed effects that capture deviations from these trends. In particular, if we require that the fixed effects sum to zero and be orthogonal to a trend,[12] then the model can be rewritten as

$$M_{apc} = \theta + \alpha\, a + \pi\, p + \kappa\, c + \check{\alpha}_a + \check{\pi}_p + \check{\kappa}_c + \varepsilon_{apc} \tag{2}$$

with the following restrictions on the parameters:

$$\sum_a \check{\alpha}_a\, a = 0 \quad \text{and} \quad \sum_a \check{\alpha}_a = 0, \tag{3}$$

$$\sum_p \check{\pi}_p\, p = 0 \quad \text{and} \quad \sum_p \check{\pi}_p = 0, \tag{4}$$

$$\sum_c \check{\kappa}_c\, c = 0 \quad \text{and} \quad \sum_c \check{\kappa}_c = 0. \tag{5}$$

In this formulation, the overall trends in the age, period, and cohort profiles are captured by the coefficients α, π, and κ, respectively, and θ is a constant. Note that while we cannot estimate all three linear trends in the age, period, and cohort profiles, the deviations from the linear trends are identified and can be estimated from the data

[11] The average sample size for a given survey year in the SCF+ data is only about 43% of the average sample size in the CPS data, causing larger sampling variation in the estimated mean income levels and income Gini coefficients at the subgroup level. Sampling variation induces classical error in our dependent variables and does not bias our results. However, the lack of precision in the SCF+ data still leads to unwelcome uncertainty in our estimation results, especially for very early and very late birth cohorts that are observed less often in our sample. We address this problem by including neighboring age groups when estimating average income levels and income Gini coefficients at the subgroup level in the SCF+ data. For example, when we compute the income Gini coefficient for birth cohorts 1959-1961 in the year 2001, we include not only households with household heads aged 40-42, but also those with households heads aged 37-40 and 43-46.

[12] This normalization of the age, period, and cohort effects is due to Deaton and Paxson (1994b).

even without knowing what the linear trends are. Moreover, differences between the linear trends in the age, period, and cohort profiles are identified, so that normalizing any one of the linear trends in equation (2) is enough to estimate this model. For example, if we set $\pi = \kappa$, as we do in our baseline normalization, then our estimated parameters are

$$\hat{\alpha} = \alpha \; + \frac{\pi \qquad \kappa}{2} \tag{6}$$

$$\hat{\pi} = \hat{\kappa} = \frac{\pi \; + \kappa}{2} \tag{7}$$

where $\hat{\alpha}$, $\hat{\pi}$, and $\hat{\kappa}$ are the estimated linear trends for age, period, and cohort, and α , π , and κ are the true linear trends in the data generating process.

II.1.3 Normalizing the linear trends

The normalization of the linear trends does not affect the predicted values of the model, and therefore, it cannot be estimated from the data. To address this issue, we follow Lagakos et al. (2018) and derive our results under three different normalizations. As extreme cases, we set either the period or the cohort trend in the income Gini coefficient to zero. As an intermediate case, we assume that the trends in period and cohort effects are equal in magnitude.[13] We treat the intermediate case as our baseline normalization.[14] In section III.4, we derive an additional normalization from estimated cohort differences using the NLSY data. Although the normalization matters quantitatively, our qualitative findings are invariant to the choice of normalization.

We consider a trend in both period and cohort effects plausible. Technological progress that increases incomes independently of age and education naturally constitutes a trend in period effects for average income. However, increasing levels of human capital, to the extent they are generated by more schooling and higher quality of edu-

[13] Lagakos et al. (2018) also consider a fourth type of normalization. In the case of labor income, one can argue that average income should not increase right before retirement due to low incentives to human capital accumulation. Thus, choosing a linear trend that makes the age profile flat before retirement age provides an alternative normalization for the linear trends. This "flat spot" strategy is also used in Heckman et al. (1998), McKenzie (2006), Huggett et al. (2011), and Bowlus and Robinson (2012). We do not use this method because it is not obvious that the age profile of average household income should be flat before retirement. Households may be still be saving and thus increasing their capital income before retirement and there may also be changes in transfer income. Similarly, there is no reason to assume that income inequality should be flat before retirement.

[14] In the case of log mean income, we always use the baseline normalization. We have also derived our results under different normalization of the trends in log mean income. The results do not vary meaningfully.

cation or cohort-specific improvements in health, are a cohort trend for average income. Jones (2002) calculates that almost a third of the growth in GDP per capita between 1950 and 1993 can be accounted for by rising educational attainment, while roughly two-thirds is accounted for by improved productivity.[15] Moreover, Hsieh et al. (2019) show that a significant share of productivity growth is explained by improved sorting of talent, which occurs between cohorts.

Assessing trends in income inequality is more difficult. To the extent that rising inequality is driven by technology or policies that increase inequality independent of household characteristics, it should be captured by period effects. A reduction in redistribution from rich to poor, for example, is a period effect that increases inequality. On the other hand, increasing inequality due to secular trends in cohort-specific characteristics, such as the distribution of human capital, sorting to education and occupations, or positive assortative mating, should be attributed to the trend in cohort effects. Since a trend in both period and cohort effects is plausible, we consider the intermediate normalization of equal period and cohort trends in income inequality a sensible baseline normalization.[16]

II.1.4 Estimation results

The additively separable age-period-cohort model is able to account for a large share of the variation in mean incomes and income Gini coefficients at the subgroup level. Focusing on the Gini coefficients for households without a college-educated household head, for example, we plot the life-cycle profiles for selected birth cohorts and the time trends for selected age groups in figure 3. Inequality increases with age and more recent cohorts experience higher levels of inequality at any given age. Moreover, age-specific inequality levels have increased over time. Importantly, however, the increase is much more pronounced among younger ages, especially during the 1970s and 1980s. The age-period-cohort model is able to account for these differential patterns in the data.

Moving beyond the model fit for a selected demographic subgroup, table 1 reports the coefficients of determination for all demographic subgroups as well as for mean

[15] Jones (2021) updates these shares to one-quarter and two-thirds for years between 1950 and 2007. The remainder is explained by the increase in the ratio of labor force to population.

[16] Huggett et al. (2011) find that more than 60% of lifetime earnings and wealth inequality is due to characteristics that are fully formed by early adulthood. It is therefore plausible that a significant share of the rising income inequality can be explained by changing cohort-level characteristics.

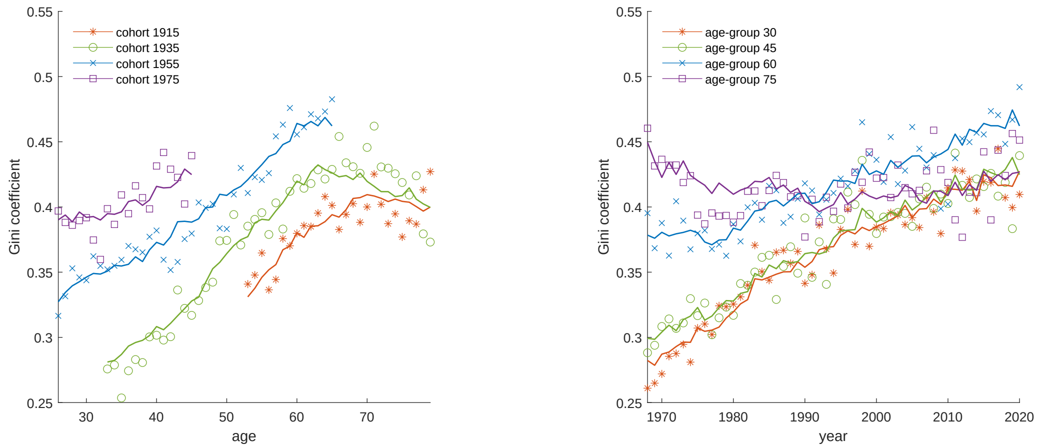

(a) Age profiles for selected birth cohorts **(b)** Time trends for selected age groups

Figure 3: Gini coefficients for households without a college-educated household head (CPS data). The solid lines show the fit of the age-period-cohort model.

income.[17] R^2 exceeds 90% in all estimated models. We further decompose R^2 into the respective parts accounted for by the linear trends in age, period, and cohort, and the deviations from these trends in the age, period, and cohort profiles using a Shapley decomposition. Nonlinear cohort effects account for between 11% and 30% of the explained variation in the estimated models, underscoring the importance of cohort differences in the process of demographic change. These cohort effects are not driven by differences across cohorts in the share of college-educated households as the model is estimated separately for college and non-college educated households. Note that the results regarding model fit do not depend on the choice of normalization for the linear trends in age, period, and cohort.

In figures 4 and 5, we plot the age, period, and cohort profiles of mean income and income inequality estimated separately on the CPS and the SCF+ data sets. The profiles are plotted after normalizing the linear trends in cohort and period effects to be equal in magnitude. As discussed in section II.1, nonlinearities in the profiles are identified from the data and do not depend on the normalization of the linear trends.[18]

We find that average income increases with age until about age 50 and decreases thereafter. Similarly, income inequality also increases over the working life and plateaus after retirement age. Compared to households without a college-educated household

[17] The table also reports the number of observations in each model, which is given by the number of year-age combinations in each data set. We include analogues of figure 3 for college-educated households and for mean income in appendix C.

[18] In appendix A, we plot the profiles for alternative normalizations of the linear trends.

Panel A: CPS	(Log) mean income		(Log) income Gini	
	College	Non-college	College	Non-college
N	2,862	2,862	2,862	2,862
R^2	0.95	0.97	0.92	0.90
Shapley decomposition of R^2				
Linear trends	0.23	0.32	0.80	0.74
Nonlinear age effects	0.55	0.43	0.05	0.07
Nonlinear period effects	0.03	0.03	0.03	0.02
Nonlinear cohort effects	0.19	0.22	0.11	0.17
Panel B: SCF+	(Log) mean income		(Log) income Gini	
	College	Non-college	College	Non-college
N	399	399	399	399
R^2	0.94	0.97	0.90	0.91
Shapley decomposition of R^2				
Linear trends	0.15	0.36	0.54	0.46
Nonlinear age effects	0.41	0.23	0.03	0.09
Nonlinear period effects	0.24	0.17	0.19	0.15
Nonlinear cohort effects	0.21	0.24	0.25	0.30

Table 1: Sample sizes and explained variation in the age-period-cohort model.

head, households with a college-educated household head experience a steeper increase in income over the working life and a sustained increase in income inequality even at older ages. We estimate very similar age profiles in both data sets.

In the period profiles, we document cyclical movement in log mean income reflecting business cycles, which can be seen more clearly in the annual CPS data. The period profiles also document a stark increase in income inequality in the early 1980s among non-college-educated households that is followed in the 1990s by a similar increase for college-educated households. Interestingly, while income inequality at the population level has increased substantially in the past two decades, we find much less of an increase or even a decrease in period effects for income inequality over this time period. Finally, when we estimate period effects going back to 1950 in the SCF+ data, we find that period effects for the income Gini coefficient show little evidence of a trend before the 1980s.

Most strikingly, we find pronounced nonlinear cohort profiles for both average income and income inequality. The cohort effects for log mean income increase up to approximately the birth cohort 1947 and become flat or decrease thereafter. In contrast, the cohort effects for income inequality follow a U-shaped profile and increase during the second half of the 20th century. These patterns are found in both the CPS and the SCF+ data.

A potential concern is that the estimated profiles are affected by the fact that we observe different cohorts in different time periods and at different ages, and that we observe some cohorts more often than others. We address this concern by re-estimating the age-period-cohort model on different restricted time windows and by comparing the resulting cohort profiles. For both data sets, we find that the estimated profiles from restricted samples align well with the ones from the full sample. We describe the procedure in more detail and show the estimated profiles in appendix B.

II.2 Constructing counterfactual moments

We use the estimated age, period, and cohort effects for log mean income and the log Gini coefficient to construct a counterfactual mean, $\tilde{\mu}_{a,p,c,e}$, and a Gini coefficient, $\tilde{g}_{a,p,c,e}$, for each subgroup as implied by our thought experiment. We give each age-education group its corresponding age effect, the period effect of the base year, and the estimated cohort effect for cohorts that are present in the base year. The cohorts that enter the economy after the base year are given the cohort effect of the youngest age-group a_0 in base year $\bar{p}$. In particular, for base year $\bar{p}$ and target year $p^{0} > \bar{p}$, we set the subgroup moments as

$$\tilde{\mu}_{a,p^{0},c,e} = \begin{cases} \exp\left(\theta_e^{\mu} + \alpha_e^{\mu} a + \pi_e^{\mu}\bar{p} + \kappa_e^{\mu} c + \check{\alpha}_{a,e}^{\mu} + \check{\pi}_{p,e}^{\mu} + \check{\kappa}_{c,e}^{\mu} + \frac{\sigma_{e,\mu}^2}{2}\right) & \text{if } c < \bar{c}_0 \\ \exp\left(\theta_e^{\mu} + \alpha_e^{\mu} a + \pi_e^{\mu}\bar{p} + \kappa_e^{\mu} \bar{c}_0 + \check{\alpha}_{a,e}^{\mu} + \check{\pi}_{p,e}^{\mu} + \check{\kappa}_{c_0,e}^{\mu} + \frac{\sigma_{e,\mu}^2}{2}\right) & \text{if } c \quad \bar{c}_0, \end{cases} \tag{8}$$

$$\tilde{g}_{a,p^{0},c,e} = \begin{cases} \exp\left(\theta_e^{g} + \alpha_e^{g} a + \pi_e^{g}\bar{p} + \kappa_e^{g} c + \check{\alpha}_{a,e}^{g} + \check{\pi}_{p,e}^{g} + \check{\kappa}_{c,e}^{g} + \frac{\sigma_{e,g}^2}{2}\right) & \text{if } c < \bar{c}_0 \\ \exp\left(\theta_e^{g} + \alpha_e^{g} a + \pi_e^{g}\bar{p} + \kappa_e^{g} \bar{c}_0 + \check{\alpha}_{a,e}^{g} + \check{\pi}_{p,e}^{g} + \check{\kappa}_{c_0,e}^{g} + \frac{\sigma_{e,g}^2}{2}\right) & \text{if } c \quad \bar{c}_0, \end{cases} \tag{9}$$

where $\bar{c}_0 := \bar{p} \quad a_0$ is the youngest cohort present in the base year, superscripts μ and g indicate the statistical moment and subscript e the education group for which the parameters have been estimated, and $\sigma_{e,\mu}^2$ and $\sigma_{e,g}^2$ are the estimated variances of the

