

More Unequal We Stand? Inequality Dynamics in the United States, 1967–2021*

Jonathan Heathcote

Fabrizio Perri

Giovanni L. Violante

Lichen Zhang

July 2023

Abstract

Heathcote et al. (2010) conducted an empirical analysis of several dimensions of inequality in the United States over the years 1967-2006, using publicly-available survey data. This paper expands the analysis, and extends it to 2021. We find that since the early 2000s, the college wage premium has stopped growing, and the race wage gap has stalled. However, the gender wage gap has kept shrinking. Both individual- and household-level income inequality have continued to rise at the top, while the cyclical component of inequality dominates dynamics below the median. Inequality in consumption expenditures has remained remarkably stable over time. Income pooling within the family and redistribution by the government have enormous impacts on the dynamics of household-level inequality, with the role of the family diminishing and that of the government growing over time. In particular, largely due to generous government transfers, the COVID recession has been the first downturn in fifty years in which inequality in disposable income and consumption actually declined.

JEL Codes: D12, D31, E21, H53, J31.

Keywords: Wages, Hours Worked, Earnings, Income, Consumption, Wealth, Inequality, Redistribution, Recessions, Surveys.

*Heathcote: Federal Reserve Bank of Minneapolis and CEPR, heathcote@minneapolisfed.org ; Perri: Federal Reserve Bank of Minneapolis and CEPR, fperri@umn.edu; Violante: Princeton University, CEPR, IFS, IZA, and NBER, violante@princeton.edu; Zhang: University of Hong Kong lichenz@hku.hk. We are grateful to Nezih Guner, Moritz Kuhn and several seminar participants for valuable comments, and to Chenming Zhang for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

The literature discussing the evolution of economic inequality in the United States is vast. There is a general sense that “inequality has been increasing,” but there is a lively debate about how much it has risen and about what to do about it. One reason for the lack of consensus is that inequality is an elusive concept. Different authors have focused on inequality in different economic outcomes (wages, earnings, disposable income, consumption, or wealth), and at different levels of analysis (individual versus household). They have also used using different measures of dispersion (percentile ratios, variances of logs, Gini coefficients), and relied on different data sets and sample restrictions. Even for researchers in the area, it is hard to know how disparate findings in the literature fit together, where there is ample consensus, and what the central unresolved questions are.

Our earlier paper, [Heathcote et al. \(2010\)](#), approached the topic using the household budget constraint as an organizing device. We started with individual wages. Then we incorporated individual hours worked to study earnings, added other household members to analyze household earnings, and moved to broader income measures by including unearned income, taxes, and transfers. Finally, we explored consumption and wealth. In that paper, we used data up through 2006. In this paper, we extend the analysis through to 2021. Thus, we are able to investigate what happened to inequality during the dramatic Great Recession and subsequent recovery, and during the short, but sharp, COVID cycle.

We explore all of the most widely-used household surveys covering income, consumption and wealth: the Current Population Survey (CPS) Annual Social and Economic Supplement (also known as the March CPS), the American Community Survey (ACS), the Consumer Expenditure Survey (CES), the Panel Study of Income Dynamics (PSID), and the Survey of Consumer Finances (SCF). These surveys have three advantages relative to administrative micro data. The first is that they are publicly available, so all our findings can be readily reproduced. A second one is that they offer comprehensive information on demographics, on virtually all components of income, and on hours worked, which allows us to separate the roles of hourly wages and labor supply in earnings dynamics. A third one is that these are household surveys, so we can link individual- and household-level measures, and move from individual income to household income and consumption.

There are, however, two main concerns with survey data. The first one is that they might not be representative. In particular, one might worry that households at the very top of the income distribution are under-represented. The second concern is that, in an-

swering surveys, respondents may under-report some forms of income.¹ To assess these limitations, we start by exploring the extent to which aggregate income measures constructed from our surveys replicate their official counterparts in national accounts. We show that survey-based per capita measures of wage and salary income line up closely with National Income and Product Accounts (NIPA) measures, and that the concordance is especially strong for the March CPS and the ACS. In contrast, there is a substantial gap between surveys and the NIPA for broader measures of per capita income, with this gap widening over time. The main source of this discrepancy is that survey estimates of self-employment and capital income are much smaller than their NIPA counterparts, and the income share of these components is growing over time.² For consumption expenditures, we show that non-health, non-housing spending in the CES tracked NIPA consumption poorly in the early years of the CES. Although the CES consumption measure still underestimates the official level in recent years, it tracks the NIPA trend and cyclical fluctuations much better.

Next, we turn to inequality. We begin with wages, focusing on a sample of working age individuals with a minimum labor force attachment: those who work at least part-time one quarter per year. Our key new findings regarding wage inequality, relative to those of our previous paper, are that wage dispersion at the bottom of the wage distribution (i.e., below the median) has remained relatively stable over the past 20 years, while inequality has kept growing at the top. Wage differentials by education, age, and race groups have stabilized since the turn of the century, while the gender wage gap has continued to narrow. Regarding occupational differences the most significant feature of the last two decades has been the rise of the wage for men in non-routine manual occupations relative to routine (manual and non manual) occupations.

With respect to inequality in individual earnings, we confirm two key results from [Heathcote et al. \(2020\)](#). First, the rise in earnings inequality at the top is mostly a secular trend, while there is a strong cyclical component to earnings dispersion at the bottom. Second, the dynamics of earnings dispersion at the top are entirely driven by long-run trends in hourly wages, while hours worked are very stable. Instead, fluctuations in earnings at the bottom of the distribution largely reflect movements in annual weeks worked, driven by unemployment and labor force participation.

Next, we shift our focus from the individual to the household. Here, we broaden

¹Another drawback of survey data relative to administrative data is their smaller sample size which limits the granularity of the analysis. See [Guvenen et al. \(2022\)](#) for a discussion.

²These forms of income are extremely concentrated at the top of the distribution, and top-income households tend to be undersampled in surveys.

the sample to all working age households, irrespective of hours worked by the adult members. Because of income pooling, the household is potentially an important source of insurance and redistribution.³

Has the role of the household as a source of redistribution changed over time? To address this question, we develop a novel decomposition of the gap between dispersion in individual earnings and dispersion in equivalized household earnings. We show that the importance of the household as an inequality-reducing mechanism has diminished steadily over the past five decades, for two reasons. First, the correlation between earnings within couples has increased over time, reducing the scope for within-household insurance. Second, the gender gap in earnings has declined, so women are now much less reliant on spousal earnings.

A second key source of insurance and redistribution across households is the government tax and transfer system. We compare inequality in market income, pre-tax income (market income plus government transfers), and disposable income (pre-tax income minus taxes paid). The dynamics at the top are quite similar across all these different income measures, and largely mirror those of inequality at the top for individual wages and earnings: slow but steady growth over the last 20 years. Conversely, at the bottom of the income distribution, the government plays a prominent role. Below the median, inequality in market income is very counter-cyclical, as labor earnings in the lowest quantiles are severely depressed by surging unemployment. However, transfers and –to a lesser extent– taxes drastically compress the level of inequality. One stark illustration of the strength of this channel is that, among the poorest 20 percent of households, average (inflation-adjusted) equivalized head-of-household earnings fell dramatically over our sample period, while household earnings rose modestly, and disposable income is now 70 percent higher than in 1967. Besides affecting long-run trends, government intervention largely eliminates the cyclicalities in inequality at the bottom, as transfers increase in recessions and create a floor under disposable income. Notably, the expansion of automatic stabilizers and ad-hoc government transfers during the COVID recession was so massive that inequality in disposable income *declined* for the first time across the eight recessions that occurred in the fifty years we study.

In the last section of the paper, we turn to consumption and wealth inequality. Here, we show that a striking finding from our earlier paper survives the addition of 15 more years of data: with all the caveats that apply to the measurement of consumption expendi-

³To fix ideas with a simple example, suppose all households are married couples where one spouse works and the other does not, and all working spouses earn the same amount. This economy would display high inequality in individual earnings, but zero inequality in household earnings.

tures in surveys, consumption inequality appears remarkably stable over time. We show that this finding applies to both the CES and the PSID data sets. With respect to wealth inequality, we find that estimates from the PSID are very much in line with those from the SCF, which are widely considered to offer a reliable picture for the United States. We also find that the distribution of wealth in the top half of the income distribution amplifies both the level and the growth in income differences between the top and the median.

Overall, we distill four key messages from this body of evidence. First, measures of wage and salary income in all the surveys we study are quite consistent with corresponding measures in the NIPA, in terms of levels, trends and cyclical fluctuations. The match is not as good for other measures of income (i.e., capital and business income) or for consumption expenditures. For wealth, the SCF and the PSID match aggregate wealth from the Flow of Funds well, while the CES does not. Second, both individual and household level inequality at the top have kept increasing since 2000 –though less rapidly than in previous years– while inequality dynamics at the bottom has been dominated by cyclical fluctuations rather than by a long run trend. Inequality above the median has been driven by wage differentials, while below the median, employment differentials are the key driver. Third, income pooling within the family and redistribution by the government have enormous impacts on the level and trend of household inequality. Over time, the role of income pooling within the household has been declining, while the role of government redistribution has been increasing. Fourth, inequality in measured consumption expenditures in both the CES and the PSID has remained remarkably stable over time.

Literature. The literature on the measurement of income inequality is vast and, as we discussed earlier, has not reached a full consensus yet.⁴

The area where there is most agreement is the one focused on individual wages and labor earnings. Both survey and administrative data point to a steady rise in wage dispersion that started in the early 1970s and exhausted its impetus in the early 2000s (for recent analyses, see [Hoffmann et al., 2020](#); [McKinney et al., 2022](#)). The divergence in trends between the top and bottom halves of the wage distribution was already identified for the early 2000s by [Autor et al. \(2008\)](#). We confirm that this observation remains true fifteen years later. We also corroborate the view that occupational tasks (abstract, routine, and manual) are a useful lens into wage inequality dynamics over the last four decades ([Autor et al., 2003](#)), but more so for men than for women –a fact that has not received much

⁴For conflicting perspectives on the size and sources of the rise in income inequality see, for example, [Piketty et al. \(2017\)](#), [Smith et al. \(2019\)](#), and [Gramm et al. \(2022\)](#).

attention.

A common feature of this body of work is that it restricts attention to workers with strong attachment to the labor force (and, at least initially, to men only). We expand the analysis to all individuals, and we illustrate the importance of participation and unemployment in driving secular and cyclical dynamics of inequality at the bottom of the earnings distribution.⁵

Our analysis at the household level connects to a large quantitative literature emphasizing the key roles of female labor supply (Juhn and Murphy, 1997; Hyslop, 2001; Blundell et al., 2016), sorting (Mare, 1991; Greenwood et al., 2014), income pooling among spouses (Altonji et al., 2022; Santos and Tertilt, 2023), and government redistribution (Heathcote et al., 2014; Borella et al., 2022) as determinants of inequality dynamics.