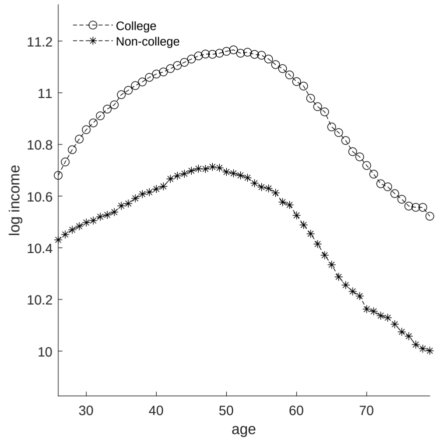

(a) Age profile: log mean income

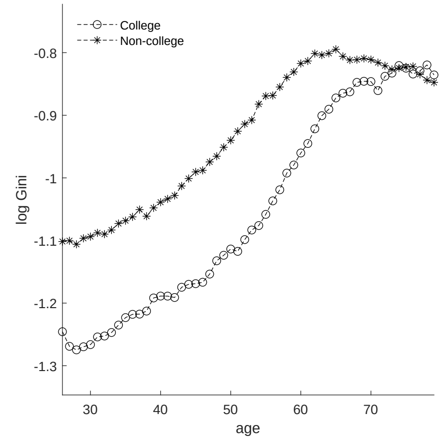

(b) Age profile: log Gini coefficient

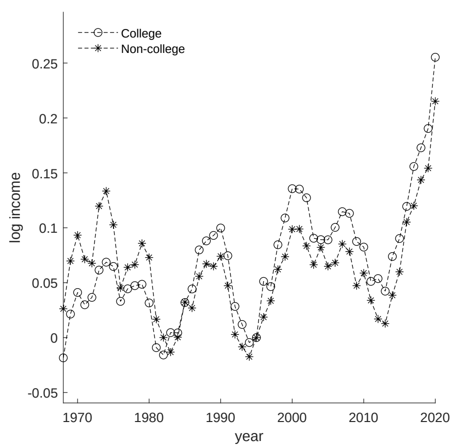

(c) Period profile: log mean income

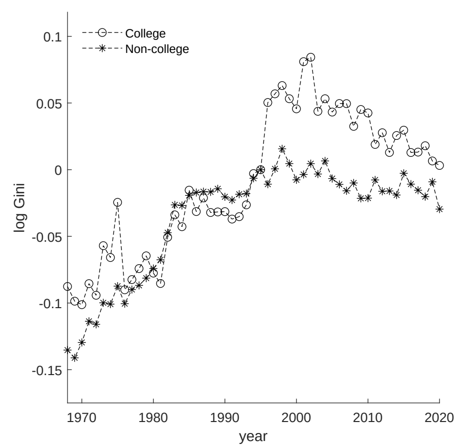

(d) Period profile: log Gini coefficient

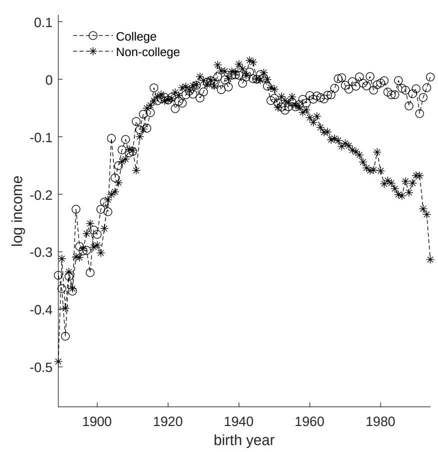

(e) Cohort profile: log mean income

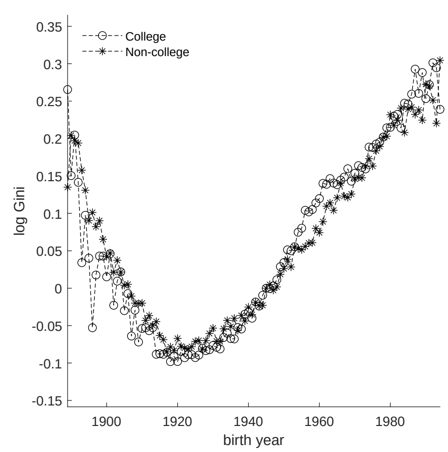

(f) Cohort profile: log Gini coefficient

Figure 4: Age, period, and cohort profiles of log mean income and log Gini coefficients in the CPS data. The period effect 1995 and cohort effect 1945 are normalized to zero.

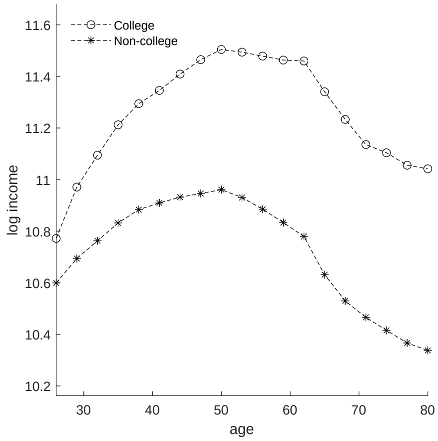

(a) Age profile: log mean income

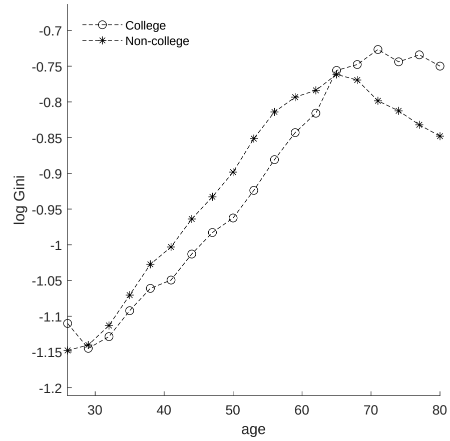

(b) Age profile: log Gini coefficient

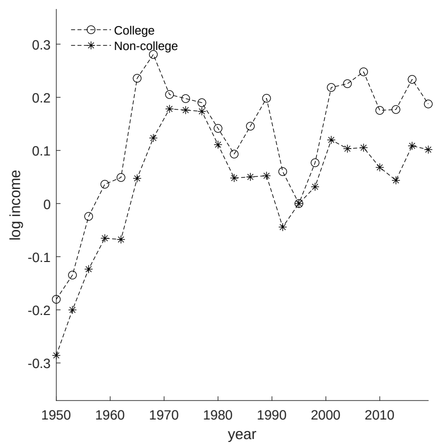

(c) Period profile: log mean income

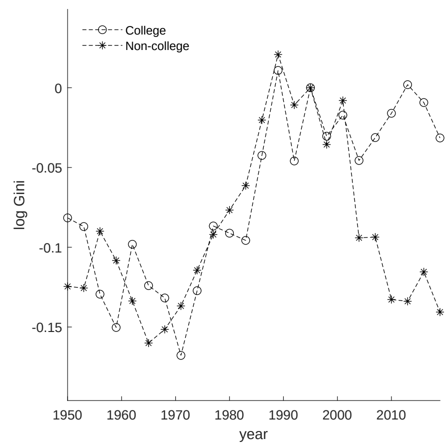

(d) Period profile: log Gini coefficient

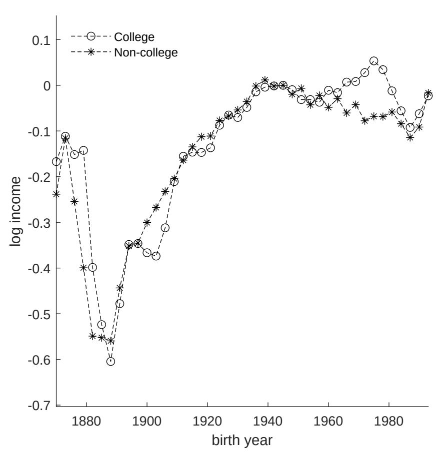

(e) Cohort profile: log mean income

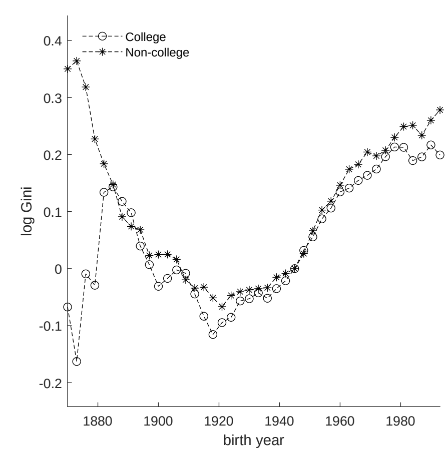

(f) Cohort profile: log Gini coefficient

Figure 5: Age, period, and cohort profiles of log mean income and log Gini coefficients in the SCF+ data. The period effect 1995 and cohort effect 1945 are normalized to zero. The period fixed effects for the years 1974, 1980, and 1986 are linearly interpolated.

error term for the log mean income and the log Gini coefficient.[19]

II.3 An aggregation methodology for the Gini coefficient

To study the evolution of income inequality at the population level, we need to aggregate the predicted subgroup means and Gini coefficients into a population-level Gini coefficient. Unfortunately, the Gini coefficient is not an aggregative inequality measure. That is, knowing the mean, Gini coefficient, and the population share of each subgroup is not sufficient to reconstruct the population-level Gini coefficient (Bourguignon, 1979). To overcome this issue, we propose a method to map the moments of subgroup distributions into the population-level Gini coefficient. The idea is to fit a parametric distribution for each set of subgroup moments and aggregate these distributions to generate a population-level income distribution.

We follow the principle of maximum entropy[20] by Jaynes (1957) and assume that income distributions at the subgroup level follow the parametric distribution that maximizes entropy subject to being supported on the positive real line and having given mean and Gini coefficient. This distribution was derived by Eliazar and Sokolov (2010), and we refer to it as the maximum entropy distribution (ME).[21]

The cumulative distribution function (CDF) of the ME distribution is given by

$$F^{\mathsf{ME}}(y;\sigma,\rho) = 1 - \frac{1}{\sigma\exp(\rho y) + (1-\sigma)} \quad \text{for } y \geq 0, \tag{10}$$

where the parameters σ and ρ are related to the mean income, μ, and the income Gini

[19]To convert predicted logarithms into levels, we take into account that the expected value is approximately given by $\exp\left(\lambda + \frac{\sigma^2}{2}\right)$, where λ is the predicted value of the logarithm and σ^2 is the variance of the expected value of the logarithm, which corresponds to the variance of the error term in the age-period-cohort model. This approximation is exact if the error term is normally distributed.

[20]The principle of maximum entropy states that "in making inferences on the basis of partial information we must use that probability distribution that has maximum entropy subject to whatever is known." Entropy is defined as $-\mathsf{E}\log(p(x_i))$, where $p(x_i)$ is the probability/density of outcome x_i. Intuitively, entropy is the expected value of uncertainty in a random variable's outcomes or the average level of information gained from observing the variable's outcomes. By using a maximum entropy distribution to model within-cohort income distributions, we make the least amount of additional assumptions on the shape of the income distribution after imposing the mean and the Gini coefficient.

[21]This is a slight abuse of language as the distribution we use here is a particular member of the class of maximum entropy distributions. In appendix H, we use lognormal and gamma as alternative distributions, which are also maximum entropy distributions but for different information constraints.

coefficient, g, as follows:

$$\mu = \frac{\log \sigma}{(\sigma - 1)\rho} \tag{11}$$

$$g = 1 + \frac{1}{\sigma - 1} - \frac{1}{\log \sigma}. \tag{12}$$

Since the expressions for μ and g are invertible for $\mu > 0$ and $0 < g < 1$, we can write the subgroup CDF as a function of the subgroup moments,

$$F_{a,p,c,e}(y) = F^{\mathsf{ME}}(y; \mu_{a,p,c,e}, g_{a,p,c,e}). \tag{13}$$

After fitting all income distributions at the subgroup level using the observed subgroup means and Gini coefficients, we construct the population-level income distribution as a weighted sum of the subgroup CDFs:

$$\Phi_p(y) = \sum_{a,e} s_{a,p,c,e} F_{a,p,c,e}(y), \tag{14}$$

where $s_{a,p,c,e}$ is population share of age-cohort-education group (a, c, e) in period p. The population-level Gini coefficient can then be computed as

$$G_p = 1 - \frac{1}{\mu_p} \int_0^{1} \big(1 - \Phi_p(y)\big)^2 dy, \tag{15}$$

where μ_p is the population-level mean income in period p.

To test our aggregation methodology, we compute the population-level income Gini coefficient for each survey year by applying our aggregation method to the observed subgroup moments and compare it with the population-level Gini coefficient computed directly from the data. Figure 6 depicts the results of this comparison. The solid black line shows the aggregated Gini coefficients, while the blue stars depict the population-level Gini coefficients computed directly from survey data. The aggregated Gini coefficients follow closely the path of the true Gini coefficients for both the CPS and SCF+ data. This observation lends confidence that our aggregation method is able to aggregate the Gini coefficients with only a limited loss of information.

To construct our main counterfactuals, we apply this aggregation methodology to the counterfactual subgroup moments and predicted population shares. Counterfactual subgroup moments are constructed as in equations (8) and (9). Similarly, we construct

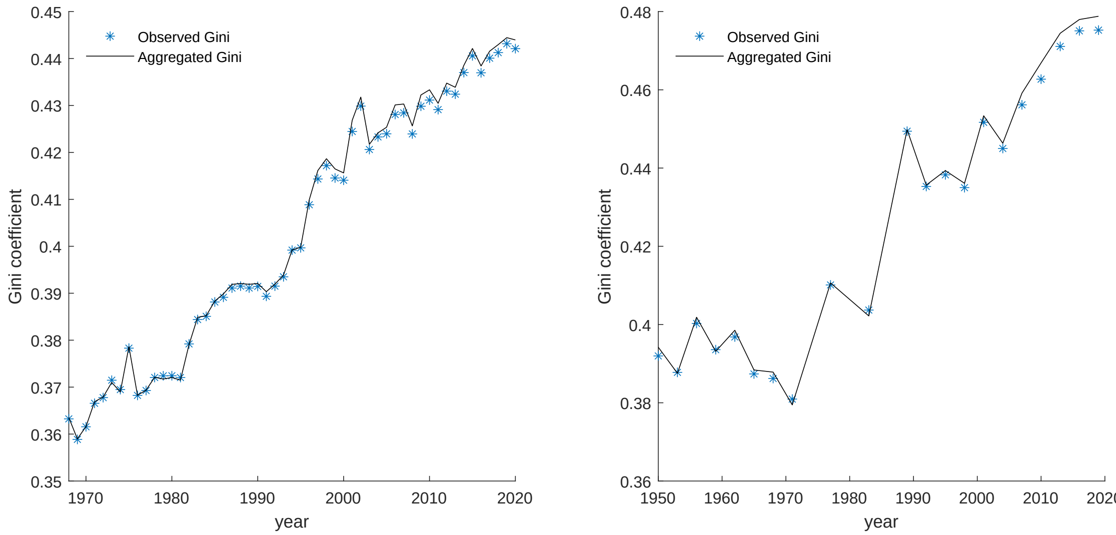

(a) CPS **(b)** SCF+

Figure 6: Aggregated Gini coefficients using ME distribution.

predicted population shares as

$$\tilde{s}_{a,p^0,c,e} = \begin{cases} \phi_{a,p^0}\,\psi_c & \text{if} \quad c < \bar{c}_0 \\ \phi_{a,p^0}\,\psi_{c_0} & \text{if} \quad c \qquad \bar{c}_0, \end{cases} \tag{16}$$

where $\phi_{a,p}$ denotes the population share of age group a in year p, which is either observed in the survey data or taken from the census forecasts; ψ_c is the college share of cohort c, which is assumed to be constant after age 26; and $\bar{c}_0$ is the youngest cohort present in the base year.