A topic where the literature is far from reaching a consensus is consumption inequality. A number of studies conclude that since the early 1980s consumption inequality has risen much less than income inequality (Johnson and Shipp, 1997; Slesnick, 2001; Krueger and Perri, 2006; Heathcote et al., 2010). All of them rely on the Interview component of the CES, which provides the most comprehensive data on household spending for a nationally representative sample. Many consumption items that are frequently purchased are, however, severely under-reported in the Interview component of the CES, but are better recorded in its Diary component. By combining Interview and Diary, Attanasio et al. (2007) find a rise in consumption inequality that is larger than the one in those initial studies, but still much below the surge in income inequality. At the opposite end of the spectrum, estimates from Aguiar and Bils (2015), who employ a demand system approach, indicate that, over the period from 1980 to 2010, consumption inequality rose even more than income inequality.⁶ Recently, Meyer and Sullivan (2023) revisit these calculations by focusing on well measured components of the CES reported at a high and stable rate relative to national accounts. They conclude that the rise in overall consumption inequality was small. Our findings are in line with theirs.⁷

A voluminous literature in macroeconomics has studied the evolution of wealth inequality over time in the U.S. (see, e.g., Bricker et al., 2016; Kuhn et al., 2020) and its determinants (for surveys, see Benhabib and Bisin, 2018; De Nardi and Fella, 2017). Much

⁵Here we are interested purely in measurement, but it is important to remind the reader that there exists a parallel literature investigating the primitive causes of the surge in wage inequality, such as technological transformations, globalization, and institutional change. For surveys on these three forces, see Acemoglu and Autor (2011), Helpman (2018), and Fortin and Lemieux (1997).

⁶See Attanasio and Pistaferri (2016) for a survey.

⁷Our finding that the rise in household disposable income dispersion is much lower than that in market income gives some context to our small estimated growth in consumption inequality.

of this research is based on the SCF, which is widely recognized as an accurate representation of both the level and the dynamics of wealth dispersion, even for the upper tail.⁸ Our contribution to this literature is to compare measures of wealth average and dispersion in the SCF to those in the PSID and the CES, where wealth is often combined with information on income and consumption expenditures to shed light on saving behavior and .⁹

Finally, only a handful of papers contain a comprehensive joint analysis of multiple dimensions of inequality, as ours does. Recent examples are [Kuhn and Ríos-Rull \(2016\)](#), which builds on [Diaz-Gimenez et al. \(1997\)](#), and [Fisher et al. \(2018\)](#). With respect to these papers, we discipline our investigation through the household budget constraint, and build step by step: from hourly wages to individual labor earnings, household earnings, household market income, household disposable income, and consumption expenditures. By doing so, we can isolate the roles played by labor supply, the family, the government, and financial markets in mediating the impact of the primitive forces driving shifts in the distribution on welfare.

The rest of the paper is organized as follows. Section 2 outlines our approach based on the budget constraint. Section 3 describes our data sources. Section 4 compares aggregates in surveys to their counterparts in official NIPA data. Section 5 studies inequality in wages and earnings at the individual level. Section 6 extends the analysis to household-level distributions of earnings, market income, disposable income, and consumption expenditures. Section 7 concludes. The Appendix contains more details on data, methodology, and findings.

2 An organizing device

The device around which we organize our analysis is the household budget constraint, reported in equation (1):

$$c + s = \sum_{i=1}^N w_i h_i + d + b^p + b^g - \tau. \quad (1)$$

⁸To further improve coverage of the upper tail, some researchers augment the SCF with members of the *Forbes* 400.

⁹For example, [Gourinchas and Parker \(2002\)](#) fit a life-cycle model with precautionary saving to data on consumption and wealth from the CES, [Vissing-Jørgensen \(2002\)](#) uses data on wealth and consumption from the CES to identify the determinants of stockholding, and [Aguiar et al. \(2020\)](#) pair wealth and consumption data from the PSID to study the nature of household hand-to-mouth status.

We start from the resources available to the household (the right-hand side of the budget constraint). Our first object of interest is w_i , which is the wage per unit of labor supplied by each household member. Inequality in wages has been widely studied, but that is just the starting point of inequality in resources. The second term is labor supply of each household member, h_i , which determines individual labor earnings, $w_i h_i$. Labor supply plays an important role in shaping earnings inequality over time –because of trends in labor force participation– and over the business cycle –because of unemployment dynamics, as emphasized by [Castaneda et al. \(1998\)](#). The next step is the summation of labor earnings across household members which, we will show, plays an important and time-varying role in shaping inequality. Beyond labor earnings, other determinants of the distribution of resources are asset income d (which includes interest, dividends, and rental income), private transfers b^p , government transfers b^g (capturing items like unemployment insurance and welfare benefits), and taxes τ . These resources can be used for consumption expenditures, c , or for saving, s .

We will focus on the following income measures: (i) labor earnings defined as wage and salary income plus self-employment income; (ii) market income, which is equal to earnings plus asset income d and private transfers b^p ; (iii) pre-tax income, which is equal to market income plus government transfers b^g ; and (iv) disposable income, which is equal to pre-tax income minus taxes τ . Household wealth at the end of the period, a , evolves according to $a = a_{-1} + s + q$, where q denotes capital gains.

In the remainder of the paper, we first describe the datasets we use and our sample selection. Next, we proceed to construct several moments describing the distribution of these different income measures across individuals and households, and show how they evolve over time.

3 Datasets

Our empirical methodology follows closely [Heathcote et al. \(2010\)](#) and [Heathcote et al. \(2020\)](#). Thus, the findings in this paper can be compared to our earlier work.

3.1 CPS

The CPS is the source of official U.S. government statistics on labor force status, and is designed to be representative of the civilian non-institutional population. The Annual Social and Economic Supplement (ASEC) applies to the sample surveyed in March, and extends the set of demographic and labor force questions asked every month to include

detailed questions on income. We use these income reports for all components of the household budget constraint except for taxes, which we impute using the National Bureau of Economic Research (NBER) TAXSIM model.

The basic unit of observation is a housing unit, so we report CPS statistics on inequality at the level of the household (rather than at the level of the family).¹⁰ We rely on CPS ASEC samples from 1967 to 2021. Our estimates are produced using the March supplement weights.

3.2 PSID

The Panel Study of Income Dynamics (PSID) is a longitudinal study of a sample of U.S. individuals (men, women, and children) and the family units in which they reside. Survey waves were annual from 1968 to 1997, and have been biennial since then.

The PSID was originally designed to study the dynamics of income and poverty. For this purpose, the original 1968 sample was drawn from two independent sub-samples: an over-sample of roughly 2,000 poor families selected from the Survey of Economic Opportunities (SEO), and a nationally-representative sample of roughly 3,000 families designed by the Survey Research Center (SRC) at University of Michigan. In 1996-97 a new sample of immigrant families was added.

We use all these three samples from survey years 1968 to 2019. Throughout the analysis, we employ the longitudinal family weights, which are designed to reduce the biases from differential sampling and response rates. Detailed income information is available in every survey year. Wealth information was collected only from 1984 and was collected at lower frequency before 1997. In 2005, the PSID began to gather information on comprehensive household expenditures, which we will use to compute measures of aggregate consumption and consumption inequality.¹¹

3.3 ACS

The American Community Survey (ACS) is a survey conducted every month and every year (from 2000 onward) by the U.S. Census Bureau, which sends surveys to about 3.5 millions households every year. The survey lacks the level of economic detail of other

¹⁰A “household” is defined as all persons, related or unrelated, living together in a dwelling unit. The “family unit” is defined as all persons living together who are related usually by blood, marriage, or adoption. For example, a household can be composed of more than one family.

¹¹See [Arellano et al. \(2017\)](#) or [Heathcote and Perri \(2018\)](#), among others, for studies which use these recent PSID consumption data to assess household consumption response to income shocks.

surveys such as the CPS, but it has a very large sample size. We use ACS samples from 2000 to 2021.

3.4 CES

The Consumer Expenditure Survey (CES) consists of two separate surveys, the quarterly Interview Survey and the Diary Survey, both collected for the Bureau of Labor Statistics by the Census Bureau. It is the only dataset that provides detailed information about household consumption expenditures in the U.S. (with the exception of the PSID over the period 2005–2019). The Diary Survey focuses on expenditures on small, frequently-purchased items (such as food, beverages and personal care items), while the Interview Survey aims at providing information on up to 95 percent of the typical household’s consumption expenditures. In this study, we make use of only the Interview Survey.

The Interview Survey is a rotating panel of households that are selected to be representative of the U.S. population. It started in 1960, but continuous data are available only from the first quarter of 1980, which is the start of our sample. In each quarter, the survey reports, for each household interviewed, detailed demographic characteristics for all household members, detailed information on consumption expenditures for the three month period preceding the interview, and information on income and hours worked over a yearly period. Each household is interviewed for a maximum of four consecutive quarters. Expenditure information is reported every quarter, income information is reported only in the first and last quarter, and wealth information is reported only in the last quarter.¹² We use CES samples from 1980 until 2022, but when looking at moments involving consumption we focus on the post 1994 period.¹³

3.5 SCF

The SCF is a triennial survey of U.S. households managed by the Board of Governors of the Federal Reserve System. The survey collects information on income (for the year preceding the survey), but focuses primarily on detailed information about household financial and non-financial assets, liabilities, and capital gains. The SCF survey has two parts: a random sample of U.S. households, and a second sample of wealthy households, identified on the basis of tax returns. The SCF provides weights for combining the two

¹²Unlike [Krueger and Perri \(2006\)](#) and [Aguiar and Bils \(2015\)](#), we do not further require that a household participate in all four interviews in order to be present for the income variables in the final interview.

¹³The definitions for some consumption components prior to 1994 differ from those for the later period.

samples. For more details about the SCF, see [Kuhn et al. \(2020\)](#). We use SCF samples from survey years 1989 to 2019.

3.6 Sample selection

In each of our five datasets, we construct three different samples, labeled Samples A, B, and C.

Sample A is the most inclusive, and is essentially a cleaned version of the raw data. We drop only records if 1) there is no information on age for either the head or the spouse, 2) either the head or spouse has positive labor income but zero annual hours (zero weeks worked in the CPS), or 3) either the head or spouse has an hourly wage less than half the corresponding Federal minimum wage in that year. In the CES, we also drop households reporting implausible consumption expenditures.¹⁴ In order to reduce measurement error in income and hours, we also exclude CES households flagged as “incomplete income reporters” ([Nelson 1994](#)), and we exclude PSID households if labor income is missing, but hours worked are positive. Sample A is designed to be representative of the entire U.S. population, and is used to compare per-capita means from micro-data to NIPA aggregates.

Sample B is further restricted by dropping a household from Sample A if no household member is of working age, which we define as between the ages of 25 and 60 (in the PSID we drop households if neither the head nor the spouse, when present, falls in this age range). Sample B is our household-level sample.

Sample C instead is an individual-level sample. To construct it, we first select all individuals aged 25-60 who belong to households in Sample B. From this group we then select those who work at least 260 hours in the year.

Throughout the paper, the price deflator used to construct real variable is the Personal Consumption Expenditure (PCE) deflator, and our base year is 2012.

Top coding affects very few observations in the PSID, but is a more serious concern in the CPS, in the ACS, and in the CES. We describe our procedures for addressing top coding in [Appendix A](#).

¹⁴Specifically, when quarterly equivalized food consumption is below \$100 in 2000 dollars. In the PSID, we categorize records as implausible when either (i) equivalized food consumption is below \$400 per year, (ii) food stamps exceed \$50,000, or (iii) food expenditures exceed ten times disposable income. In such cases, we drop households, but only when computing moments involving food consumption.

3.7 Summary

Table 1 summarizes basic demographic characteristics of households in Sample A for all the data sets.

Table 1: Sample A, Summary Statistics, 2018

	CPS	ACS	PSID	CES	SCF
Number of households	66,929	1,215,264	8,422	14,793	5,813
Average household size	2.44	2.45	2.14	2.21	2.44
Wages/salaries below \$10,000 (%)	30.5	30.4	34.2	38.9	31.8
Wages/salaries above \$200,000 (%)	4.6	4.4	3.9	4.3	4.9
Average age	51.7	52.3	53.9	52.6	50.4
White (non-hispanic) (%)	79.6	77.2	79.5	82.9	66.6
College (%)	37.7	35.6	37.9	37.6	37.5
Female (%)	52.8	53.1	53.4	53.0	53.2

Note: All statistics are weighted, except for the number of households. Dollars are in 2012 prices. The statistics for “Average age” and percentages “White,” “College,” and “Female” refer to the set of individuals who are either heads of household or the spouses of heads.

In terms of demographics, the CPS and the ACS are very similar. The PSID and the CES have smaller and older households, while the SCF has fewer households with a White head or spouse. The fraction of households with a college-educated head or spouse and the fraction of female head or spouse is similar across all datasets.

4 Macro Facts in Micro Data

The feedback between distributions and aggregates is central in much recent macroeconomic research. Micro data that, once aggregated, correctly reproduce macroeconomic time series are essential inputs for this agenda. In this section, we compare several aggregates constructed from our five datasets with the corresponding NIPA series. We start with income measures, and then move to consumption and wealth. Details on how we produce each figure are in [Appendix A](#).

4.1 Income

Wages and Salaries In [Figure 1](#) we report per capita wages and salaries from the NIPA, together with the respective series that result from aggregating household level records

from our five datasets (Sample A).¹⁵ Note, first, that aggregate wages and salaries from the CPS (the solid blue lines) line up closely with the corresponding NIPA series (dotted black), both in terms of levels, long run trends, and cyclical fluctuations. This suggests that for researchers interested in aggregate dynamics and distributional issues related to wage and salary income, the CPS is a reliable reference.¹⁶ The figure also shows that the SCF and the ACS agree quite closely with the NIPA. The PSID series displays a higher level of labor income (relative to that of the NIPA and the CPS) and exhibits slower growth in labor income for the most recent years of the sample.¹⁷ It turns out that, mechanically, the higher level of wages and salaries per capita in the PSID is due to the PSID having a smaller average household size than the other datasets (as noted in Table 1 above).¹⁸ Finally, note that the CES displays a lower level of per capita wages and salaries, relative to the other datasets. As noted in Table 1 above, this is partly explained by the CES having a much higher fraction of households with labor income below 10000\$.

Pre-tax income Next, we turn to pre-tax income. Now, the issue of conceptual comparability between the NIPA and surveys becomes more important. We therefore consider two different definitions of pre-tax income in the NIPA. The first, which we label NIPA-, includes only the additional income categories in the NIPA (besides wages and salaries) that are also recorded in surveys, namely, business income, capital income, FICA contributions, and a range of public transfers, but not Medicare or Medicaid. The second, which we label NIPA+, is identical to NIPA Personal Income, and includes additional income categories that are not reported in surveys. These categories include Medicaid, Medicare, imputed rents for owner-occupied housing, and employer contributions for employee pension and insurance funds.

In Figure 2 we plot pre-tax per capita income for NIPA-, NIPA+ and for our Sample A in the CPS.¹⁹ The gaps between the CPS and NIPA- series reflect under-reporting for the income categories that are in both the CPS and NIPA-, while the additional differences

¹⁵In Figure 1 and Figure B.2 CES observations in year t report averages computed in survey year $t + 1$. This is because in CES a significant fraction of the income reported by households in survey year t refers to income earned in year $t - 1$

¹⁶In the rest of the paper, we will focus on a broader measure of labor earnings, which include wage and salary income and self-employment income. As we will shortly see, self-employment income tends to be under-reported in household surveys, relative to the national accounts.

¹⁷This fact has also been documented in other studies, like Cynamon and Fazzari (2017) and Lippi and Perri (2023)

¹⁸In Figure B.1, in Appendix B, we show that *per household* wages and salaries in the PSID are close to those in the CPS, but household size is notably lower starting in the 1980s.

¹⁹Figure B.2 in Appendix B plots series of pre-tax income for the other surveys.

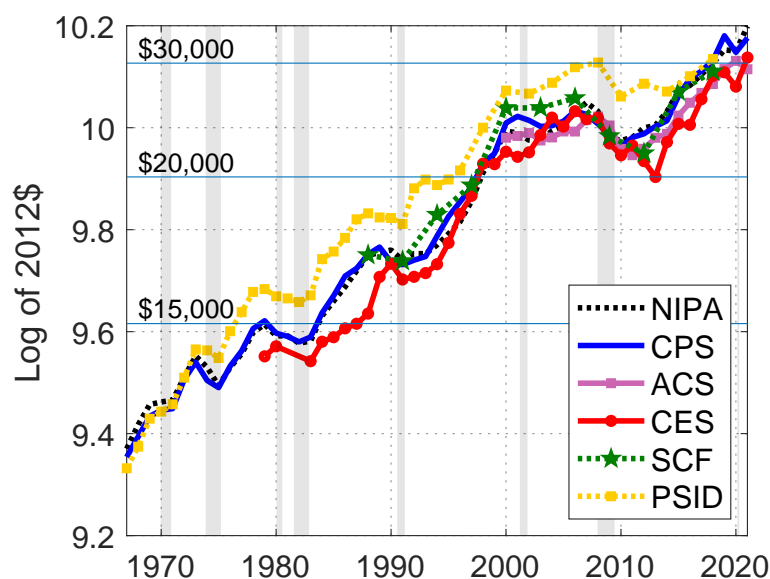


Figure 1: Wage and salary income per capita
NIPA, CPS, ACS, CES, SCF and PSID. Sample A. NBER recessions are shaded in gray.

between the CPS and NIPA+ series reflect the fact that NIPA+ includes some income components that are not reported in the CPS. Focusing first on the comparison between the CPS and NIPA-, we note that the two series are similar in terms of trends and cyclical fluctuations, but CPS income is about 20 percent lower than its NIPA counterpart, and this gap grows slightly over time. Turning now to NIPA+, the gap with the CPS widens even more markedly over the past two decades, indicating the rising importance of income components that are missing from the CPS. For example, in 2000 the gap was about 25 percent, while by 2021, it had risen to about 35 percent.

In Figure 3, we analyze in more detail the discrepancies in pre-tax income between the NIPA+, NIPA-, and CPS measures. For three components of income – self-employment income, capital income and transfers – the figure plots the ratio between the CPS measure and the corresponding NIPA+ and NIPA- measures. There are quantitatively relevant gaps in all the three components. CPS transfers were about 15 percent smaller than NIPA-transfers in the early part of the sample, with the shortfall rising to about 25-30 percent in the later part of the sample, and growing especially large in recessions.²⁰ The gap between CPS and NIPA+ transfers is larger and grows even faster, reflecting the large and rapidly growing size of Medicare and Medicaid. These two lines suggest that under-reporting

²⁰For this comparison, COVID era Economic Impact Payments are included in both CPS and NIPA transfers.

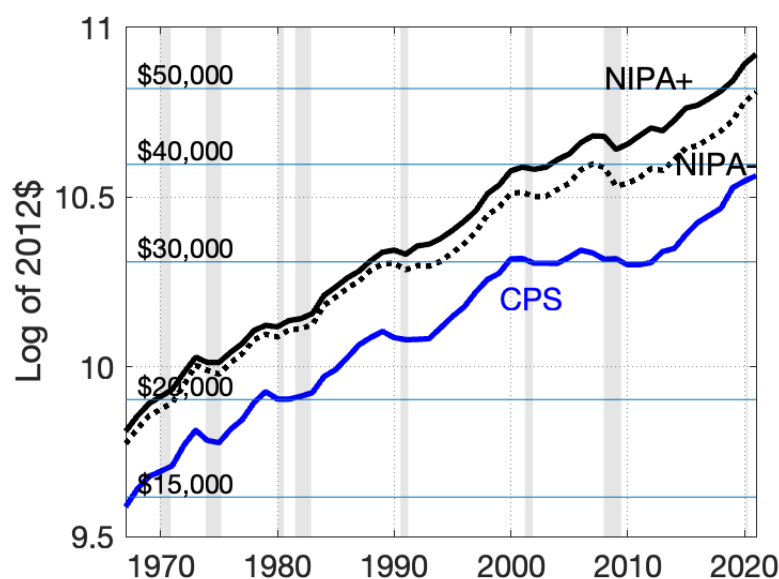


Figure 2: Pre-tax income per capita
NIPA+, NIPA- and CPS. Sample A. NBER recessions are shaded in gray.

of transfers might lead one to over-estimate inequality in survey data (as argued, for example, by [Gramm et al., 2022](#)).²¹

The figure illustrates that there is an even larger discrepancy between the measures of self-employment and capital income in the CPS and the NIPA. This missing income reflects a mix of (i) the CPS under representing wealthy households that account for a large share of income in these categories, and (ii) survey respondents under-reporting these income components. In Section 6.4, we will use these gaps to construct imputed measures of this missing income based on a procedure similar to the one in [Piketty et al., 2017](#), and we will show how our inequality measures are affected.

4.2 Consumption expenditure

Previous work that has used survey data from the CES to study consumption inequality (see, among others, [Krueger and Perri, 2006](#); [Blundell et al., 2008](#); [Aguiar and Bils, 2015](#)) has highlighted that aggregate expenditures from the CES do not line up well with corresponding NIPA figures. In this section, we update this comparison to 2021, and we also

²¹One survey that is explicitly designed to capture transfer payments is the Survey of Income and Program Participation (SIPP). See [Ben-Shalom et al. \(2012\)](#) and [Guner et al. \(2023\)](#) for SIPP-based estimates of the size of transfers to low income households, and their impact on poverty.

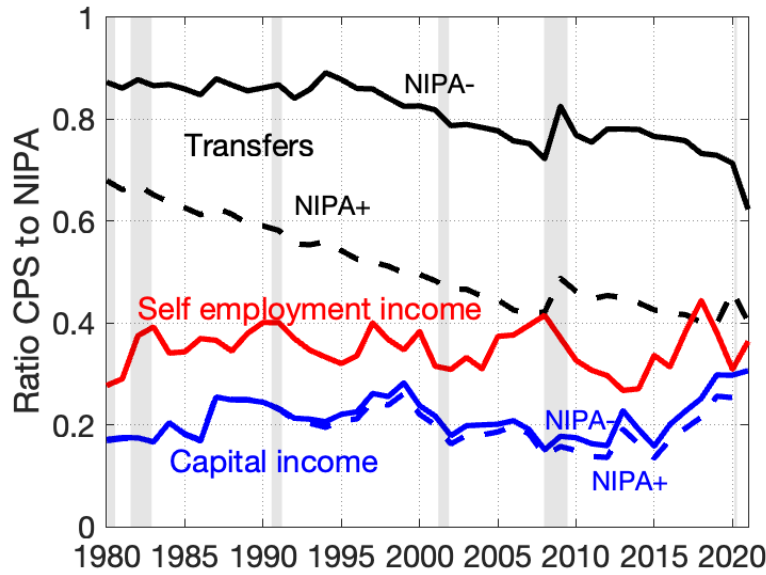


Figure 3: Income components in CPS relative to NIPA counterparts
Sample A. NBER recessions are shaded in gray.

integrate comprehensive aggregate consumption expenditures from the PSID.²² Figure 4 reports aggregate consumption expenditures on non-durable and durable categories from the CES, the PSID and the NIPA.²³ The left panel of the figure shows that in terms of non-durable expenditure, the CES still displays a considerable gap relative to the NIPA (and to the PSID, which is much closer to the NIPA). However, the growth in CES expenditures over the most recent period (2004-2018) tracks the growth in the NIPA more closely than in earlier years, and more closely than the PSID.²⁴ The right panel shows that for durable consumption expenditures, the match between the NIPA and the CES is excellent, both in terms of levels and in terms of cyclical fluctuations, while the PSID seems to overstate the level of durables spending.