III Results

III.1 The role of demographic change in the past

Figure 7 shows how demographic change in the past has affected the evolution of income inequality. We show the results for both the CPS and the SCF+ data for purposes of comparison. The blue stars show the actual evolution of income inequality in the survey data. The solid black line shows the aggregated Gini coefficients, using the

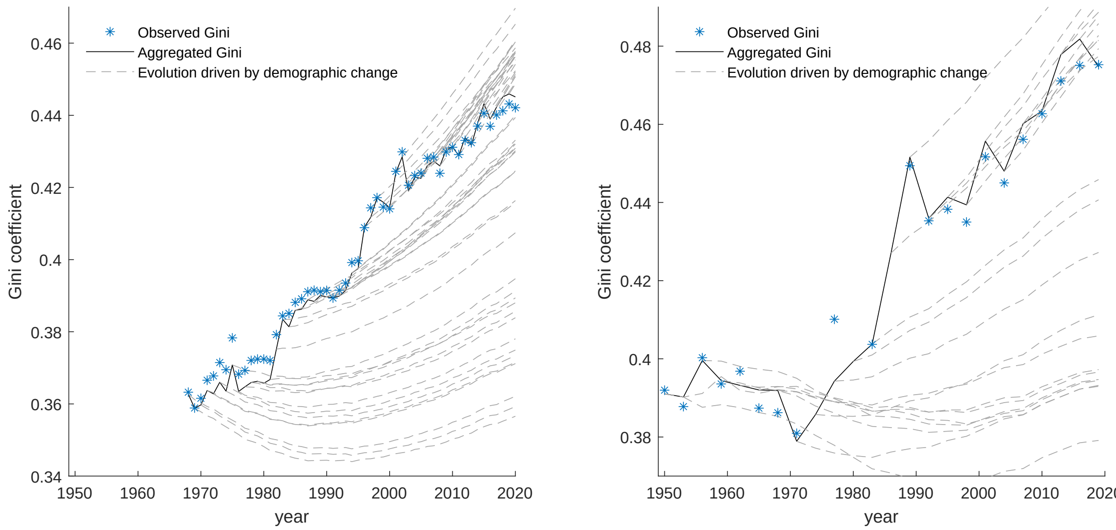

(a) CPS **(b)** SCF+

Figure 7: Counterfactual evolution of income inequality in the past. The blue stars show the observed evolution of income inequality. The dashed lines starting from different base years show the evolution of income inequality driven by demographic change.

predicted values from the age-period-cohort model.[22] Starting from each possible base year, a dashed gray line plots how income inequality would have evolved if demographic change unfolded as it actually did but the economic environment was held fixed in the base year.[23]

We find that demographic change had an important effect on the evolution of income inequality in the past. Moreover, it turns out that demographic change has become more important over time. In particular, we find that demographic change has little effect on income inequality if the economic environment is held fixed in the 1950s, 1960s, or 1970s. The counterfactual trajectories of income inequality for these base years show a slight decrease and then recover back to the initial level. However, the slopes of the counterfactual trajectories are steeper for more recent base years. For example, if

[22]In contrast to figure 6, differences between our aggregated time series and the observed population-level Gini coefficients now stem from two sources. First, as in figure 6, using parametric income distributions introduces error if incomes at the subgroup level do not exactly follow the assumed parametric distribution. Second, using predicted values from our age-period-cohort models to fit the parametric distributions introduces additional error if the estimated model does not explain all the variation in the data. Overall, we match the shape of the time series well, and the differences between the aggregated and the observed population-level Gini coefficients are small.

[23]To compute the counterfactual evolution starting from base years 1974, 1980, and 1986, for which we do not have survey waves in the SCF+ data, we linearly interpolate the period fixed effects and the population share of each age-by-education group.

the economic environment is held fixed after the mid-1990s, then our counterfactuals not only show an increase in the income Gini coefficient, but demographic change can actually account for the entire observed increase in income inequality. These results are consistent across both data sets.

How much of the actual increase in the income Gini coefficient can be accounted for by demographic change depends on the normalization of the linear age, period, and cohort trends. A stronger positive cohort trend, which also implies a stronger positive trend in age effects, increases the effect of both cohort replacement and population aging on income inequality. However, as we show in appendix D, a large share of the observed increase in income inequality since the 1990s can be accounted for by demographic change even if we assume no trend in the cohort profile and we instead allow the period effects to exhibit a strong positive trend. A potential concern is that even a "no cohort trend" normalization is not conservative enough and that the true linear trend in cohort effects is negative. To address this concern, we assess the plausibility of a positive cohort trend in section III.4 using additional data for selected birth cohorts from the NLSY surveys. We show that these data support the estimated cohort effects under our baseline normalization.

In figure 7, we use age, period, and cohort effects that are estimated on the full sample. As an additional exercise, we compute vintage predictions in appendix E, in which for each base year, we only use data up until that year. We obtain similar results, especially for the SCF+ data for which the vintage predictions are almost identical to the past counterfactuals depicted in figure 7.

In our main analysis we measure aggregate inequality using the Gini coefficient. However, the Gini coefficient does not capture all changes in the distribution of household income. It is therefore important to know whether our main findings also apply to other inequality measures. In appendix F, we study the effect of demographic change on the evolution of the share of income received by households in the top five percentiles and find similar results.

III.2 The role of demographic change in the future

The important role of past demographic change, especially in the most recent decades, raises the question whether demographic change will further increase income inequality in the future. To address this question, we choose the latest survey wave as the base year and plot the evolution of income inequality under predicted demographic change until the year 2060 in figure 8. We plot the evolution of the Gini coefficient in the future

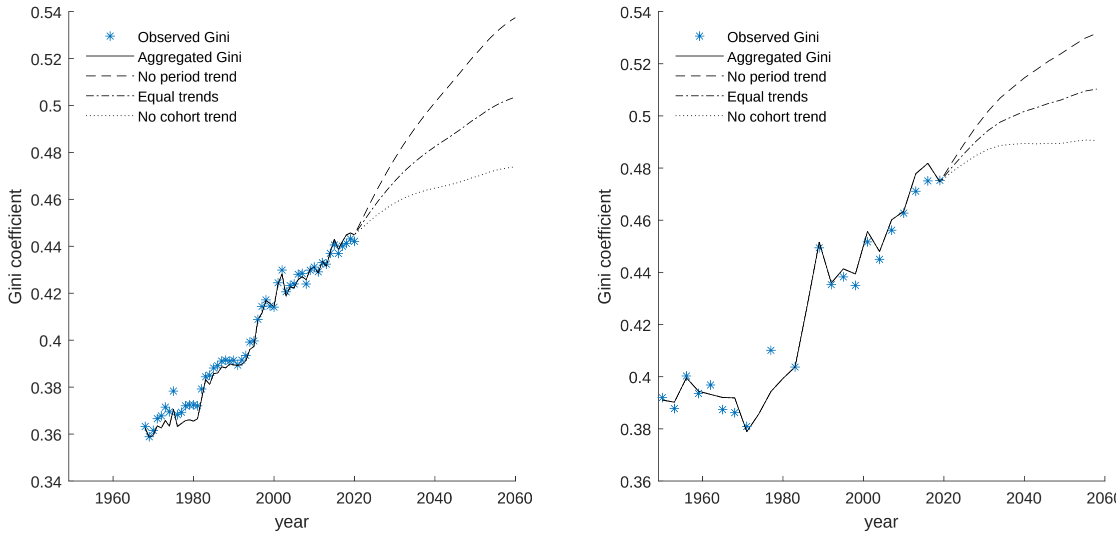

(a) CPS

(b) SCF+

Figure 8: Counterfactual evolution of income inequality in the future.

under three different normalizations, depicted as dashed lines.

We find that demographic change will lead to an increase in income inequality over the next four decades. This is the case in both the CPS as well as the SCF+ data, irrespective of whether we attribute trends to cohort effects or period effects; however, the increase is more dramatic in the former case. The difference between the three specifications is due to two interlinked factors. First, assuming no period trend in income Gini coefficient in the past implies a stronger positive cohort trend. Thus, cohort replacement will put a stronger upward pressure on overall income inequality. Second, the estimated age profile of within-cohort income Gini coefficient is steeper if we assume no period trends, which implies that projected population aging has a greater effect on income inequality.

III.3 Decomposing the effects of demographic change

We have shown above that demographic change has mattered for the evolution of income inequality in the recent past and will likely further increase income inequality in the future. These results, however, do not reveal whether projected demographic change will increase income inequality predominantly through the effect of population aging or that of cohort replacement. To investigate the respective contributions of these two channels, we construct additional counterfactuals by shutting down either the population aging or the cohort replacement channel. These counterfactuals are computed under the baseline normalization where we assume equal trends in period and cohort effects in both log mean income and log Gini coefficient.

To isolate the effect of cohort replacement, we fix the marginal age distribution of the population. In particular, for each target year p^0, we construct the counterfactual population-level CDF as

$$\tilde{\Phi}_{p^0}(y) = \sum_{a,e} \tilde{s}_{a,p^0,c,e} \tilde{F}_{a,p^0,c,e}(y), \tag{17}$$

where counterfactual income distributions, $\tilde{F}_{a,p^0,c,e}(y)$, are constructed as before, and population shares, $\tilde{s}_{a,p^0,c,e}$, are constructed as

$$\tilde{s}_{a,p^0,c,e} = \begin{cases} \phi_{a,p}\, \psi_c & \text{if } \; c < \bar{c}_0 \\ \phi_{a,p}\, \psi_{c_0} & \text{if } \; c \quad \bar{c}_0, \end{cases} \tag{18}$$

where $\phi_{a,p}$ is the population share of age group a in the base year $\bar{p}$, and ψ_c is the college share of cohort c.

To isolate the effect of population aging, we remove all cohort differences and allow only population shares of different age groups to change. In particular, we first equalize college shares across birth cohorts by setting the college share in each cohort equal to the aggregate college share in the base year. We then equalize cohort effects by setting cohort effects equal to a common cohort effect, $\bar{\kappa}_{p,e}$, in each education group e. The common cohort effects are chosen such that the predicted aggregate Gini coefficient in the base year remains unchanged. We thus set means and Gini coefficients as

$$\tilde{\mu}_{a,p^0,e} = \exp\left(\theta^\mu_e + \alpha^\mu_e a + \pi^\mu_e \bar{p} + \check{\alpha}^\mu_{a,e} + \check{\pi}^\mu_{p,e} + \bar{\kappa}_{p,e} + \frac{\sigma^2_{e,\mu}}{2}\right) \tag{19}$$

$$\tilde{g}_{a,p^0,e} = \exp\left(\theta^g_e + \alpha^g_e a + \pi^g_e \bar{p} + \check{\alpha}^g_{a,e} + \check{\pi}^g_{p,e} + \bar{\kappa}_{p,e} + \frac{\sigma^2_{e,g}}{2}\right) \tag{20}$$

and form subgroup CDFs by plugging these moments into equation (13). Finally, we construct the population-level CDF as

$$\tilde{\Phi}_{p^0}(y) = \sum_{a,e} \tilde{s}_{a,p^0,e} \tilde{F}_{a,p^0,e}(y), \tag{21}$$

where $\tilde{s}_{a,p^0,e}$ is constructed as

$$\tilde{s}_{a,p^0,e} = \phi_{a,p}\, \psi_{p}, \tag{22}$$

where $\phi_{a,p}$ is the population share of age group a in year p and ψ_p is the college share in the population in the base year $\bar{p}$.

Figure 9 plots the future evolution of income inequality under these counterfactuals. The dashed line shows the full effect of demographic change under the baseline normalization. The two dotted lines show the isolated effects of population aging and cohort replacement. The main observation from this figure is that cohort replacement is driving most of the increase in income inequality in the future, while population aging has a small positive effect.

The finding that population aging will hardly affect income inequality in the future raises the question of whether this lack of effect was also the case in the past. In figure 10, we compare the increase in income inequality over different 21-year periods in our baseline counterfactual to the obtained increase if we shut down either the cohort replacement or the population aging channel.

We find that population aging increased income inequality in the 1950s and in the time period after 1980. The effect of population aging peaked in the 21-year period starting around the year 2000, when it explained about one-third of the full effect of demographic change.[24] We can also see that the effect of population aging has returned to almost zero by the time period starting with the latest survey wave, which is consistent with the findings in figure 9. Overall, cohort replacement accounts for most of the increase in income inequality driven by demographic change.

III.4 Normalizing the cohort trends using NLSY data

In our main analyses, we estimate cohort effects in income inequality using an age-period-cohort model. We interpret the estimated cohort effects as summary statistics for all cohort-level differences in income-relevant characteristics. In the absence of detailed data on the changing characteristics of birth cohorts throughout the 20th century, our approach offers a viable strategy for studying how population aging and cohort replacement are affecting the evolution of income inequality. However, the results

[24] As the aggregation of age, period, and cohort effects into a population-level Gini coefficient is highly nonlinear, the individual effects of population aging and cohort replacement need not sum up exactly to the full effect of demographic change in our counterfactual. In practice, however, we find that the sum of the effects of the individual channels is not too far from the full effect of demographic change.

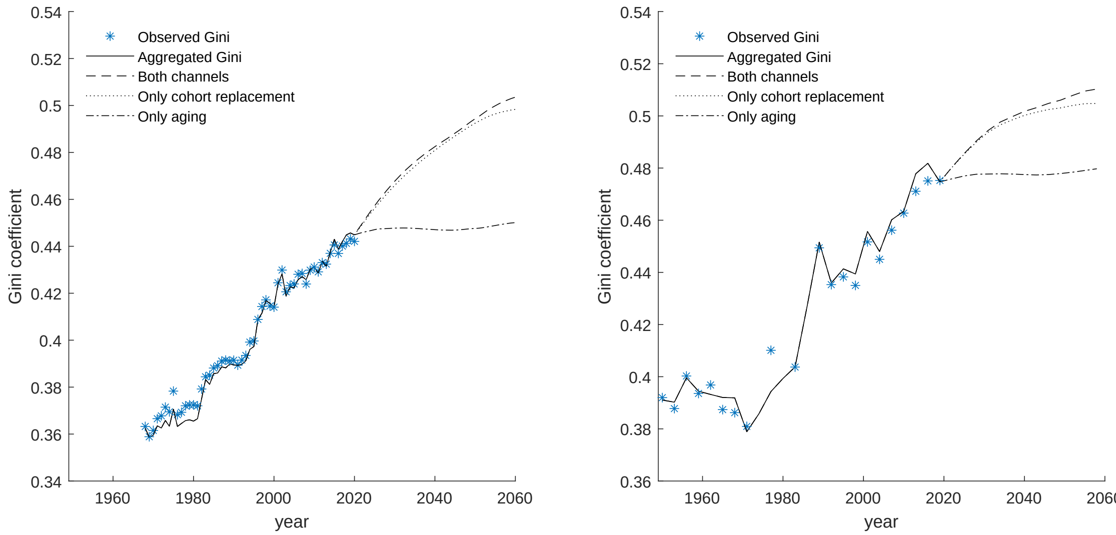

(a) CPS

(b) SCF+

Figure 9: Decomposition of the effect of demographic change in the future.

rely on a normalization of the linear trends in the age-, period-, and cohort profiles, and while we have argued for the reasonableness of the chosen normalization, we have so far relied on theoretical rather than empirical arguments.

In this section, we use rich data on the birth cohorts 1957-1961 and 1980-1984 from the National Longitudinal Surveys of Youth (NLSY) to confirm that recent birth cohorts are more unequal and to derive a normalization for the cohort trend that is based on observed cohort-level differences. Using the methodology developed in DiNardo et al. (1996), we study how inequality among the NLSY79 cohorts (those born between 1957-1961) changes if we impose on them the distribution of characteristics of the NLSY97 cohorts (those born between 1980-1984). Specifically, we adjust the distributions of birthplace (US versus abroad), parental education, marital status, spousal education, as well as education and race, both interacted with cognitive test scores. Note that adjusting the joint distribution of education and cognitive test scores accounts for trends in selection into education. We run this analysis separately for college-educated and non-college-educated respondents to conform with our age-period-cohort-model.

Table 2 reports the results. We find an increase of 0.08 and 0.12 in the log Gini coefficient for college-educated and non-college-educated households, respectively. Since these differences in the Gini coefficient reflect cohort-level differences in characteristics that should be stable after early adulthood, they correspond to the differences in cohort effects between the two cohorts. For comparison, the age-period-cohort model under the

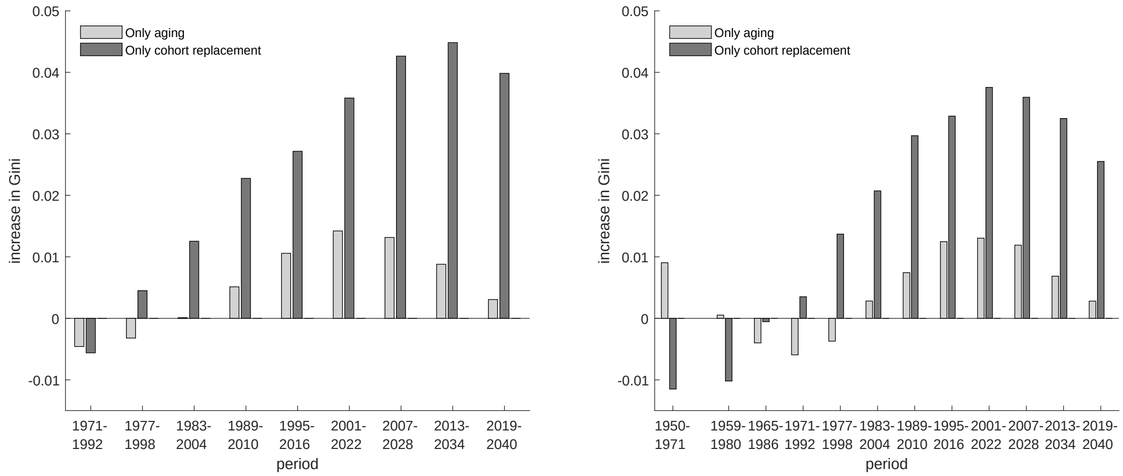

(a) CPS **(b)** SCF+

Figure 10: Decomposition of the effect of demographic change in the past and future.

baseline normalization estimates these differences to be 0.12 and 0.14 in the CPS data, and 0.08 and 0.10 in the SCF+ data. Thus, the results from this analysis corroborate our finding that more recent birth cohorts have distributions of characteristics that induce higher inequality in household income compared to earlier cohorts. Moreover, the estimated cohort differences using detailed NLSY data are also quantitatively similar to what we estimate with the age-period-cohort model under our preferred normalization.

In appendix G, we use the results derived on the NLSY data to derive a normalization for the linear trends in the age-period-cohort model. We derive the effects of demographic change under this new normalization and show that they are very similar to our main results.

	college	non-college
NLSY	0.08	0.12
APC-model (CPS, baseline)	0.12	0.14
APC-model (SCF+, baseline)	0.08	0.10

Table 2: Difference in the cohort effects of cohorts 1959 and 1981 in the NLSY data and in the estimated age-period-cohort model.

IV A re-weighting analysis

Our findings suggest a more important role for demographic change in explaining the evolution of household income inequality than generally found in the literature (e.g., Kuhn et al. (2020)). Previous studies of compositional effects typically rely on using a re-weighting method following DiNardo et al. (1996), and explicitly account for differences in only a limited number of characteristics. For comparison, we re-implement our thought experiment using a re-weighting method. In this exercise, we account for changes in terms of the age structure and educational attainment of the US population. These are the most common characteristics used in studies of compositional effects on inequality (e.g., see Lemieux (2006), Kuhn et al. (2020), Hoffmann et al. (2020)), and we observe these characteristics in both surveys. Moreover, educational attainment is fairly constant after age 26 which simplifies the implementation of the thought experiment.

Let $F_{YX,p}$ be the CDF for the joint distribution of income Y and characteristics X in year p, and let $F_{X,p}$ be the CDF for the marginal distribution of characteristics X in year p. Given a pair of base year $\bar{p}$ and target year p' with $p' > \bar{p}$, our goal is to compute a counterfactual income distribution for the target year, $\tilde{F}_{YX,p,p'}$, by fixing the conditional distribution of income in the base year, $F_{Y|X,p}$, and implementing a distribution of characteristics, $\tilde{F}_{X,p'}$, as implied by our thought experiment. This counterfactual income distribution can be obtained under the re-weighting approach by suitably re-weighting the cross-sectional data in the base year. In our implementation, the characteristics X include age and a dummy for college education.

We construct two sets of re-weighting factors. The first set of factors re-weight the data in the base year to match the marginal distribution of age in the target year. The second set of re-weighting factors adjusts the share of college-educated household heads in each age group to account for the fact that any given age corresponds to a more recent birth cohort in the target year compared with the base year. For birth cohorts older than 26 in the base year, we assume that their college share remains fixed. For birth cohorts that turn 26 only after the base year, we assume that they have the same college share as the youngest cohort in the base year.

For each pair of base and target year $(\bar{p}, p')$, the first set of age-specific re-weighting factors, $\phi_{p,p'}(a)$, is given by the ratio of the marginal density of age in the target year to the density in the base year:

$$\phi_{p,p'}(a) = \frac{dF_{A,p'}(a)}{dF_{A,p}(a)}. \tag{23}$$

Using the fact that birth cohort equals year minus age, the second set of age and education-specific re-weighting factors, $\psi_{p,p'}(a,e)$, can be written as

$$\psi_{p,p'}(a,e) = \begin{cases} \frac{dF_{E|A,\bar{p}}(e|a-(p'-\bar{p}))}{dF_{E|A,\bar{p}}(e|a)} & \text{for } a \geq a_0 + p' - \bar{p} \\ \frac{dF_{E|A,\bar{p}}(e|a_0)}{dF_{E|A,\bar{p}}(e|a)} & \text{for } a < a_0 + p' - \bar{p} \end{cases} \tag{24}$$

where $F_{E|A,\bar{p}}(e|a)$ denotes the conditional distribution of education, E, given age, A, in the base year, $\bar{p}$, and a_0 is set to 26. Since age and education take discrete values in our data sets, we can use relative frequencies as estimators for densities.[25]

We can study the compositional effects of demographic change by multiplying the sample weight of each observation in the base year by the product of the two re-weighting factors, which results in the desired counterfactual distribution of income for the target year. In particular, we can compute the counterfactual Gini coefficient $\tilde{G}_{p,p'}$ as

$$\tilde{G}_{p,p'} = \frac{\sum_i^{N^{\bar{p}}} \sum_{j=i+1}^{N^{\bar{p}}} w_i w_j |y_i - y_j|}{\sum_i^{N^{\bar{p}}} w_i \sum_i^{N^{\bar{p}}} w_i y_i}, \tag{25}$$

where $(y_i)_{i=1}^{N^{\bar{p}}}$ is the vector of income observations from the base year $\bar{p}$, and $w_i = \omega_i \phi_{p,p'}(a_i) \psi_{p,p'}(a_i, e_i)$, where ω_i is the sample weight for observation i in the base year.

Because we only change the weights and not the underlying observations in the base year, the income distributions within population subgroups are held fixed and the overall income Gini coefficient can only change as a result of changing shares of different subgroups. Hence, as long as cohort differences in income distributions are captured by age and education of the household head, this exercise corresponds to the thought experiment of fixing the economic environment in its state in the base year and only allowing the compositional effects of demographic change to shape income inequality in the subsequent years.

In figure 11, we plot the actual evolution of the income Gini coefficient together with counterfactual evolutions for different base years. The blue stars again show the observed evolution of income inequality in the survey data. The dashed lines starting from different base years show the evolution of income inequality driven by demographic change. In contrast to our main results, the re-weighting analysis suggests that demographic change has contributed little to the observed increase in income inequality.

[25] DiNardo et al. (1996) suggest estimating densities using a regression analysis and Bayes theorem. Their approach is equivalent to ours when densities are estimated using a fully saturated model. Estimating a fully saturated model is feasible in our context because the common support assumption holds.

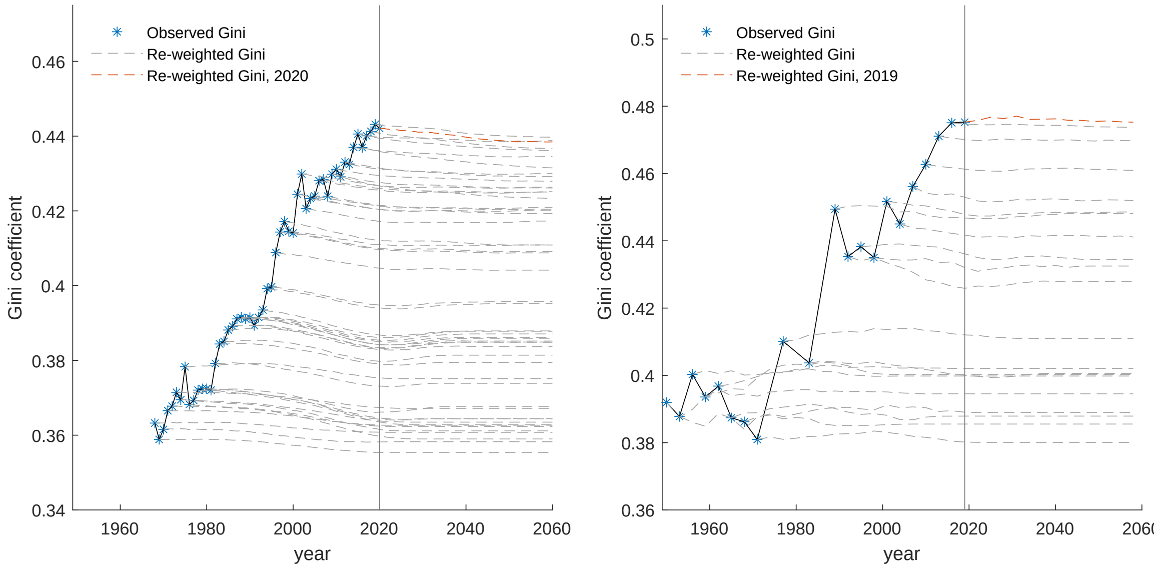

(a) CPS **(b)** SCF+

Figure 11: The effect of demographic change as suggested by the re-weighting analysis. The blue stars show the observed evolution of income inequality. The dashed lines starting from different base years show the evolution of income inequality driven by demographic change. The dashed red lines highlight the counterfactuals following the most recent survey wave in each data set.

IV.1 Shortcomings of the re-weighting analysis

The re-weighting method gets the effect of demographic change right only if the conditional distribution of unobserved characteristics does not change over time. This is typically called the ignorability or conditional independence assumption (Fortin et al., 2011). In our case, this means that all characteristics that affect a cohort's income distribution must be summarized by that cohort's share of college-educated household heads. This is unlikely to be true. To the extent that subgroup income distributions evolve according to the age-period-cohort model in section II.1, the estimated cohort profiles clearly show that differences in age and educational attainment alone cannot satisfactorily explain cohort-level differences in income distributions.

To complicate matters, population aging requires the researcher to increase the weights on older households in order to match the evolution of the population age structure. However, increasing the weight on older households is problematic because older individuals in any given cross-section belong to earlier birth cohorts. Any attempt to match the evolution of the age structure therefore implies up-weighting households whose unobserved characteristics are furthest from those of the typical household in the target year. Hence, the re-weighting analysis fails to capture the effect of cohort replacement and confounds the effect population aging.[26]

V Conclusion

In this paper, we study how demographic change affects the evolution of household income inequality in the United States. We consider a thought experiment in which the economic environment is held fixed but demographic change is allowed to take place. In this thought experiment, demographic change affects inequality not only because population aging increases the share of older households, but also because older birth cohorts are gradually replaced by younger birth cohorts that have different distributions of income-relevant characteristics. Moreover, we use the thought experiment to study how projected demographic change will affect inequality in the near future.

The main contributions of this paper are to highlight the importance of cohort differences and to implement a parametric methodology that can account for both population aging and cohort replacement. We model the evolution of subgroup income distributions by an additive age-period-cohort model, which allows us to account for cohort differences in both observed and unobserved characteristics. Using household income data for the United States, we estimate life-cycle profiles and cohort differences in mean incomes and income Gini coefficients. We document important cohort differences in income distributions that are not accounted for by differences in age and educational attainment. We then use the estimated age and cohort effects to predict how the moments of subgroup income distributions evolve under demographic change when the economic environment is held fixed. Finally, we use these counterfactual moments, together with predicted population shares, to study the effect of demographic change on aggregate inequality.

We find that demographic change plays an important role in the evolution of household income inequality in the United States—both in the past and in the future. We argue that the compositional effects of demographic change can account for all of the increase in income inequality over the past two decades. Moreover, we predict that demographic change will further increase inequality in the near future, with our estimates suggesting an increase in the income Gini coefficient of between one and six percentage

[26] Interestingly, the literature on adjusting Gini coefficients for differences in age structures following Paglin (1975) often ignores that income distributions can depend on birth cohorts, and therefore, it derives age-adjusted Gini coefficients that are also confounded by cohort effects.

points by the year 2040.

To derive these results, we impose a number of restrictions on the data. First, we assume additively separable age, period, and cohort profiles for the logarithms of mean income and the income Gini coefficient at the level of the demographic subgroup. In the main text, we discuss a simple income process that gives rise to additive age, period, and cohort effects for these moments. Second, we restrict linear trends in period and cohort effects to be weakly positive. We show that detailed information on income-relevant characteristics of selected cohorts is consistent with this assumption. Finally, our thought experiment abstracts from general equilibrium effects and is therefore a *simple counterfactual treatment* (Fortin et al., 2011).

An important insight from this paper is that changes in aggregate inequality are not always indicative of contemporaneous changes in the economic environment. Instead, changes in aggregate inequality can result from the gradual replacement of older cohorts whose characteristics were shaped several decades earlier. In the case of the United States, we argue that cohorts born in the second half of the 20th century have become progressively more unequal in their income-relevant characteristics, which in turn affects the evolution of aggregate income inequality in the first half of the 21st century.

References

Acemoglu, Daron and David Autor, "Skills, tasks and technologies: Implications for employment and earnings," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

Aguiar, Mark and Erik Hurst, "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 2013, *121* (3), 437–492.

Altonji, Joseph G, Prashant Bharadwaj, and Fabian Lange, "Changes in the characteristics of American youth: Implications for adult outcomes," *Journal of Labor Economics*, 2012, *30* (4), 783–828.

Attanasio, Orazio P., "Cohort Analysis of Saving Behavior by U.S. Households," *Journal of Human Resources*, 1998, *33* (3), 575–609.

Auclert, Adrien, Hannes Malmberg, Frederic Marténét, and Matthew Rognlie, "Demographics, Wealth, and Global Imbalances in the Twenty-First Century," *Working paper*, 2021.

Autor, David H, Lawrence F Katz, and Melissa S Kearney, "Trends in US wage inequality: Revising the revisionists," *The Review of Economics and Statistics*, 2008, *90* (2), 300–323.

Bourguignon, Francois, "Decomposable Income Inequality Measures," *Econometrica*, 1979, *47* (4), 901–920.

Bowlus, Audra J and Chris Robinson, "Human capital prices, productivity, and growth," *American Economic Review*, 2012, *102* (7), 3483–3515.

Burtless, Gary, "Effects of Growing Wage Disparities and Changing Family Composition on the US Income Distribution," *European Economic Review*, 1999, *43* (4-6), 853–865.

Card, David and Thomas Lemieux, "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis," *Quarterly Journal of Economics*, 2001, *116* (2), 705–746.

Daly, Mary C and Robert G Valletta, "Inequality and Poverty in United States: the Effects of Rising Dispersion of Men's Earnings and Changing Family Behaviour," *Economica*, 2006, *73* (289), 75–98.