²²From 2004 onward, the PSID has collected data on consumption expenditures that are comparable with those in the CES.

²³The aggregate comprises categories of consumption expenditure that are comparable across data sets (see the Appendix for details).

²⁴Over the period 1994-2004, NIPA consumption per capita grew at 1.7% per year, while CES consumption grew at only 0.4% per year. Over the period 2004-2018, NIPA and CES consumption both grew at around 1% per year, while the PSID displays no growth.

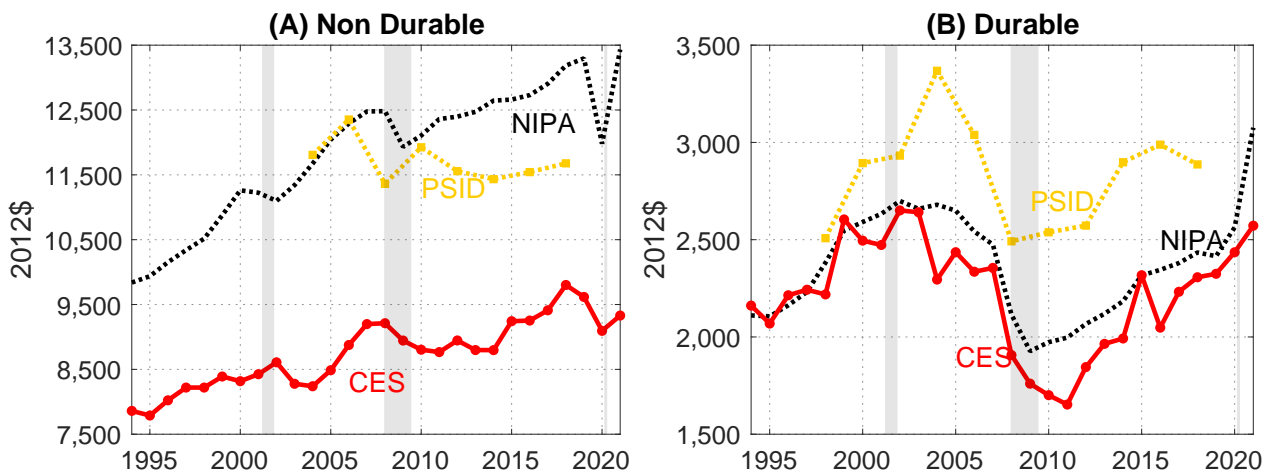


Figure 4: Consumption expenditure per capita

CES, PSID and NIPA. Sample A. (A): Non-durable goods and services. (B): Durables. NBER recessions are shaded in gray.

4.3 Wealth

Figure 5 compares measures of household net worth in the three surveys that report comprehensive measures of net worth (the SCF, the PSID and the CES) with a comparable measure from the Flow of Funds (FOF).²⁵ As is well known, average wealth in the SCF is the closest to the FOF. Wealth in the PSID is slightly lower, but the PSID replicates closely the trend and cyclical fluctuations observed in the FOF. Finally, wealth in the CES is much lower than in all the other surveys, suggesting that the CES misses a significant portion of household wealth. All four measures of wealth show a marked hump corresponding to the house price boom and bust of 2000-2008.

5 Individual-level inequality

We start our analysis of inequality in the United States by studying the distribution of individual labor income. We begin with individual hourly wages, and zoom in on different points in the distribution. We also distinguish between the role of observable demographics versus residuals. Next, we turn our attention to earnings, the product of hourly wages and hours worked, and argue that labor supply dynamics are key to understanding the divergence between wage and earnings dispersion in the bottom half of the distribution.

²⁵This data is produced by the Federal Reserve Board of Governors as part of the Financial Accounts of the United States.

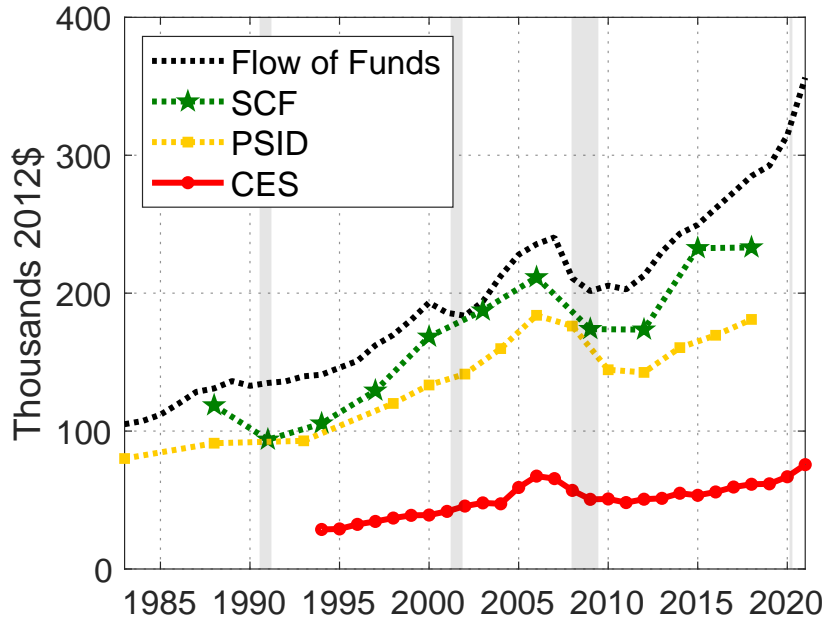


Figure 5: Household wealth per capita
Flow of Funds, SCF, PSID and CES. Sample A. NBER recessions are shaded in gray.

5.1 Wages

In a competitive labor market, the hourly wage equals the marginal product of an additional hour worked. As such, it reflects both the price of skills and the quantity of efficiency units produced by an hour of work. In the presence of search frictions or monopsony power, hourly wages deviate from the marginal product of labor and also contain a component of rent accruing either to the worker, or to the firm, or to both.

Our empirical measure of hourly wages is constructed for every individual in Sample C in the CPS as annual earnings divided by annual hours worked.

Figure 6 displays two standard measures of dispersion in hourly wages by gender: the variance of log wages and the Gini coefficient. For men, both measures of inequality increase quite steadily until the early 2000s, and then decelerate. For women, the rise in wage dispersion starts later, around 1980, and also slows down in the 2000s. Since 1980, the evolution of wage inequality has been strikingly similar across genders. However, the level of wage inequality for women is systematically lower than for men.

Quantitatively, the overall rise in U.S. wage inequality since the late 1960s is substantial. The variance of log wages rises by over 20 log points, and the Gini by over 10 points.

These two statistics capture dispersion across the entire population and, as a result,

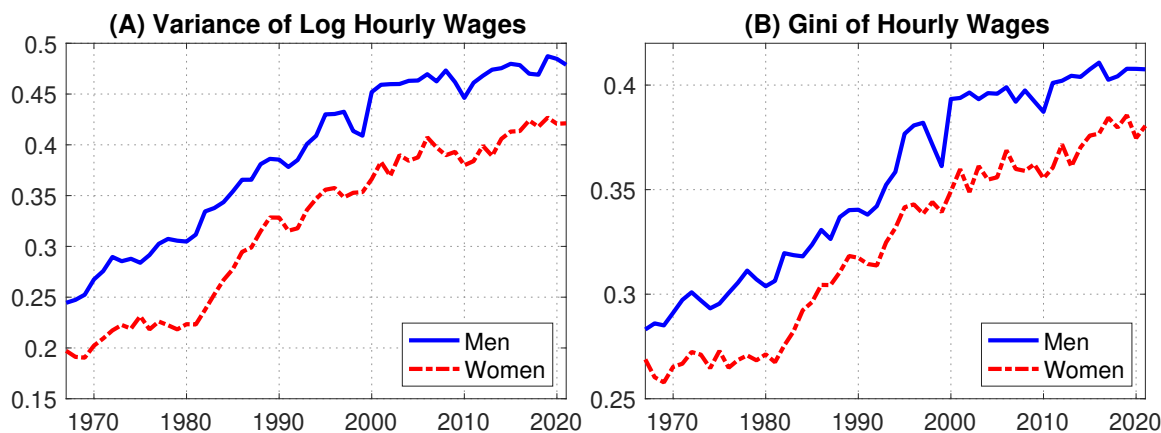


Figure 6: Inequality in hourly wages
CPS, Sample C. (A): Variance of log hourly wages. (B): Gini coefficient for hourly wages

they can hide diverging dynamics of inequality in different parts of the distribution. Figure 7 inspects the evolution of wage inequality in the top and bottom halves of the distribution. Our preferred measures are the top 10% to middle 10% ratio and the middle 10% to bottom 10% ratio. The former is the ratio between the average wage in the top 10% of the distribution (i.e., all individuals between the P90 and the P100) and the average wage in the middle 10% of the distribution (i.e., all individuals between the P45 and the P55). The middle to bottom ratio is defined analogously.

We prefer these measures to standard percentile ratios primarily because they will later allow us to jointly examine the dynamics of earnings, weekly wages, and weeks worked. Having said that, we note that Panel (B) in Figure 7 can be used to compare the top to middle ratio to the P95-P50 and P98-P50 ratios, and the middle to bottom ratio to the P50-P5 ratio.

Panels (A) and (C) show that the rise in inequality above the median was quite modest until the mid 1980s, when it accelerated significantly. For example, from 1985 to 2020, the gap between wages of the top 10% and the middle 10% surged by nearly 50% for both men and women. The comparison between our preferred measures of inequality (Panel A) and percentile ratios (Panel B) suggests that the vast majority of the increase in wages at the top, responsible for the rise in the top to middle ratio, has occurred above the 95th percentile of the wage distribution.

Panel (A) shows that in the bottom half of the population distribution, wage inequality as measured by the middle to bottom ratio was surprisingly stable over the entire sample period. But Panel (D), which splits the sample by gender, illustrates that the absence of an

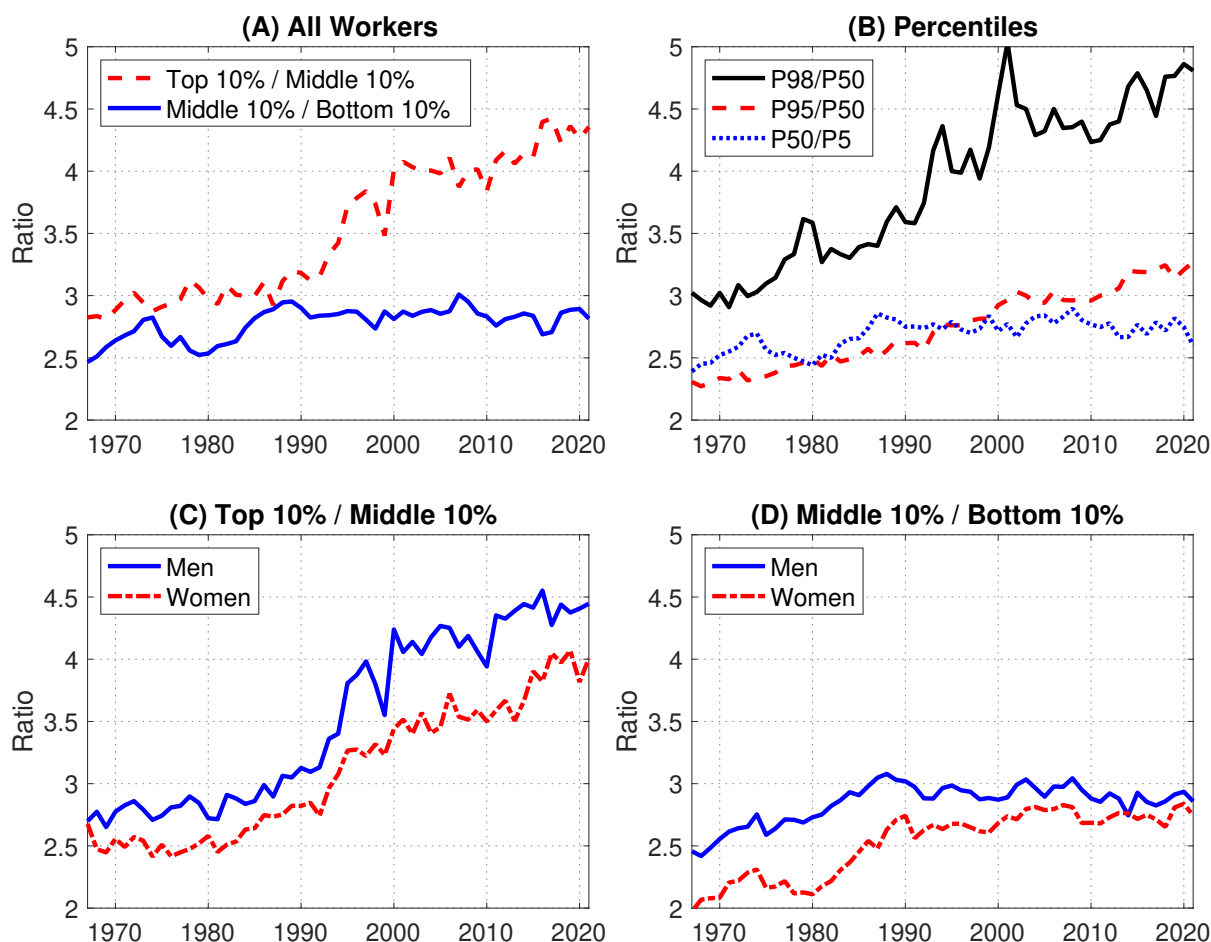


Figure 7: Wage inequality in different parts of the distribution

CPS, Sample C. Individuals ranked by hourly wage. (A): Ratio of average wage of top 10% to average wage of middle 10% and ratio of average of middle 10% to average of bottom 10%. (B): Ratio of 98th percentile wage to median wage, ratio of 95th percentile to median, and ratio of median to 5th percentile. (C): Top to middle ratios by gender. (D): Middle to bottom ratios by gender.

upward trend until 1990 is the result of a composition effect. The middle to bottom ratio rises for both men and women over this period, but rising female participation –implying progressively more weight on women in the sample population– combined with a lower level of wage dispersion for women offsets the upward trend in gender-specific inequality when analyzing the entire sample. Since the late 1980s, however, the middle to bottom ratio has been extremely stable for both women and men.²⁶

²⁶The dynamics of the minimum wage do not seem to be an important determinant of wages at the bottom of the distribution in the aggregate: the correlation between the bottom percentile (P5, P10, and P20) values of the hourly wage distribution and the real federal minimum wage (FRED series FEDMINNFRWG divided by our price index) is close to zero.

5.2 Observables and residuals

In order to understand the sources of the rise in U.S. wage inequality, it is useful to distinguish the roles of some key observable demographics: education, age, gender, race and occupation type. We perform this decomposition in Figure 8.

We define the education premium as the ratio between the average hourly wage of workers with at least 16 years of schooling, and the average wage of workers with less than 16 years of schooling. As has been extensively documented, the college wage premium declines until the late 1970s for both genders, and then starts rising steadily until the mid-2000s (Panel A). Since then, the college premium has remained flat. To put this rise in perspective, consider that in 1980, U.S. male college graduates earned 45% more than high school graduates, while by the early 2000s they earned roughly 90% more.²⁷

The experience (age) wage premium (Panel B) is defined as the ratio between the average hourly wage of 45-to-55-year-olds and the hourly wage of 25-to-35-year-olds. The experience premium for both men and women rises steadily from the mid 1970s to the mid 1990s, with an increase of at least 20% for men and 10% for women. After 1995, it flattens out.

The plot of the gender wage premium, defined as the ratio of the average hourly wage of men to the average hourly wage of women, (in Panel C) shows that, on average, men earned over 40% more per hour than women in 1975, but only 20% more in 2020. Convergence was especially strong in the 1980s, slowed down in the 1990s, and has picked up again from around 2000.

Panel (D) of Figure 8 plots the hourly wage race gap between White and Black workers. Wages of Black men have converged toward those of White men at a very slow pace over the past 50 years, with the most visible gains occurring in the early-to-mid-1970s. Today, the average gap is stable at around 25%. For women, after some convergence in the early 1970s, the race wage gap has widened. In 1980 Black working women earned only 7% less than White women, but for most of the remaining sample period, this gap widened, reaching 15%. Interestingly, the very tight labor market of the last two years has delivered some strong relative wage gains for Black women, but it is too early to assess whether these gains will persist.

The bottom two panels plot occupational premia. Our categorization of occupations into different groups, based on the type of task performed by workers employed in those

²⁷The ratio between hourly wage of workers with an advanced degree (master's or higher) and that of workers with a bachelor's degree follows a pattern similar to the one for the college premium: it rises until the early 2000s and then it levels off at around 38% for men and 32% for women.

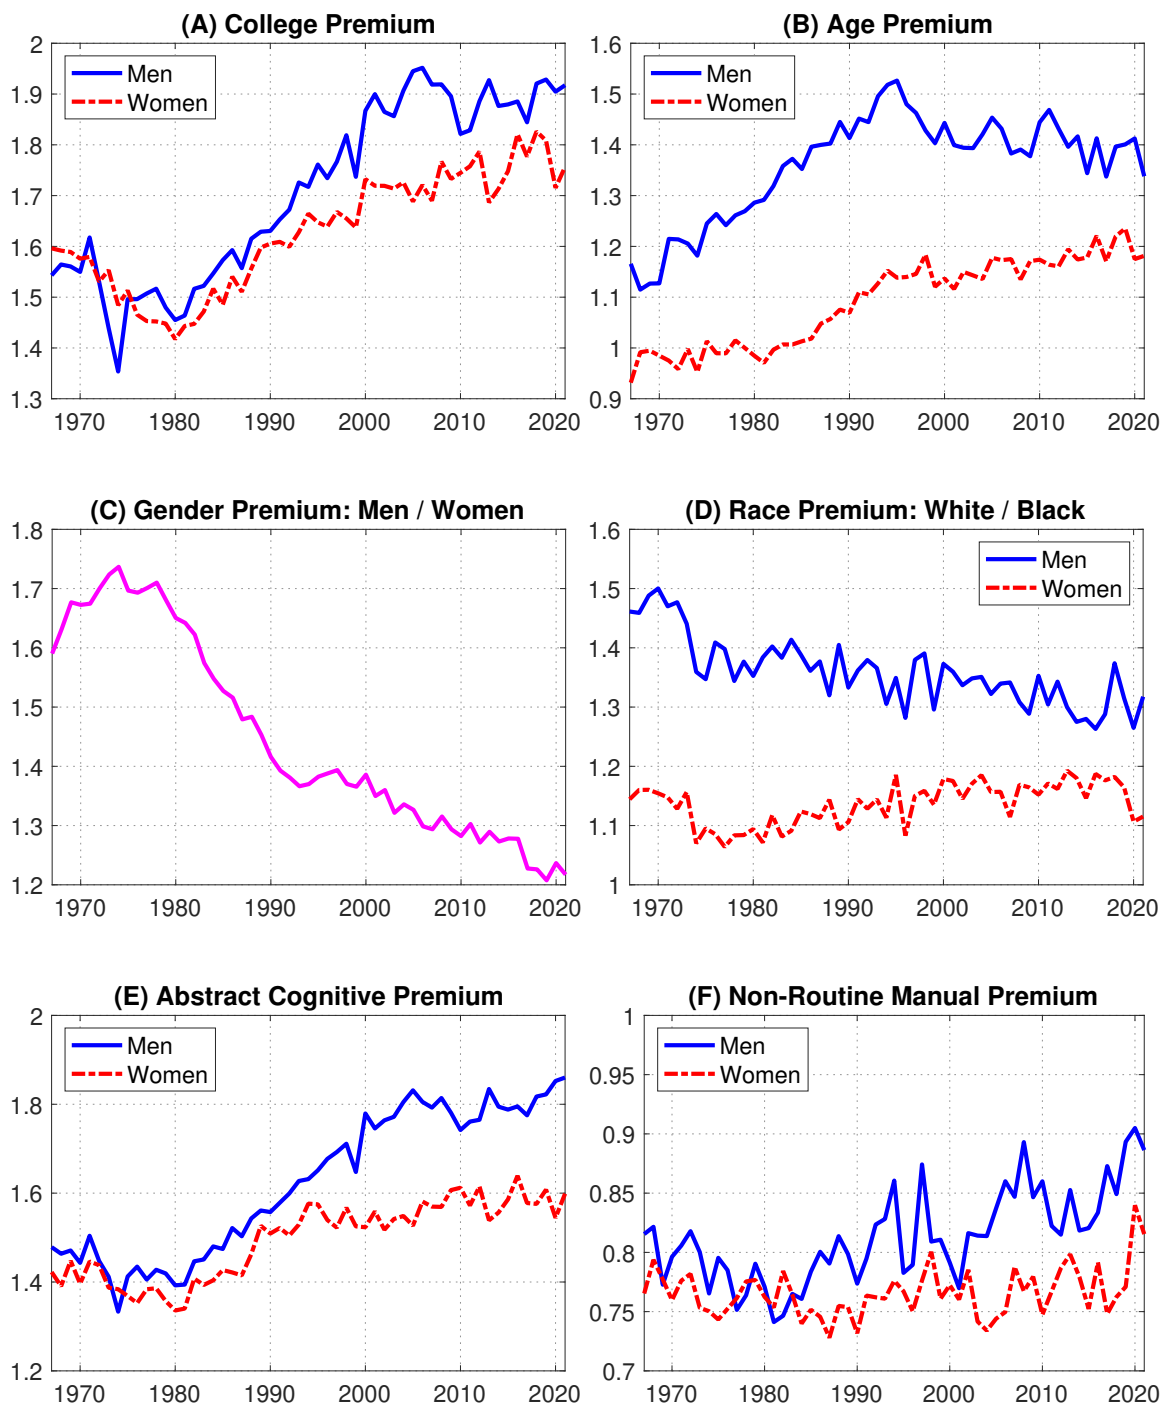


Figure 8: Wage premia based on observable characteristics

CPS, Sample C. (A): Ratio of average wage of workers with 16+ years of schooling to average of workers with less than 16 years. (B): Ratio of average wage of 45-to-55 year-olds to average of 25-to-35 year-olds. (C): Ratio of average wage of men to average wage of women. (D): Ratio of average wage of Whites to average wage of Blacks. (E): Ratio of average wage in Abstract Cognitive occupations to average in Routine Cognitive and Routine Manual (RCM) occupations. (F): Ratio of average wage in Non-Routine Manual occupations to average wage in RCM occupations.