Deaton, Angus and Christina Paxson, "Intertemporal Choice and Inequality," *Journal of Political Economy*, 1994, *102* (3), 437–467.

Deaton, Angus S and Christina Paxson, "Saving, growth, and aging in Taiwan," in "Studies in the Economics of Aging," University of Chicago Press, 1994, pp. 331–362.

DiNardo, John, Nicole M Fortin, and Thomas Lemieux, "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semi-Parametric Approach," *Econometrica*, 1996, *64*, 1001–1044.

Eika, Lasse, Magne Mogstad, and Basit Zafar, "Educational Assortative Mating and Household Income Inequality," *Journal of Political Economy*, 2019, *127* (6), 2795–2835.

Eliazar, Iddo and Igor M Sokolov, "Maximization of Statistical Heterogeneity: From Shannon's Entropy to Gini's index," *Physica A: Statistical Mechanics and its Applications*, 2010, *389* (16), 3023–3038.

Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux, "Unconditional quantile regressions," *Econometrica*, 2009, *77* (3), 953–973.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo, "Decomposition Methods in Economics," in "Handbook of Labor Economics," Vol. 4, Elsevier, 2011, pp. 1–102.

Gourinchas, Pierre-Olivier and Jonathan A Parker, "Consumption over the life cycle," *Econometrica*, 2002, *70* (1), 47–89.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos, "Marry Your Like: Assortative Mating and Income Inequality," *American Economic Review*, 2014, *104* (5), 348–53.

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner, "Lifetime earnings in the united states over six decades," *American Economic Journal: Applied Economics*, 2022, *14* (4), 446–479.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante, "Two Views of Inequality over the Life Cycle," *Journal of the European Economic Association*, 2005, *3* (2–3), 765–775.

_ , _ , and _ , "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, 2014, *104* (7), 2075–2126.

Heckman, James J, Lance Lochner, and Christopher Taber, "Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents," *Review of economic dynamics*, 1998, *1* (1), 1–58.

Heikkuri, Vesa-Matti and Matthias Schief, "Subgroup Decomposition of the Gini Coefficient: A New Solution to an Old Problem," *working paper*, 2022.

Hendricks, Lutz and Todd Schoellman, "Student Abilities during the Expansion of US Education," *Journal of Monetary Economics*, April 2014, *63*, 19–36.

Hobijn, Bart, Todd Schoellman, and Alberto Vindas, "Structural transformation by cohort," Technical Report, Technical Report, Arizona State University 2019.

Hoffmann, Florian, David S Lee, and Thomas Lemieux, "Growing income inequality in the United States and other advanced economies," *Journal of Economic Perspectives*, 2020, *34* (4), 52–78.

Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow, "The Allocation of Talent and US Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Huggett, Mark, Gustavo Ventura, and Amir Yaron, "Sources of Lifetime Inequality," *American Economic Review*, 2011, *101* (7), 2923–54.

Jaynes, Edwin T, "Information Theory and Statistical Mechanics," *Physical Review*, 1957, *106* (4), 620.

Jedwab, Remi, Paul Romer, Asif M Islam, and Roberto Samaniego, "Human capital accumulation at work: Estimates for the world and implications for development," *American Economic Journal: Macroeconomics*, 2023, *15* (3), 191–223.

Jones, Charles I, "Sources of US economic growth in a world of ideas," *American Economic Review*, 2002, *92* (1), 220–239.

_ , "The past and future of economic growth: A semi-endogenous perspective," Technical Report, National Bureau of Economic Research 2021.

Juhn, Chinhui, Kevin M Murphy, and Brooks Pierce, "Wage inequality and the rise in returns to skill," *Journal of Political Economy*, 1993, *101* (3), 410–442.

Kahn, Lisa B, "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy," *Labour Economics*, 2010, *17* (2), 303–316.

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins, "Income and Wealth Inequality in America, 1949–2016," *Journal of Political Economy*, 2020, *128* (9), 3469–3519.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, "Life Cycle Wage Growth across Countries," *Journal of Political Economy*, 2018, *126* (2), 797–849.

Larrimore, Jeff, "Accounting for United States Household Income Inequality Trends: The Changing Importance of Household Structure and Male and Female Labor Earnings Inequality," *Review of Income and Wealth*, 2014, *60* (4), 683–701.

Lee, Donghoon and Kenneth I Wolpin, "Intersectoral labor mobility and the growth of the service sector," *Econometrica*, 2006, *74* (1), 1–46.

Lemieux, Thomas, "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?," *American Economic Review*, 2006, *96* (3), 461–498.

Low, Hamish, Costas Meghir, and Luigi Pistaferri, "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 2010, *100* (4), 1432–1467.

McKenzie, David J, "Disentangling age, cohort and time effects in the additive model," *Oxford bulletin of economics and statistics*, 2006, *68* (4), 473–495.

Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz, "The Short- and Long-Term Career Effects of Graduating in a Recession," *American Economic Journal: Applied Economics*, 2012, *4* (1), 1–29.

Paglin, Morton, "The Measurement and Trend of Inequality: A Basic Revision," *American Economic Review*, 1975, *65* (4), 598–609.

Porzio, Tommaso, Federico Rossi, and Gabriella Santangelo, "The human side of structural transformation," *American Economic Review*, 2022, *112* (8), 2774–2814.

Preda, Vasile, Silvia Dedu, and Carmen Gheorghe, "New classes of Lorenz curves by maximizing Tsallis entropy under mean and Gini equality and inequality constraints," *Physica A: Statistical Mechanics and its Applications*, 2015, *436*, 925–932.

Raaum, Oddbjørn and Knut Røed, "Do Business Cycle Conditions at the Time of Labor Market Entry Affect Future Employment Prospects?," *Review of Economics and Statistics*, 2006, *88* (2), 193–210.

Rothstein, Jesse, "The Lost Generation? Scarring after the Great Recession," *Working paper*, 2019.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek, "IPUMS USA: Version 10.0," *Minneapolis, MN: IPUMS*, 2020.

Schwandt, Hannes and Till Von Wachter, "Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets," *Journal of Labor Economics*, 2019, *37* (S1), 161–198.

Storesletten, Kjetil, Christopher I Telmer, and Amir Yaron, "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 2004, *51* (3), 609–633.

United States Census Bureau, "2017 National Population Projections Tables: Main Series," 2017. Data retrieved from the United States Census Bureau, `https://www.census.gov/data/tables/2017/demo/popproj/2017-summary-tables.html`.

Appendix

A Age, period, and cohort profiles under different normalizations

Figures 12 to 15 show the estimated age, period, cohort profiles under alternative normalizations suggested by Lagakos et al. (2018). The green lines show the profiles under no trend in period effects normalization, the red line shows the profiles under no cohort trend normalization and the blue line shows the profiles under the intermediate case of assuming equal trends in period and cohort profiles. Hence, the profiles plotted in blue correspond to those shown in section II.1.4.

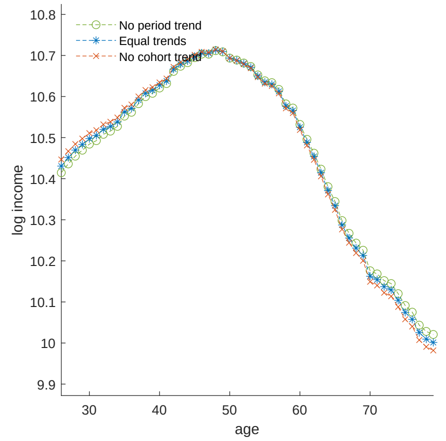

(a) Age profile: log mean income

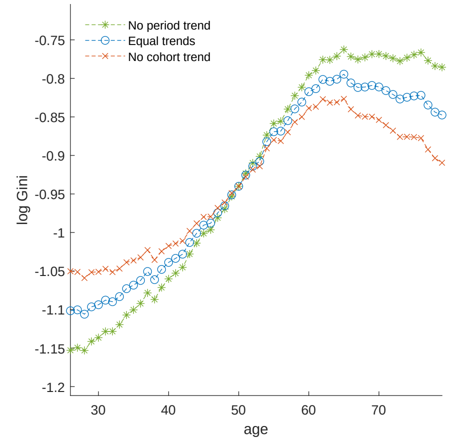

(b) Age profile: log Gini coefficient

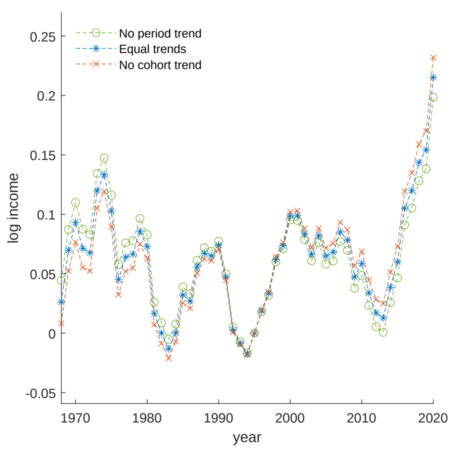

(c) Period profile: log mean income

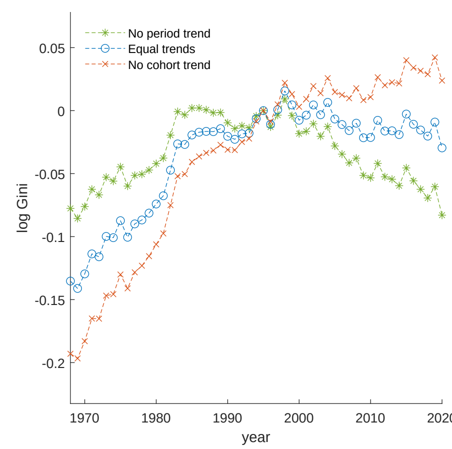

(d) Period profile: log Gini coefficient

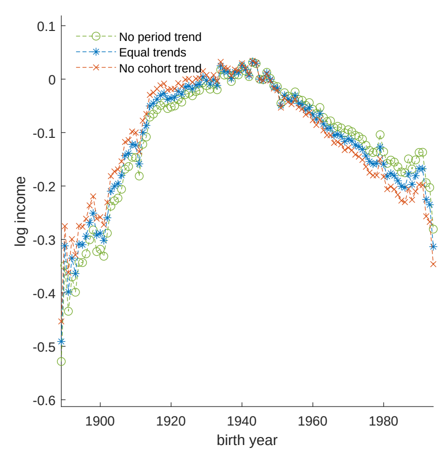

(e) Cohort profile: log mean income

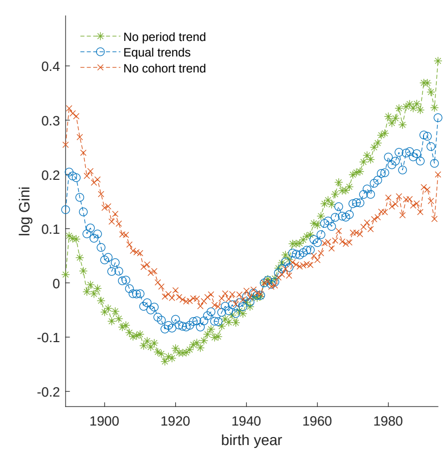

(f) Cohort profile: log Gini coefficient

Figure 12: Age-, period-, and cohort profiles of log mean income and income Gini coefficients for households with non-college educated household head in the CPS data under different normalizations for the linear trends. The period effect 1995 and cohort effect 1945 are normalized to zero.

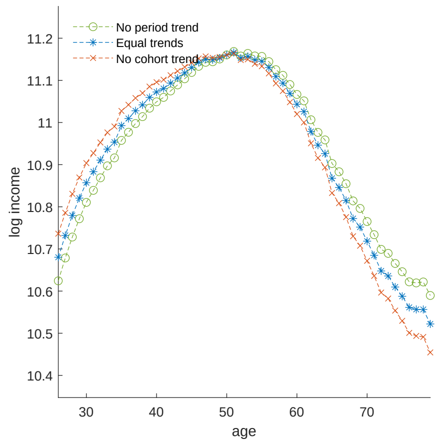

(a) Age profile: log mean income

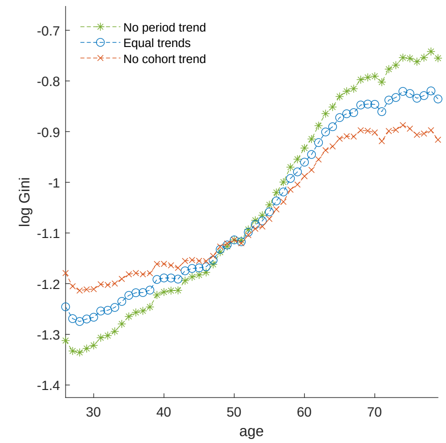

(b) Age profile: log Gini coefficient

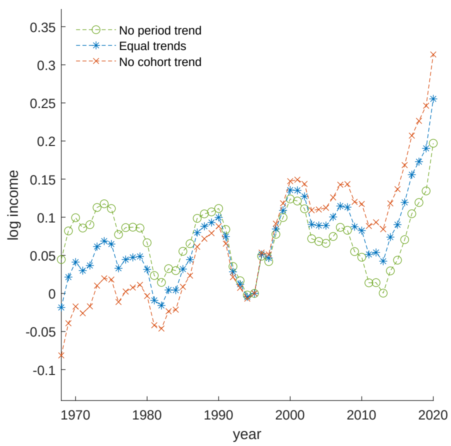

(c) Period profile: log mean income

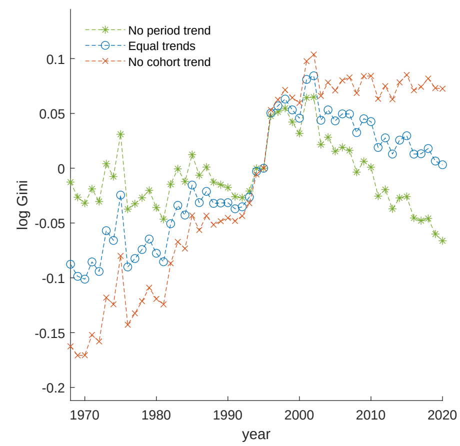

(d) Period profile: log Gini coefficient

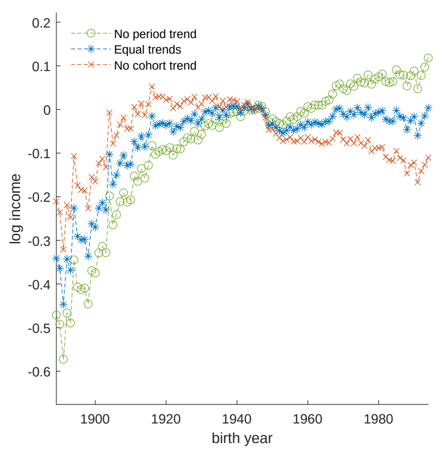

(e) Cohort profile: log mean income

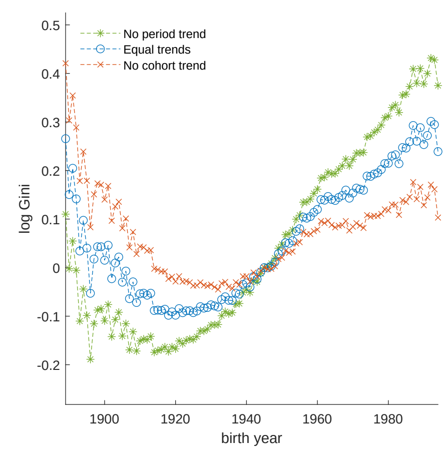

(f) Cohort profile: log Gini coefficient

Figure 13: Age-, period-, and cohort profiles of log mean income and log Gini coefficients for households with college educated household head in the CPS data under different normalizations for the linear trends. The period effect 1995 and cohort effect 1945 are normalized to zero.