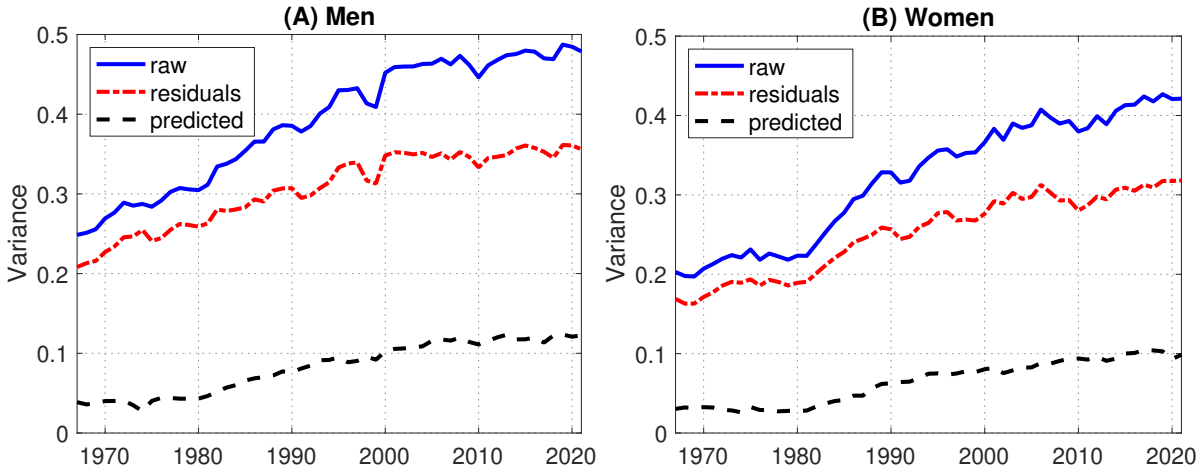


Figure 9: Variances of raw, predicted, and residual log hourly wages CPS, Sample C. (A): Men. (B): Women. Residuals are computed from standard log wage regressions, run separately for each year and for each gender, with controls for education, occupational classification, age, and race.

jobs, follows the classification of [Acemoglu and Autor \(2011\)](#). The abstract cognitive (AC) group includes Professional, Managerial, and Technical occupations. The routine cognitive and routine manual (RCM) group includes Clerical and Sales occupations as well as Production occupations and Operators. The non-routine manual (NRM) group comprises all Service occupations. Panels (E) and (F) illustrate the polarization in the wage distribution that occurred in the space of tasks. Since 1980, wages of both AC jobs (the highest paid ones) and NRM jobs (the lowest paid ones) have increased relative to wages of routine occupations, which are concentrated in the middle of the wage distribution.²⁸ We note that while this phenomenon is very apparent for men, it is much less visible for women.

Figure 9 displays raw, predicted and residual log wage inequality for men and women. The predicted component is computed by running a year-by-year regression of log wages on standard demographics, while residuals are the difference between raw log wages and their predicted component.²⁹ The figure shows that the increase in inequality until 1980 is entirely due to the residual component. Since then, the predicted component (mostly college and task premia) accounts for roughly 1/3 of the rise in the variance of log wages.

²⁸See [Cortes et al. \(2017\)](#) for an analysis of the causes of the falling wages of these routine jobs

²⁹The regression includes two dummies for education (college degree and above, less than college), three occupation types (AC, RCM, and NRM), three for race (White, Black, Other), and an age polynomial (age and age squared).

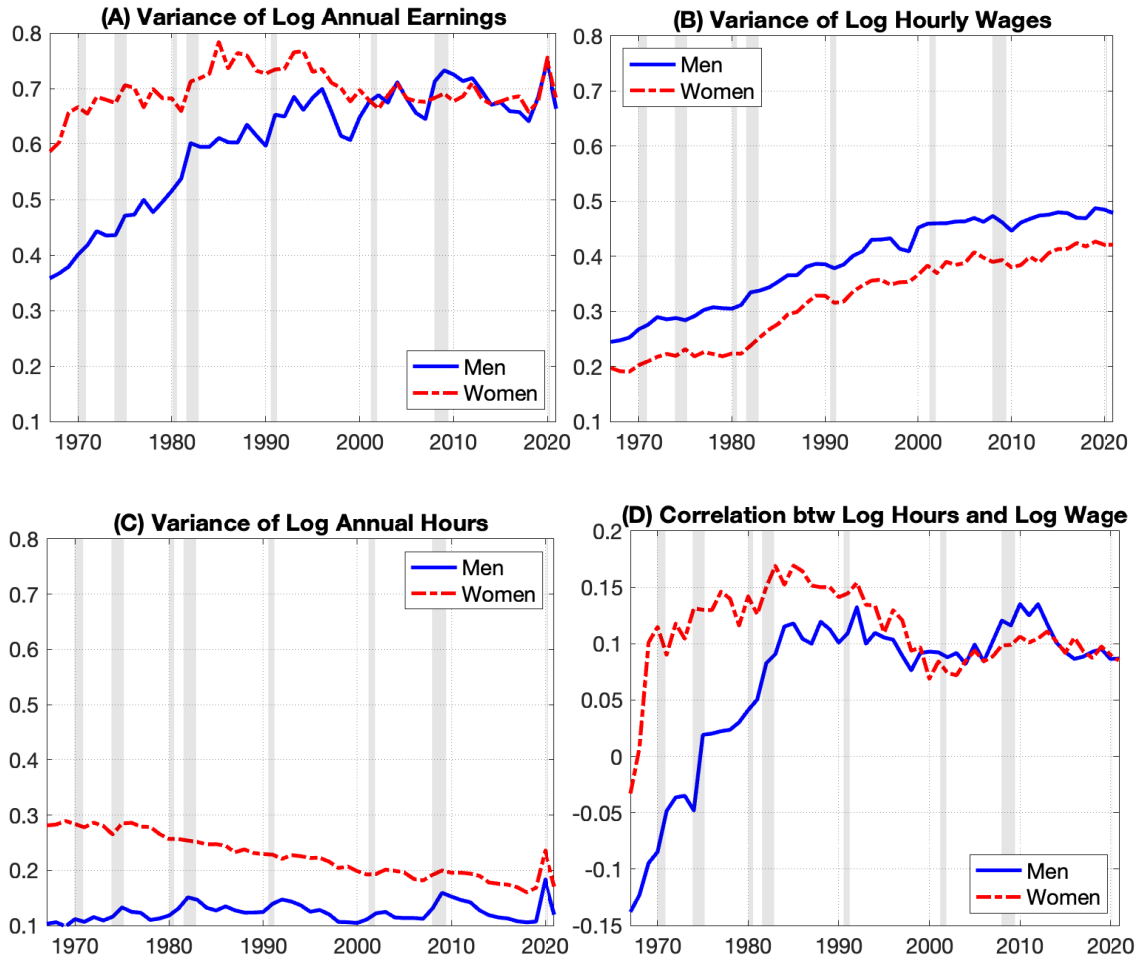


Figure 10: Decomposition of the variance of log earnings
 CPS, Sample C. (A): Variance of log annual earnings. (B): Variance of log hourly wages. (C): Variance of log annual hours worked. (D): Correlation between log hours worked and log hourly wages. NBER recessions are shaded in gray. Note that $A = B + C + 2D\sqrt{BC}$

5.3 Earnings

We now turn to the analysis of individual earnings inequality. Individual earnings are the product of hourly wages and annual hours worked. The dynamics of earnings inequality can, therefore, differ from those of wage inequality because of changes in the distribution of hours worked and in the correlation between wages and hours.

Figure 10 decomposes the evolution of the variance of log earnings into the variance of log wages, the variance of log hours, and the correlation between wages and hours. For men, the variance of log earnings displays a very sharp rise until 2000, and a markedly slower increase afterwards. The surge in the first three decades is due to a combination

of expanding hourly wage inequality, an increasing wage-hours correlation, and flat dispersion of hours worked. For women, the variance of earnings grows until the late 1990s, but more modestly compared to that of men, and then declines. The initial rise is the product of widening wage dispersion and a rising correlation between wages and hours, which largely offset the decline in the variance of hours due to increasing female labor force attachment. The decline in the variance of female earnings in the second half of the sample reflects the fact that the increase in wage inequality slows down and the wage-hours correlation flattens, but dispersion in hours worked keeps decreasing at the same pace.

Figure 10 suggests that inequality in individual earnings has stabilized since the late 1990s. However, it is important to remember that this plot is constructed using Sample C, which drops individuals with zero or low hours. We now turn to Sample B, which does not impose any lower bound on hours worked. Figure 11 decomposes annual earnings into weekly wages versus weeks worked for men in Sample B. A comparison of Panels (A) and (B) shows that the stark increase in earnings inequality in the top half of the earnings distribution throughout the sample period is entirely driven by differential growth in real wages at the top vis-a-vis the middle. Top wages doubled over the past half century, whereas middle wages were essentially flat. Weeks worked at the top and in the middle are very stable over the entire period, at around 50 weeks per year for both groups.

The story at the bottom is quite different. A comparison of Panels (B) and (C) shows that earnings inequality for men at the bottom of the distribution rises throughout the sample period, in contrast to the message suggested by Panel (A) of Figure 10 (which excludes the zero earnings men). This increase in earnings inequality for men at the bottom is driven mostly by a decline in weeks worked by low-earnings workers. Over this period, wages at the bottom of the distribution fell 20%, with all the decline occurring in the 1970s. Weeks worked, instead, declined by 80%, displaying a very pro-cyclical pattern.³⁰ This decline in hours worked and wages for low earnings men explains the rise in the wage-hours correlation shown in Panel (D) of Figure 10. Note that wages and hours both decline until the late 1980s, whereas since then wages have remained stable while hours

³⁰In [Heathcote et al. \(2020\)](#), we decompose average weeks worked by men in the bottom 20% of the earnings distribution as the product of two terms: average weeks worked conditional on working a positive number of weeks, and the fraction of men who work a positive number of weeks. We show that the fraction of men working positive weeks accounts for almost all of the long run decline in weeks worked. Thus, a large part of the increase in inequality at the bottom of the male earnings distribution in the United States over the past 55 years reflects declining weeks worked, which in turn is explained by an increase in the fraction of men out of the labor force.

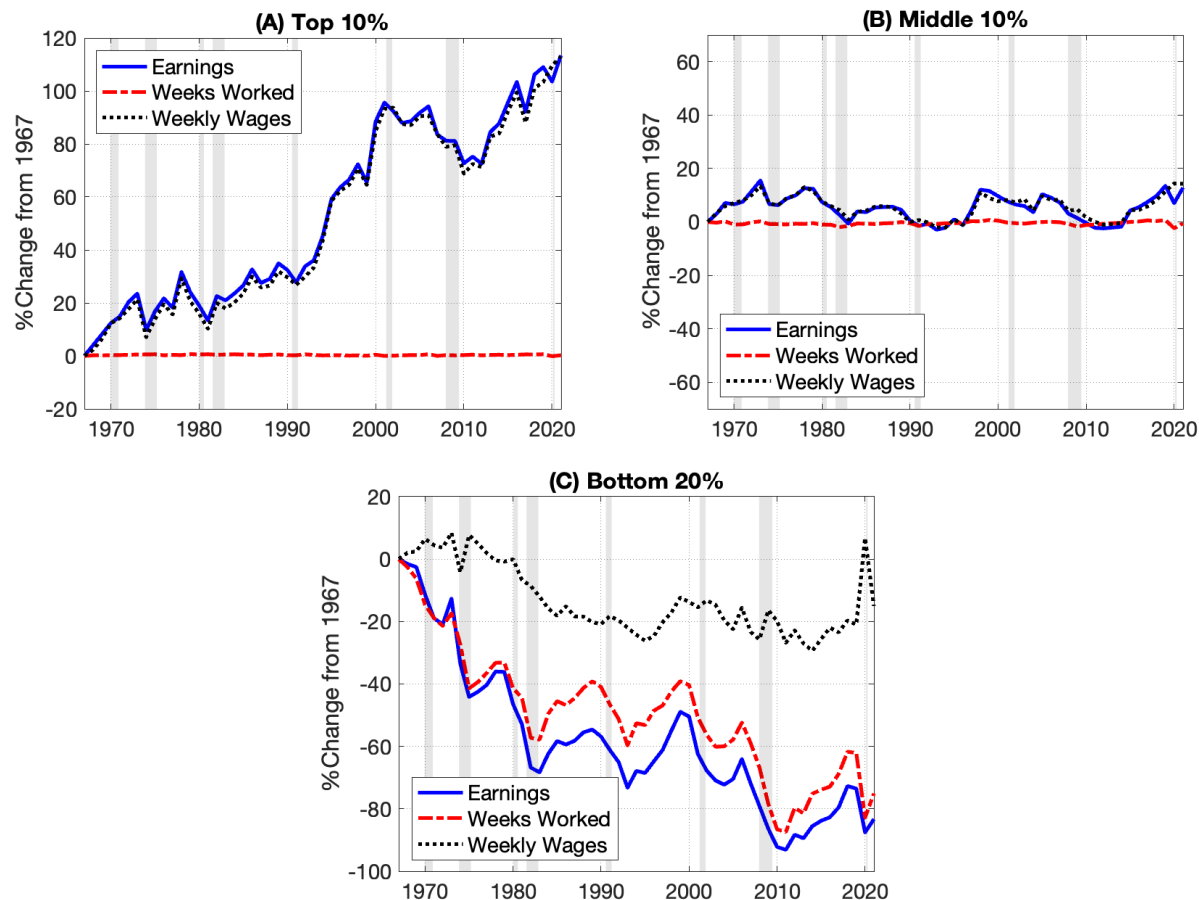


Figure 11: Earnings, wages, and weeks worked in three slices of the earnings distribution CPS, Sample B. Men ranked by annual earnings. (A): Top 10%. (B): Middle 10%. (C): Bottom 20%. NBER recessions are shaded in gray.

have kept falling.³¹

6 Household-level inequality

Thus far, we have focused on inequality at the individual level. But most Americans live in multi-person households, and pool resources within the household. Indeed, data on consumption and wealth are only collected at the household level. What is the impact of

³¹Figure B.3 in the Appendix compares the ratio of earnings of the middle 10% to the bottom 20% for two samples of men. The first is men in our Sample C, which imposes a lower bound on hours worked. The second sample is all working age men, including those with zero annual earnings. Comparing the two lines plotted, it is clear that inequality in the broader sample keeps rising, and that the stabilization in middle-to-bottom earnings inequality in Sample C is an artifact of a declining male employment rate, which implies that over time an increasing number of low-earnings men are selected out of the sample.

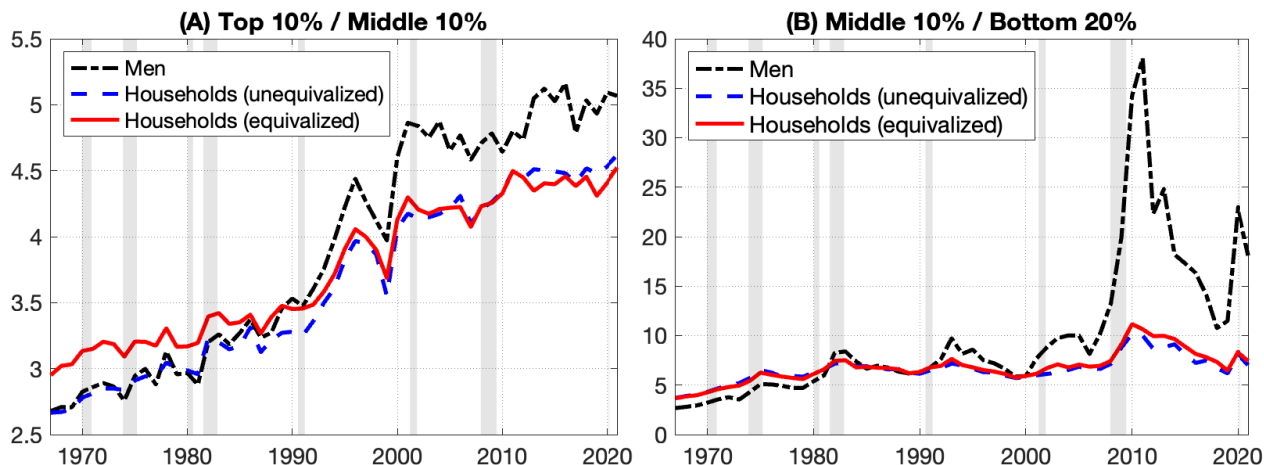


Figure 12: Individual versus household earnings inequality

CPS, Sample B. Male earnings, household earnings, and equivalized household earnings. Inequality measures computed for income-measure-specific distributions. (A): Ratio of average earnings of top 10% to average earnings of the middle 10%. (B): Ratio of average earnings of the middle 10% to average earnings of the bottom 20%. NBER recessions are shaded in gray.

within household income pooling on inequality, and how has the role of the household changed over time?

We focus here on Sample B, which contains all households in which at least one household member is between ages 25 and 60. Recall that Sample B does not impose a lower bound on hours worked, so it includes some households with zero household earnings.

6.1 Household versus individual earnings

Figure 12 compares the dynamics of earnings inequality in the CPS for three different earnings measures. The first is male earnings: we take all men aged 25-60 in Sample B, and rank them by individual labor earnings. The second measure is household earnings: for each household in Sample B, we sum earnings for all individuals in the household. The third measure is equivalized household earnings, which is household earnings divided by an equivalence scale designed to measure of available household resources per person in a way that accounts for possible economies of scale within the household. We use the OECD equivalence scale, which puts weight 1.0 on the first adult in the household, weight 0.7 on each additional adult, and weight 0.5 on each child.³² For each of these earnings measures, we rank households by earnings, and compute average earn-

³²Thus, a single adult household with household earnings of \$100,000 has the same equivalized household income as a married couple household with two children with \$270,000 household earnings.

ings of the top 10 percent, the middle 10 percent, and the bottom 20 percent. The plot shows the top to middle, and middle to bottom earnings ratios.

Comparing inequality in male earnings to household earnings, we note that the dynamics at the top of the respective distributions are similar, though the increase in male earnings dispersion is larger than the increase at the household level. The dynamics at the bottom are more different: there are large spikes in inequality in the male earnings distribution during the Great Recession and the COVID recession that are much less apparent in household earnings. It is also important to note that here, we have not plotted inequality for female earnings. More than 20 percent of working age women have zero earnings, so the ratio of middle to bottom female earnings is infinite. Thus the distribution for female earnings looks very different than the distribution for household earnings.

Comparing household earnings to equivalized household earnings shows that whether or not one equalizes makes very little difference to these measures of inequality. Equalizing does slightly reduce the measured rise in inequality at the top, because it amplifies inequality at the start of the sample period, when middle income households were larger, on average, than top income households.

We now explore a decomposition that allows us to better understand the relationship between inequality in individual earnings and inequality in household earnings, and how the role of the household in shaping inequality has changed over time.

For each working age individual i (both men and women, including individuals with zero annual earnings) belonging to a household in our Sample B, we compute two measures of earnings: individual earnings, y_i , and earnings per working-age adult in the household, $\bar{y}_i = \sum_{i \in H_i} y_i / N_i$, where H_i denotes the set of working-age adults in i 's household and N_i is the cardinality of H_i .

Note that for single adult households, $\bar{y}_i = y_i$. Note also that if we take a sample of working age individuals, the variance of earnings per household adult will always be smaller than the variance of individual earnings. We define a measure of household income pooling as

$$HP_t = \frac{\text{var}(y_{it}) - \text{var}(\bar{y}_{it})}{\text{var}(y_{it})}. \quad (2)$$

Note that $HP_t \in [0, 1]$ measures the share of the variance of income that is reduced by income pooling. At one extreme, if all individuals belonged to single adult households, there would be zero household pooling. There would also be zero pooling if all individuals lived as couples in which each of the two individuals in every couple had identical individual earnings. At the other extreme, if all individuals belonged to one large house-

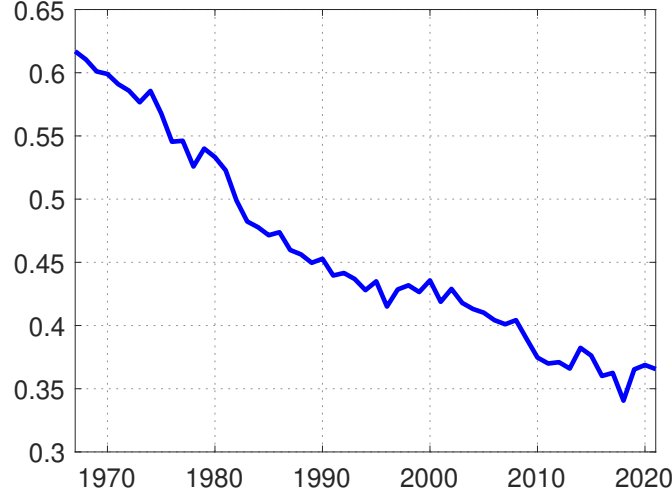


Figure 13: Household income pooling index, HP_t
CPS, Sample B

hold, the household would deliver perfect income pooling, $HP_t = 1$.

Figure 13 plots the extent of household pooling for working age adults in our Sample B. It shows a large decline in the extent of pooling over time. Around the start of our sample, household pooling was reducing the variance of individual earnings inequality by over 60 percent. By the end of our sample, it was reducing earnings inequality by only 35 percent. Why has the household become so much less important as a mechanism for reducing inequality? To understand this, we now decompose our household pooling measure into terms that highlight the impact of various demographic and labor market trends.