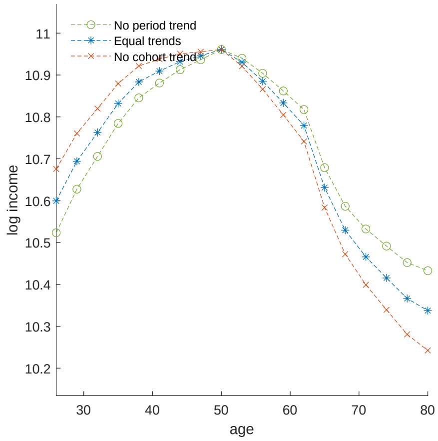

(a) Age profile: log mean income

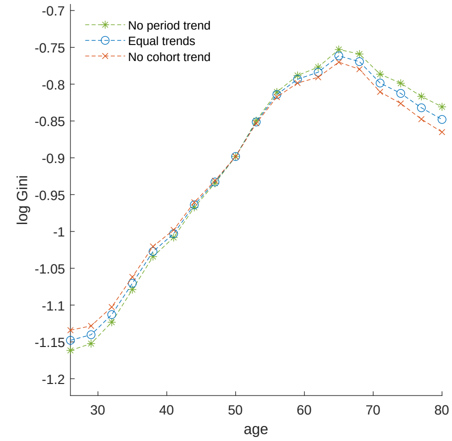

(b) Age profile: log Gini coefficient

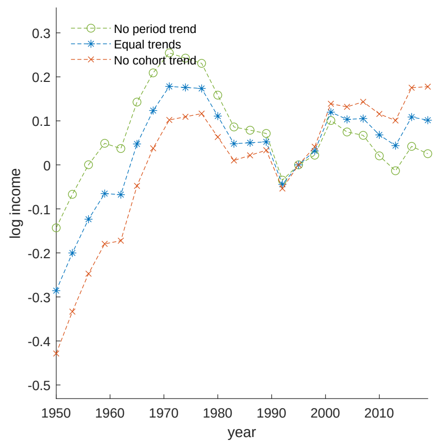

(c) Period profile: log mean income

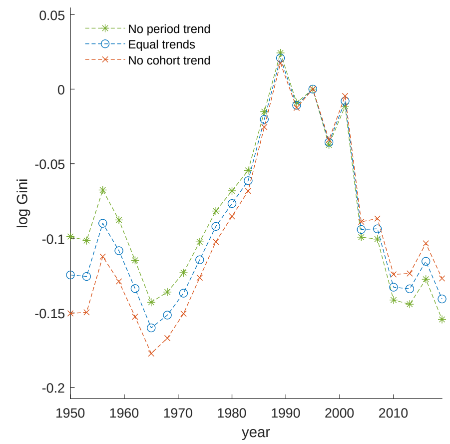

(d) Period profile: log Gini coefficient

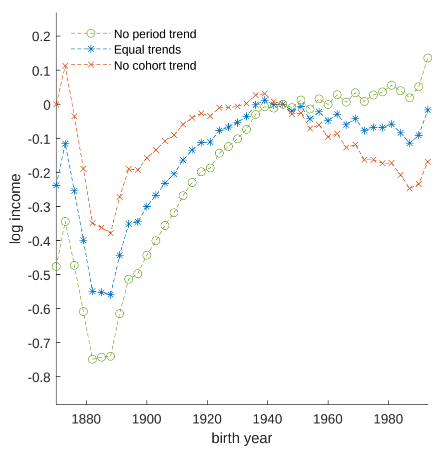

(e) Cohort profile: log mean income

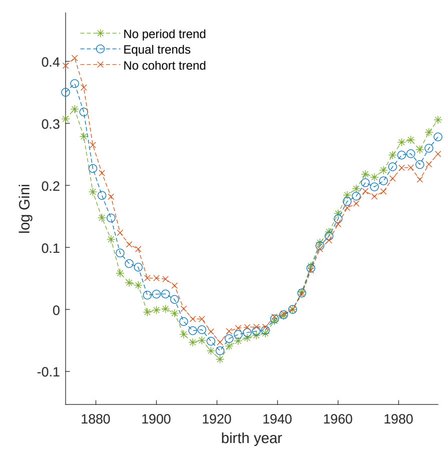

(f) Cohort profile: log Gini coefficient

Figure 14: Age-, period-, and cohort profiles of log mean income and log Gini coefficients for households with non-college educated household head in the SCF+ data under different normalizations for the linear trends. The period effect 1995 and cohort effect 1945 are normalized to zero.

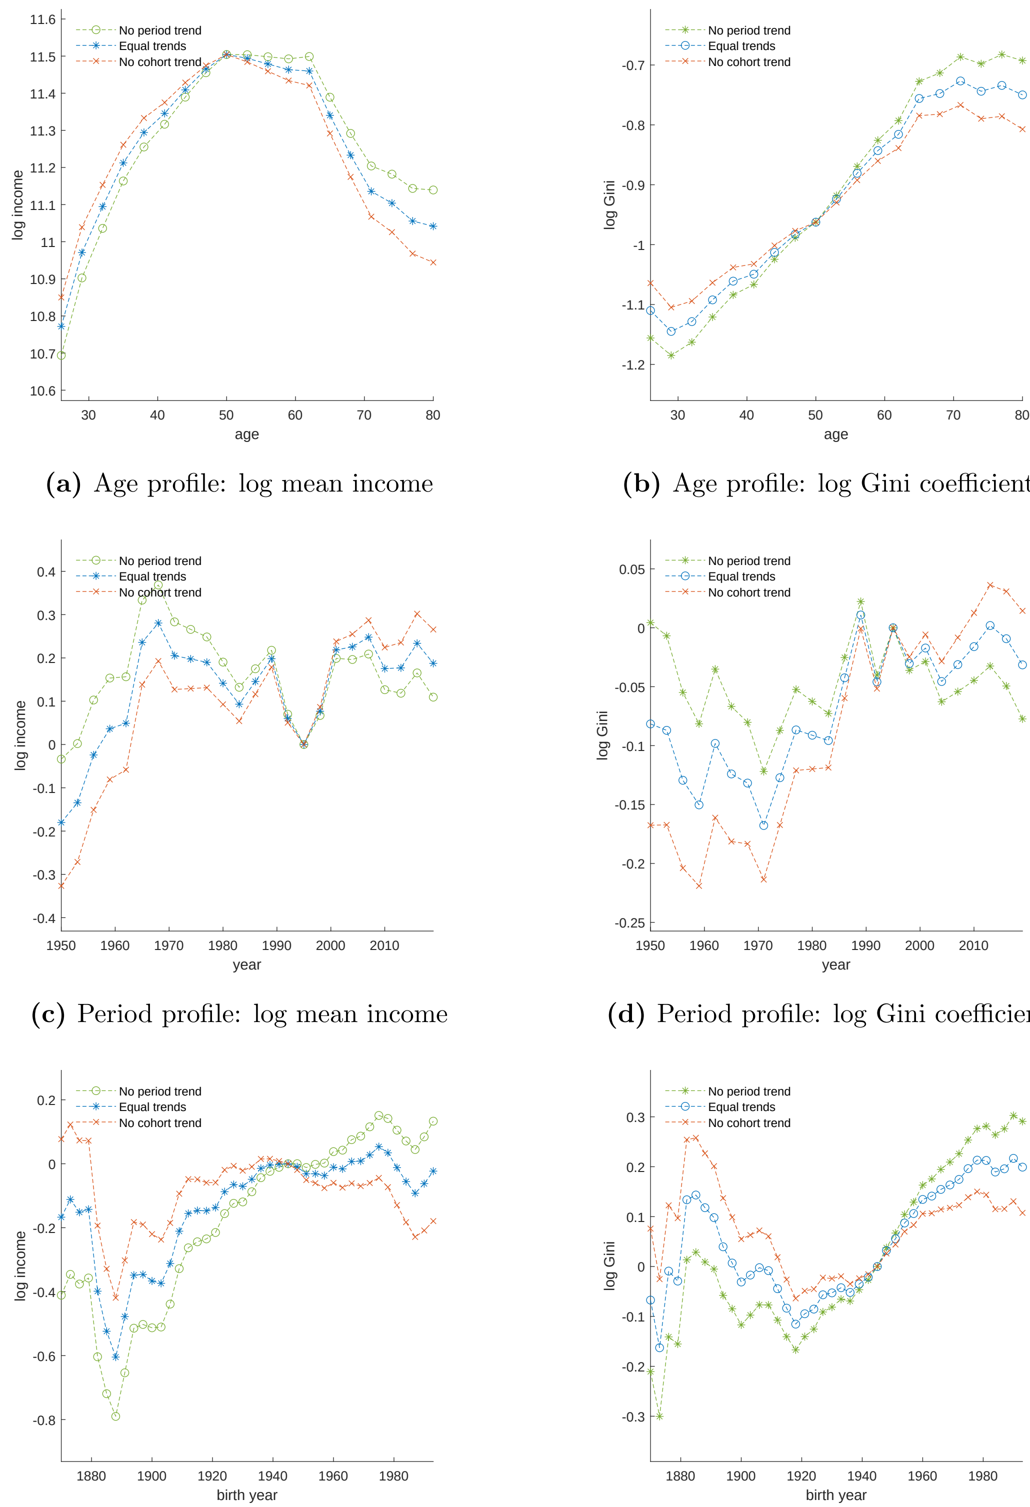

(a) Age profile: log mean income

(b) Age profile: log Gini coefficient

(c) Period profile: log mean income

(d) Period profile: log Gini coefficient

(e) Cohort profile: log mean income

(f) Cohort profile: log Gini coefficient

Figure 15: Age-, period-, and cohort profiles of log mean income and log Gini coefficients for households with college educated household head in the SCF+ data under different normalizations for the linear trends. The period effect 1995 and cohort effect 1945 are normalized to zero.

B Robustness of estimation results from age-period-cohort model

We estimate our age-period-cohort model on a balanced sample in age and survey waves. As a consequence, only a fraction of birth cohorts are observed at all ages, while most cohorts are observed when they are either relatively young or relatively old. It also means that for some cohorts we have much fewer observations that for others. One may thus be concerned that this panel structure affects our estimation results. For example, if the age profile of income inequality differed across decades, the age-period-cohort model would partially attribute the induced variation to cohort effects. In this case, estimating the age-period-cohort models on different sub-periods would result in different age and cohort profiles.

We address this concern by re-estimating the age-period-cohort model on different restricted time windows and by comparing the resulting age and cohort profiles. Reassuringly, we find that profiles estimated on restricted samples are similar to the profiles estimated on the full sample. In figures 16, 17, 18, and 19, we plot the age and cohort profiles from estimating the age-period-cohort model on different sub-periods consisting of 10 consecutive waves in the CPS and of 5 consecutive waves in the SCF+ data. For each sub-period, we re-estimate all parameters of the model such that the age and cohort profiles can take different shapes. In each case, we normalize the trend in the age profile to be equal to the trend estimated in the full sample under the normalization of equal cohort and period trends. This normalization, however, does not force the linear slopes of the period and cohort profiles in the models estimated on the sub samples to be identical to the ones in the full model, nor does it constrain the nonlinear age, period, and cohort effects. Reassuringly, we find that the overall shapes of the age and cohort profiles do not depend on the sub-period used to estimate the age-period-cohort model.

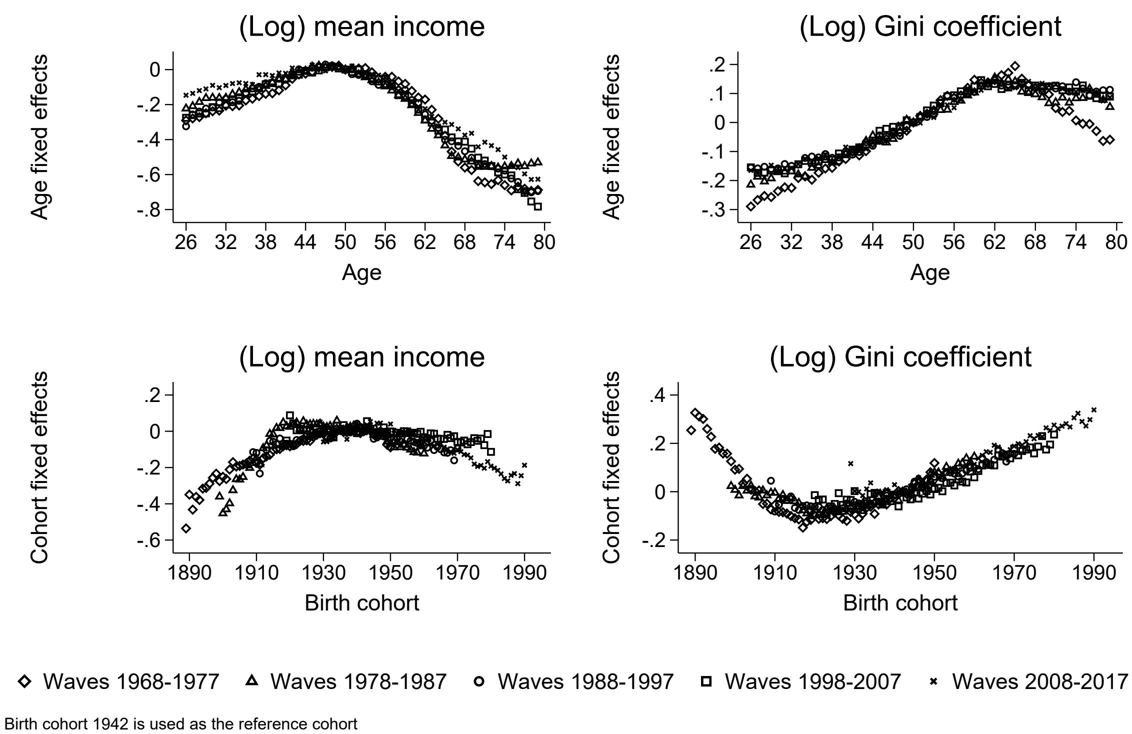

Figure 16: CPS, Non-college

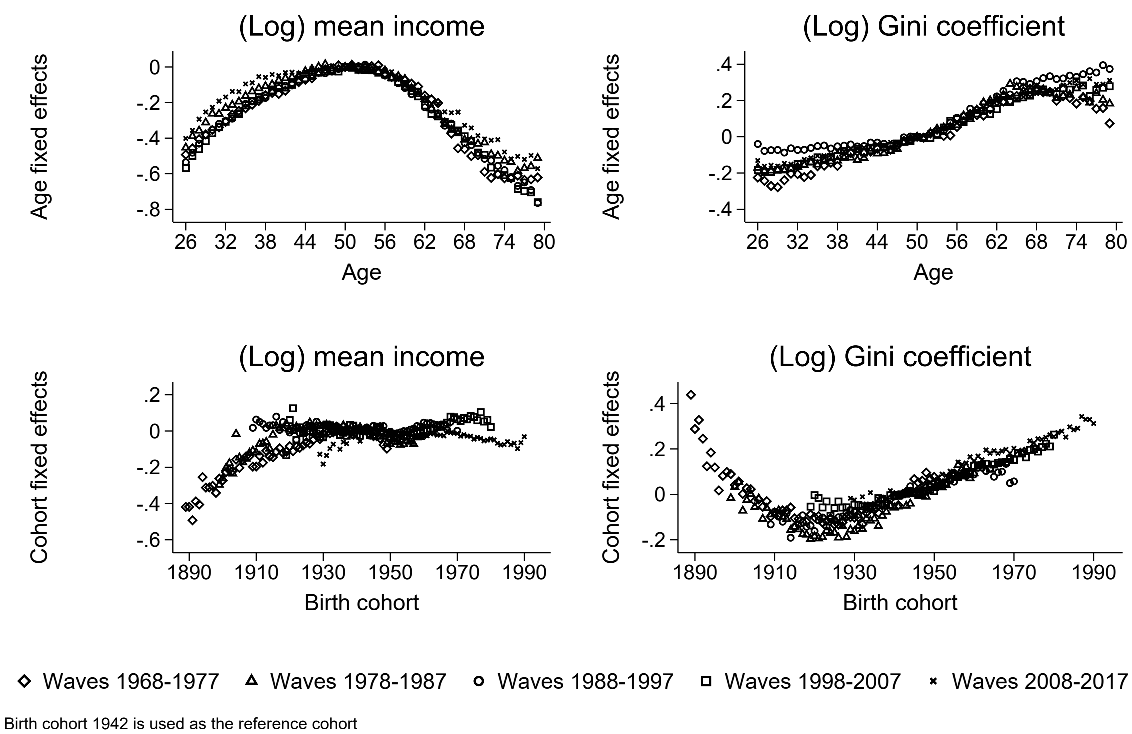

Figure 17: CPS, College

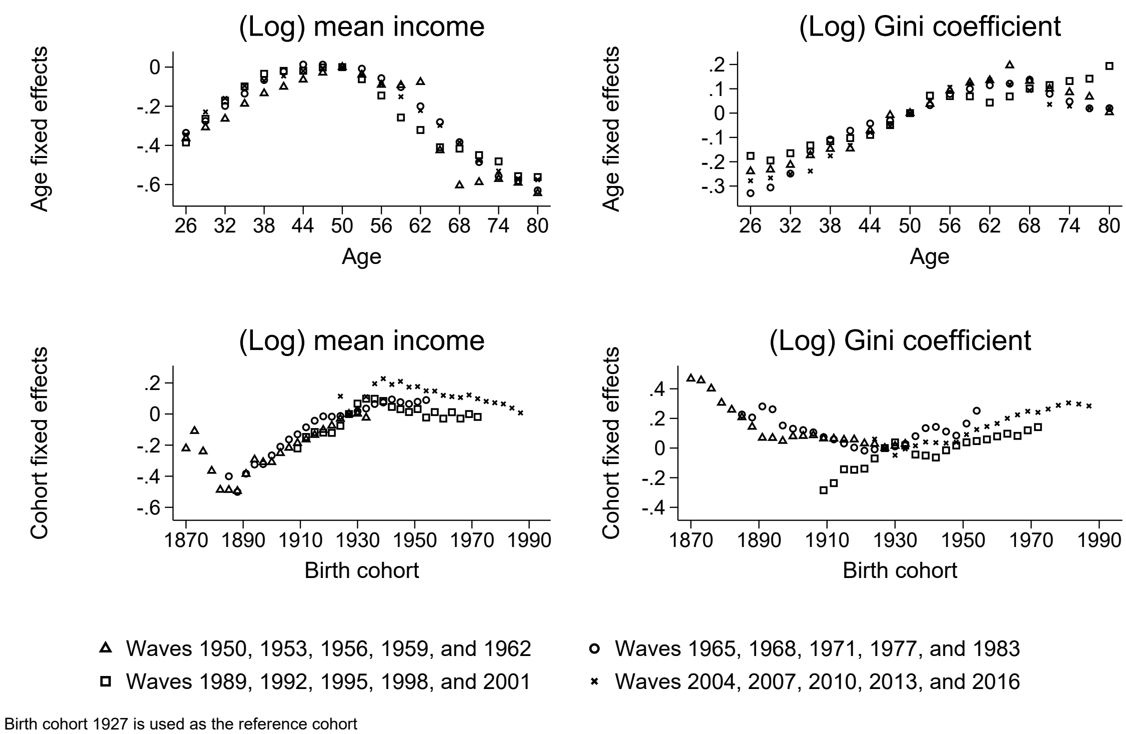

Figure 18: SCF+, Non-college

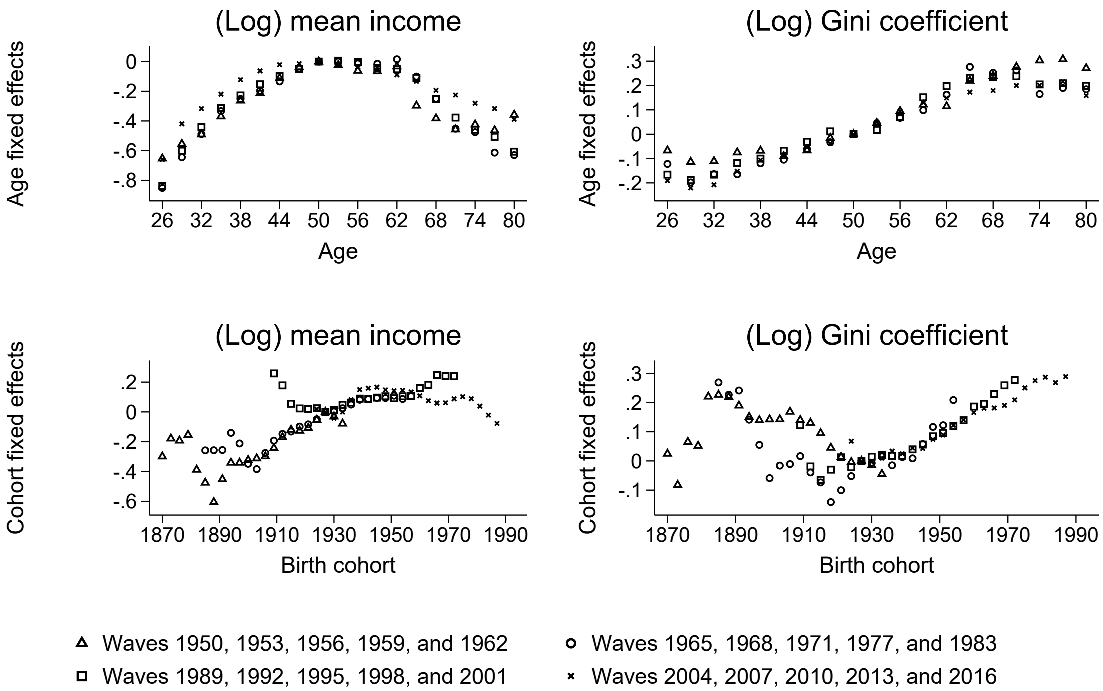

Birth cohort 1927 is used as the reference cohort

Figure 19: SCF+, College

C Fit of the age-period-cohort model

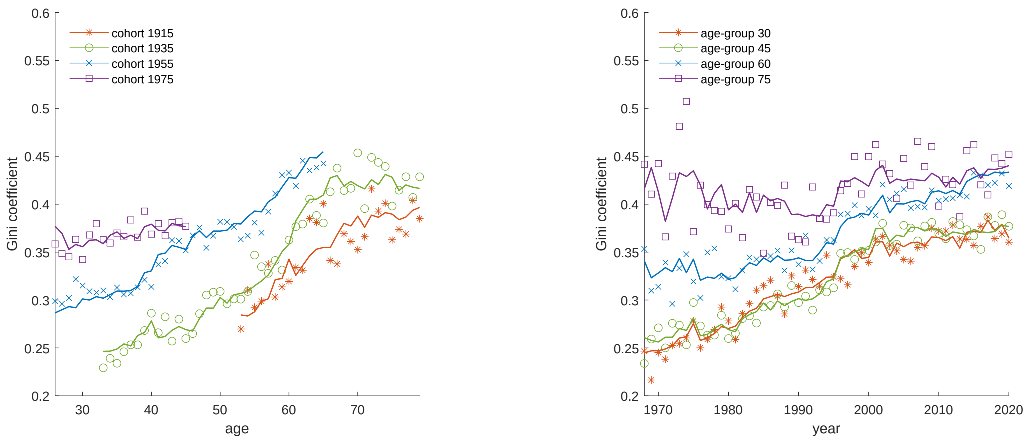

(a) Age profiles for selected birth cohorts

(b) Time trends for selected age groups

Figure 20: Gini coefficients for households with a college-educated household head (CPS data). The solid lines show the fit of the age-period-cohort model.

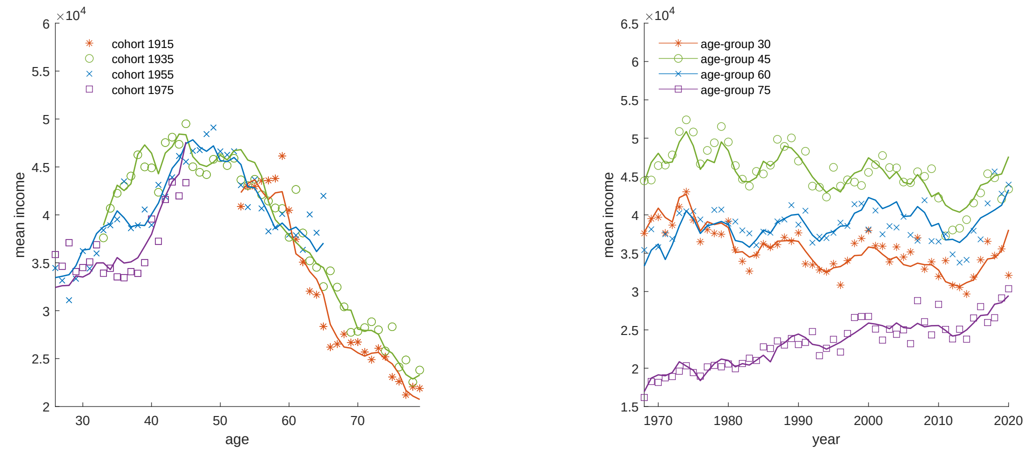

(a) Age profiles for selected birth cohorts

(b) Time trends for selected age groups

Figure 21: Mean incomes for households without a college-educated household head (CPS data). The solid lines show the fit of the age-period-cohort model.

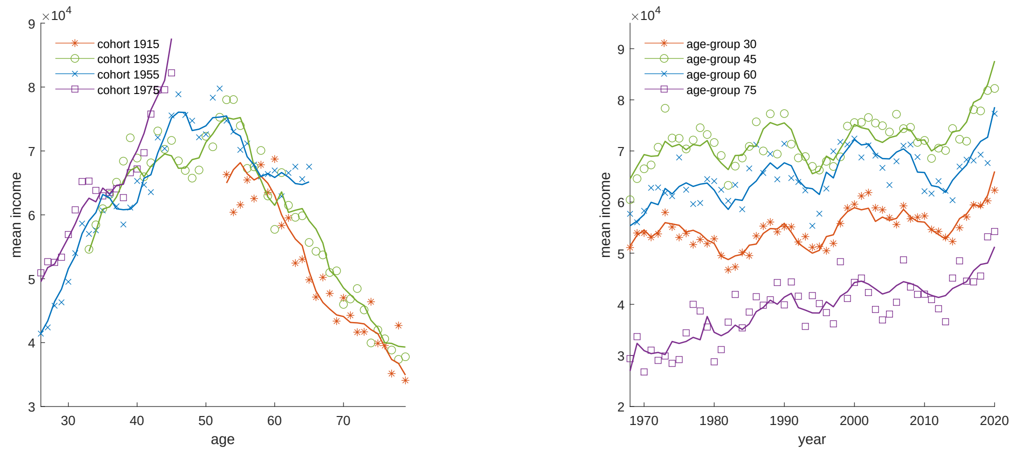

(a) Age profiles for selected birth cohorts

(b) Time trends for selected age groups

Figure 22: Mean incomes for households with a college-educated household head (CPS data). The solid lines show the fit of the age-period-cohort model.

D The role of demographic change in the past under alternative normalizations

In figure 7 in the main text, we show that under our baseline normalization demographic change explains a large share of the observed increase in income inequality in the past – especially since the 1990s. How much of the actual increase in the income Gini coefficient following a given base year can be accounted for by demographic change depends on the normalization of the linear age, period, and cohort trends. In figure 23 we plot the actual increase in the income Gini over an 21 year period following each base year together with the increase in the counterfactuals over the same time period. We again see that the role of demographic change has become more important starting in the 1970s. The role of demographic change is always stronger if we assume that there are no period trends in income inequality over the sample period and the cohort profile is therefore estimated to have a positive trend. Nevertheless, we find that a large share of the observed increase in income inequality since the 1990s can be accounted for by demographic change even if we impose that there is no linear trend in cohort effects and we instead allow the period effects to exhibit a strong positive linear trend.

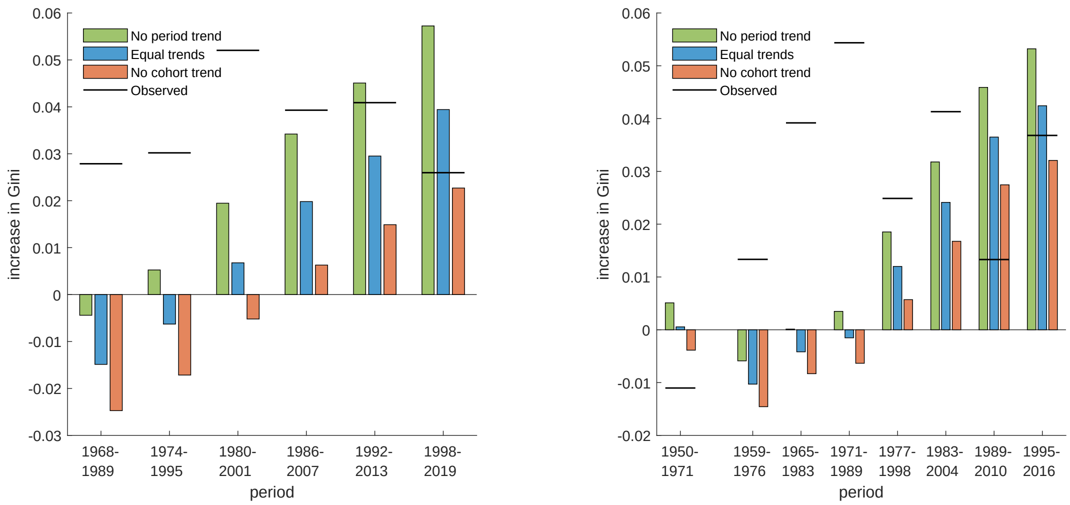

(a) CPS **(b)** SCF+

Figure 23: Increase in income Gini driven by demographic change over periods of 21 years.

E Vintage predictions

In the baseline counterfactuals, we use age, period, and cohort effects that are estimated on the full sample. In particular, even when we consider a base year in the past, we estimate the relevant cohort fixed effects from all available years in our data set – including all years after the base year. If the age-period-cohort model describes the data generating process well, this approach is innocent and will increase precision while not biasing the estimates. In appendix B we show that estimating the age-period-cohort model on restricted sub-periods does not appear to affect the estimated nonlinear cohort effects by much.

The baseline counterfactuals do not, however, necessarily correspond to the predictions that we would have made, had we written the paper in the respective base year. Besides the fact that we would not have obtained exactly the same estimates, the normalization of the linear trends also depends on the time period covered in the data. In figure 24, we recompute figure 7 from the main text with the exception that for each base year, we now we only use data up until that year to estimate the age, period, and cohort effects that we use for the predictions. Moreover, for each base year, we force the linear trends in period and cohort profiles to be of equal size. This exercise corresponds to computing vintage predictions, which show the predictions that we would have made, had we written the paper in the respective base year.

Panel (b), which shows the vintage predictions for the SCF+ data, looks remarkably similar to the corresponding panel in figure 7 in the main text. While we would not have predicted any increase in income inequality due to demographic change had we written the paper in the early 1970s, we would have predicted a steep increase had we written the paper in the early 1990s. In fact, as we find in figure 7 in the main text, the predicted increase in income inequality since the 1990s accounts for all of the observed increase since then.

The findings are somewhat different for the CPS data. While we also would have predicted all of the observed increase in income inequality had we written the paper in the 1990s, we find that we would have predicted a significant increase even if we had written the paper in the early 1970s. This stands in contrast to the findings in the main text and the vintage predictions derived from the SCF+ data. This discrepancy stems from the normalization of the linear trends.

Because the SCF+ data covers a longer time period that includes the 1950s and 1960s during which income inequality did not increase, the linear trends in period and cohort profiles are comparatively small and the choice of normalization matters little. Hence, restricting the data set by dropping later years hardly affects how we normalize the linear trends. The CPS data on the other hand covers a time period throughout which income inequality has increased and the linear trends in period and cohort profiles, which we force to be equal, are therefore steeper. Moreover, we estimate the profile of period effects to be steeper at the beginning of the sample period relative to the end. Hence, restricting the data set by dropping later years combined with the equal trends assumption results in steeper cohort profiles for earlier base years. Consequently, demographic change drives up income inequality more strongly for earlier base years.

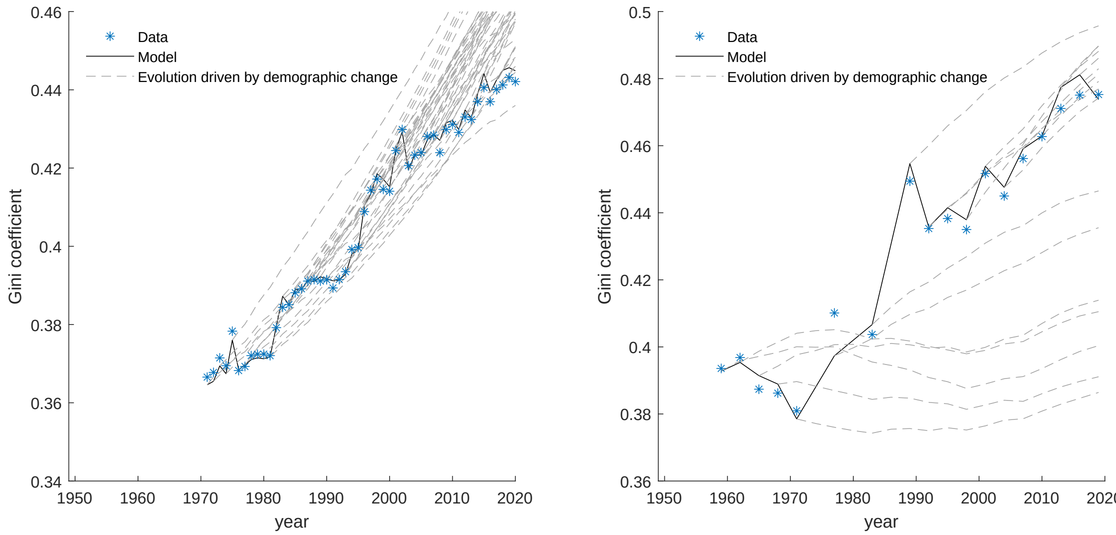

(a) CPS data

(b) SCF+ data

Figure 24: Vintage predictions

F The effect of demographic change on the top five percent income share

In our main analyses in section III we measure aggregate inequality using the Gini coefficient. However, there is increased attention on the evolution of top income shares. It is therefore important to know whether our main findings also apply to these inequality measures. In this appendix, we study the effect of demographic change on the evolution of the share of income received by households in the top five percentiles. We conduct this analysis using the SCF+ data without dropping the top 1 percent of households. This is feasible because the SCF+ data is not topcoded.

To derive the evolution of the top 5 percent income share in the thought experiment, we follow the same steps as discussed in section II. That is, we construct counterfactual mean incomes and Gini coefficients for population subgroups using the same estimated age, period, cohort effects as in the main analysis. To derive the aggregate income distribution, however, we use a different parametric distribution to model subgroup income distributions. We do this because the maximum entropy distribution introduced in section II.3 has a thinner right tail than observed income distributions. As a consequence, if we use the same parametric distribution as in section II.3, we underestimate the top 5 percent income share in the aggregate income distribution as can be seen in figure 25.

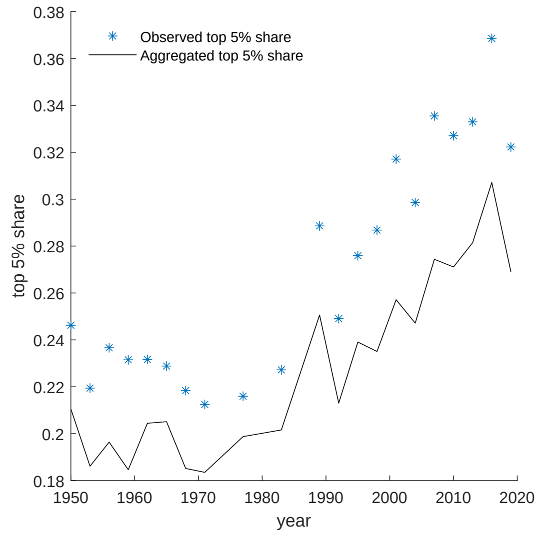

Figure 25: Aggregated top 5 percent income shares using maximum Shannon entropy distribution.

To match the top five percent income share in the aggregate income distribution, we change the notion of entropy from the standard Shannon entropy to Tsallis entropy, which is a generalization of Shannon entropy.[27] Maximum Tsallis entropy distributions have been used in applied sciences to model phenomena that follow a power law. We model the subgroup income distributions with the parametric distribution that maximizes Tsallis entropy for given mean and Gini coefficient.[28] This distribution has a free parameter, q, which governs the thickness of the tail. We find that with $q = 0.67$, we are able to match the top five percent income share in the aggregate income distribution quite well.

In figure 26, we show the evolution of the top five percent income share in our thought experiment under the baseline normalization of equal trends in period and cohort effects. The blue stars show the evolution of the top 5% income share in the SCF+ data, and the solid black line shows the evolution of the top 5% income share in the aggregate distribution derived from the predicted values of the age-period-cohort model. The dashed lines show the evolution of the top 5% income share in our thought experiment corresponding to different base years. Our findings for the top 5% income share are in line with our findings for the Gini coefficient in section III. Demographic

[27] Tsallis entropy is defined as $H_T(f) = \frac{1}{q-1}\left(\int_{-\infty}^{\infty} f^q(x)dx\right)$, where f is a density function and q is a parameter. In the limit $q \to 1$, this definition coincides with the standard Shannon entropy, $H(f) = \int_{\infty}^{\infty} f(x)\ln f(x)dx$.

[28] This distribution is solved for in Preda et al. (2015).

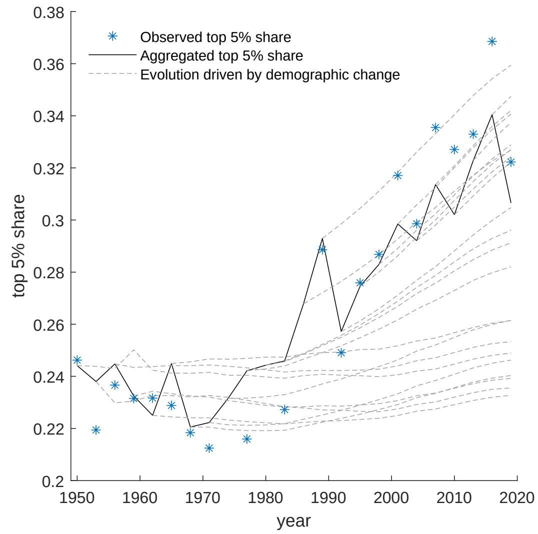

Figure 26: The evolution of top 5 percent income share as driven by demographic change in the thought experiment using SCF+ data.

change has little effect on income inequality if we fix the economic environment to a base year in the 1950s, 1960s, or the 1970s. If we fix the economic environment to a base year in the mid-1990s, however, we find that demographic change can account for most of the observed increase in top 5% income share since then.

G Normalizing the cohort trends using NLSY data

In this section, we use the results in table 2 to derive a new normalization for the linear trends in the age-period-cohort model. That is, we choose a linear trend for the cohort effects that, together with the estimated nonlinear cohort effects, produces the differences of 0.08 for college-educated households and 0.12 for non-college-educated households in the log Gini coefficients between cohorts 1959 and 1982. The idea is that if we know the difference in cohort effects between two cohorts, then we can back out the value for the linear trend in the cohort profile. And knowing the trend in the cohort profile allows us to estimate the linear trends in age and period profiles. We derive our main results under this new normalization and depict them in figures 27 and 28. The qualitative results remain the same while quantitatively the results are between those derived under the baseline normalization and the no cohort trend normalization.

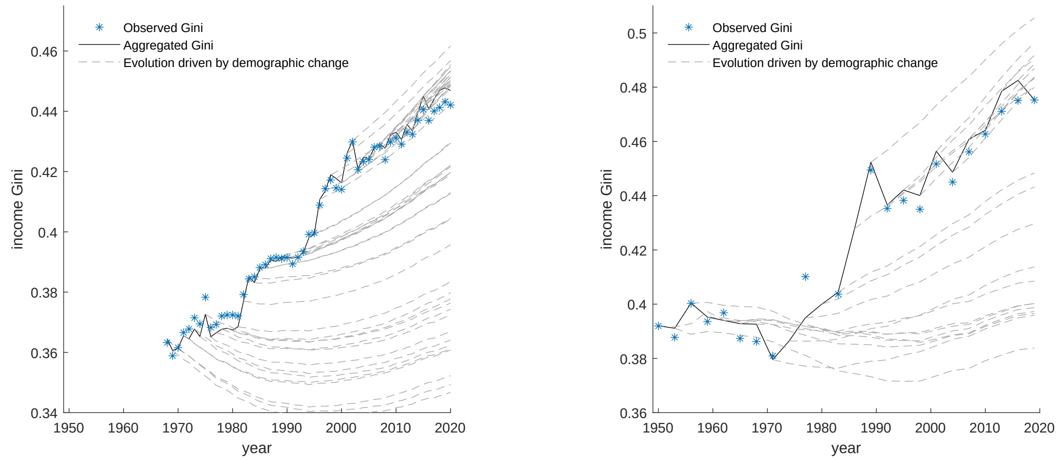

(a) CPS **(b)** SCF+

Figure 27: The effect of demographic change on income inequality in the past under the normalization derived from the NLSY data.

H Aggregation of Gini coefficients with Maximum Entropy distribution

Figure 6 in the main text shows that we can match the population-level Gini coefficients extremely well by using a parametric distribution to describe subgroup income distributions. We choose the distribution that is supported on the positive real line and maximizes entropy given our estimated moments, mean income and income Gini coefficient. Here, we consider lognormal and gamma distributions as alternative two-parameter distributions to describe incomes at the subgroup level and compare them to the distribution used in the main text.

The lognormal and the gamma distributions are maximum entropy distributions for given mean and variance of log income, and mean income and mean logarithmic deviation of income, respectively. Figure 29 shows that using these alternative income distributions and targeting their respective characterizing moments at the subgroup level results in a worse fit for the population-level Gini coefficient.

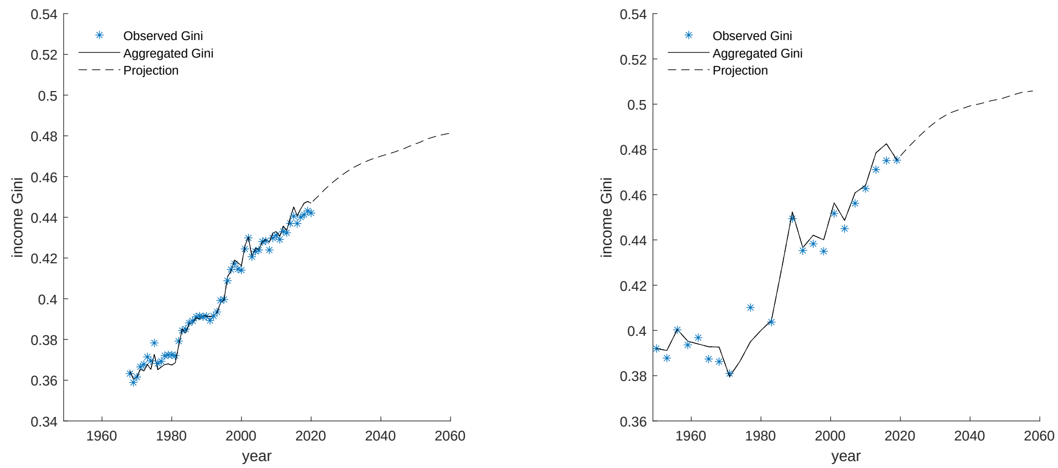

(a) CPS **(b)** SCF+

Figure 28: The effect of demographic change on income inequality in the future under the normalization derived from the NLSY data.

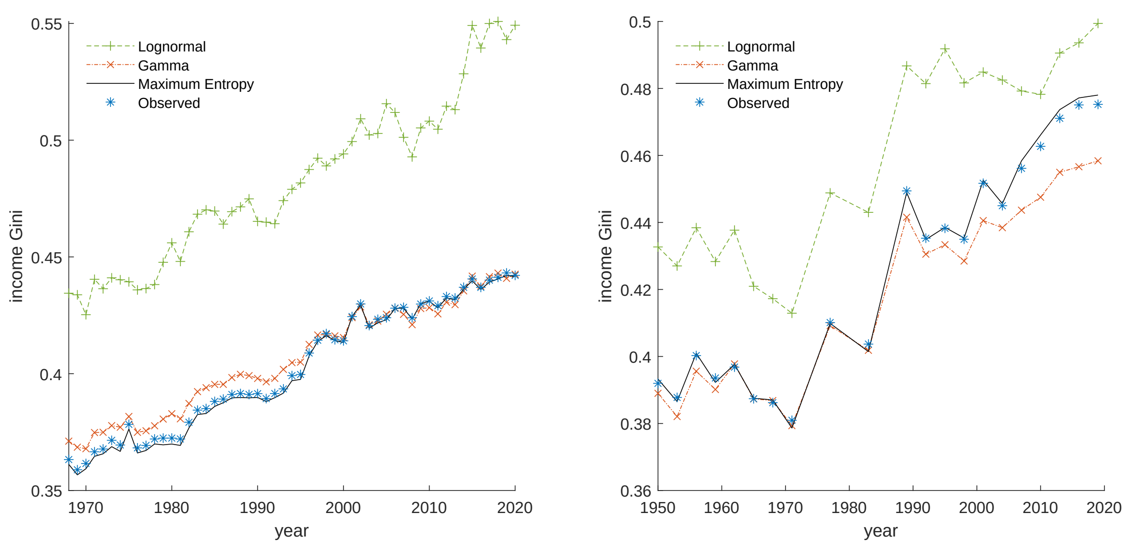

(a) CPS **(b)** SCF+

Figure 29: Aggregated income Gini coefficients using different parametric distributions.