To simplify slightly, we restrict our sample to households that contain either one working age adult, or two working age adults, one of which is male and one female.³³ We define household earnings as the total earnings of these adults. Let superscripts s and c denote individuals belonging to single adult and two adult (couple) households, respectively. Let π_t^s denote the fraction of adults in our sample that are in single households at t . By the law of total variance,

$$\text{var}(y_{it}) = \pi_t^s \text{var}(y_{it}^s) + (1 - \pi_t^s) \text{var}(y_{it}^c) + \text{var}_{k \in \{s,c\}}(Y_t^k),$$

where Y_t^k denotes the conditional mean of individual income for individuals in house-

³³Thus, we exclude households with three or more working age adults, and households with two working age adults of the same gender. The excluded households account for around 8 percent of total earnings in Sample B. Note that two-adult households in our sample need not be married (89 percent of them are), and one adult households in our sample need not be unmarried.

holds of type k . A similar decomposition applies for $var(\bar{y}_{it})$. Noting that $var(y_{it}^s) = var(\bar{y}_{it}^s)$ and $Y_t^s = \bar{Y}_t^s$ and $Y_t^c = \bar{Y}_t^c$, the difference in variance between individual and per-household-adult earnings is

$$var(y_{it}) - var(\bar{y}_{it}) = [var(y_{it}^c) - var(\bar{y}_{it}^c)](1 - \pi_t^s).$$

Let superscripts m and f denote males and females. The variance of per-adult earnings for individuals in couples can be decomposed as

$$var(\bar{y}_{it}^c) = var\left(\frac{y_{it}^{c,m} + y_{it}^{c,f}}{2}\right) = \frac{1}{4}var(y_{it}^{c,m}) + \frac{1}{4}var(y_{it}^{c,f}) + \frac{1}{2}cov(y_{it}^{c,m}, y_{it}^{c,f}),$$

where $y_{it}^{c,m}$ ($y_{it}^{c,f}$) denotes the earnings of the man (woman) in the couple to which individual i belongs.

The variance of individual earnings for individuals in couples can be decomposed into the average within-gender variances for coupled men and coupled women, plus the across-gender variance in average wages:

$$var(y_{it}^c) = \frac{1}{2}var(y_{it}^{c,m}) + \frac{1}{2}var(y_{it}^{c,f}) + \frac{1}{2}(Y_t^{c,m} - Y_t^{c,f})^2,$$

where $Y_t^{c,m}$ ($Y_t^{c,f}$) denotes the average earnings of men (women) in couples at date t .

Combining these expressions, we can write

$$\begin{aligned} var(y_{it}) - var(\bar{y}_{it}) &= \underbrace{\frac{1}{4}(var(y_{it}^{c,m}) + var(y_{it}^{c,f}))}_{(1) \text{ within-gender inequality}} - \underbrace{\frac{1}{2}cov(y_{it}^{c,m}, y_{it}^{c,f})}_{(2) \text{ sorting}} \\ &\quad + \underbrace{\frac{1}{2}(Y_t^{c,m} - Y_t^{c,f})^2}_{(3) \text{ gender gap}} - \underbrace{[var(y_{it}^c) - var(\bar{y}_{it}^c)]\pi_t^s}_{(4) \text{ household composition}}. \end{aligned} \quad (3)$$

Each of the four terms on the right-hand side has a straightforward interpretation.

The first (within-gender inequality) is the reduction in earnings variance from household pooling that would obtain if all individuals were in couples, if men and women had the same average earnings, and if there was zero correlation between earnings within the household.³⁴

³⁴In such a world, $HP_t = \frac{1}{4}(var(y_{it}^{c,m}) + var(y_{it}^{c,f}))/var(y_{it})$. Thus, if $var(y_{it}^{c,m}) = var(y_{it}^{c,f}) = var(y_{it})$, the household pooling coefficient would be 0.25.

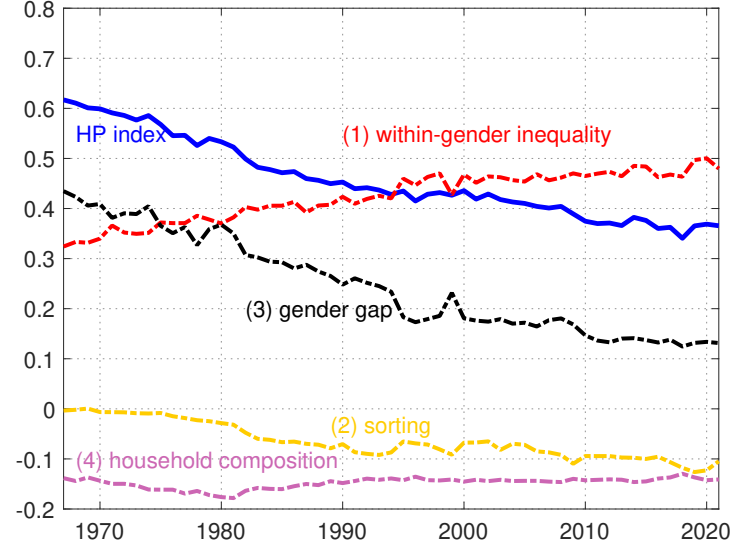


Figure 14: Household income pooling decomposition

CPS, Sample B. Households with one or two (one male and one female) working age adults. Components plotted are defined in equation (3).

The second (sorting) is the extent to which that first pooling mechanism is made smaller (larger) because of positive (negative) within-couple correlation in earnings.

The third (gender gap) is the additional reduction in earnings variance from household pooling that reflects the gender earnings gap.

The final term (household composition) captures the fact that all the above channels apply only to couples, and the household provides no insurance to singles.

Figure 14 plots the evolution of the different terms over time, all relative the variance of individual earnings $var(y_{it})$. The contribution of within-gender inequality is rising over time, reflecting the fact that the variance of earnings for married individuals is rising faster than the variance for all individuals. Thus, the decline in household pooling must be driven by the other terms. The sorting term starts around zero and is generally declining, reflecting the fact that the covariance between male and female earnings within couple households is initially near zero, but rises over time. This increase in correlation in earnings within the household reduces the coefficient of household pooling, but the effect is quantitatively rather small.

The most important driver of reduced household pooling is a narrowing gender earnings gap. In the late 1960s, the gender earnings differential was much larger than it is now, so the pooling of earnings within couples had a very large inequality-reducing effect: many married women would have been very poor without their husbands' income. Over time, average earnings for women have risen much faster than those for men, and

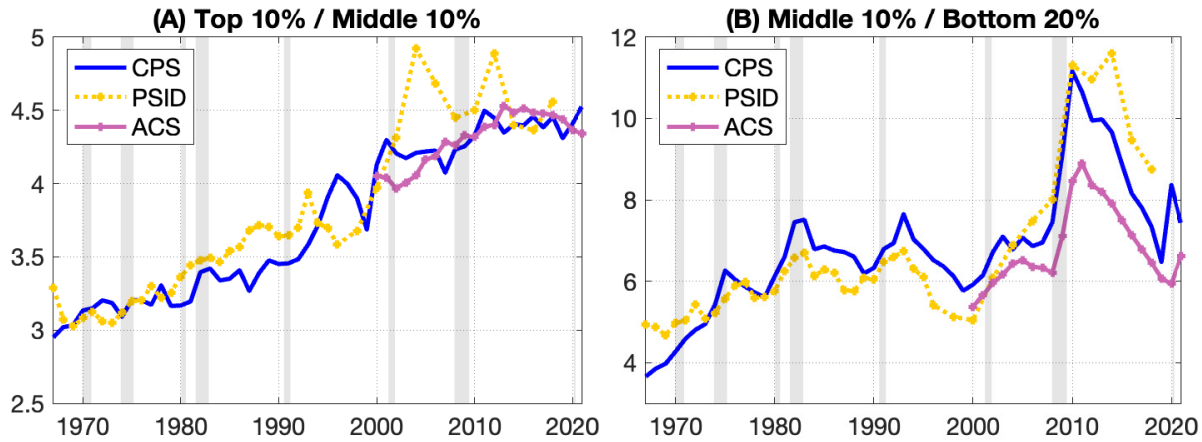


Figure 15: Comparing household earnings inequality across data sets
CPS, PSID, and ACS, Sample B. Households ranked by equivalized household earnings. (A): Ratio of average equivalized earnings of the top 10% to the average of the middle 10%. (B) Ratio of the average equivalized earnings of the middle 10% to the average of the bottom 20%. NBER recessions are shaded in gray.

thus income pooling within couple households has become much less important as a mechanism reducing inequality.

The existence of singles necessarily reduces household pooling, so the household composition term is negative. Over time, there has been a rise in the share of working age individuals living in single adult households, with most of that increase occurring before 1980. All else equal, that trend would have reduced household pooling. But the figure indicates that the impact of household composition on household pooling has been quite stable over time. The reason is that the extent of income pooling *within* couple households – the term in square brackets in the household composition term – has declined over time, due to the sorting and gender gap forces just explained. Thus, single adult households are now less disadvantaged from an income-pooling perspective, which offsets the fact that there are more of them.

6.2 Comparison of earnings inequality across data sets

Before discussing inequality in different income measures at the household level, we first check whether the data sets used in this paper tell a consistent story with respect to earnings. As shown in Figure 15, the CPS, the PSID, and the ACS align closely. This means that researchers can estimate individual income dynamics from the PSID and make comparisons to cross-sectional moments from the much larger CPS and ACS samples. In the

Appendix we also report earnings inequality at the top and the bottom using the CES and the SCF, as well as pre-tax income inequality across the CPS, the PSID, the ACS, the CES, and the SCF. See Figures B.6 and B.7 for more details.

6.3 Inequality in different income measures

From now on, we focus on equivalized measures of household income. Figure 16 offers a big picture summary of how different income measures have evolved across the distribution. This figure is produced using Sample B in the CPS. We rank households year by year by equivalized total market income (all income excluding government transfers and before taxes). We take the richest 10 percent of households by this ranking, along with the middle 10 percent and bottom 20 percent of households. For each of these groups, we compute average equivalized market income, average equivalized market income plus government transfers (pre-tax income), and average equivalized disposable income (pre-tax income minus taxes paid).³⁵

The figure shows trend growth in market income at the top and the middle of the income distribution, but no trend growth in market income at the bottom. Another interesting pattern is that market income at the bottom is much more cyclical than income higher up the distribution, with sharp declines in recessions followed by strong growth in expansions.

At the top and in the middle of the distribution, government transfers are quite small, and as a result, pre-tax income tracks market income closely. But at the bottom, transfers are large, and have risen over time. Thus, pre-tax income at the bottom exhibits positive trend growth, even though market income has been stagnant. Note also that pre-tax income at the bottom is much less cyclical than market income: transfers tend to expand during recessions, automatically stabilizing household incomes.

If we move from pre-tax income to disposable income, the figure shows that taxes reduce income much more at the top than at the bottom, reflecting the fact that the U.S. tax system is quite progressive. The TAXSIM model treats the three rounds of Economic Impact Payments (stimulus checks) during 2020 and 2021 as tax credits. That explains why disposable income grows fast at the bottom in the last two years of our sample, as net taxes paid by the group turn negative.

Figure 16 gives the impression that all measures of income inequality have been widening over time, but the levels of income for these different groups are so different that it is

³⁵Recall that we use the NBER TAXSIM model to compute taxes (Feenberg and Coutts, 1993). Our tax estimates include both federal and state level taxes. See Appendix A for more details.

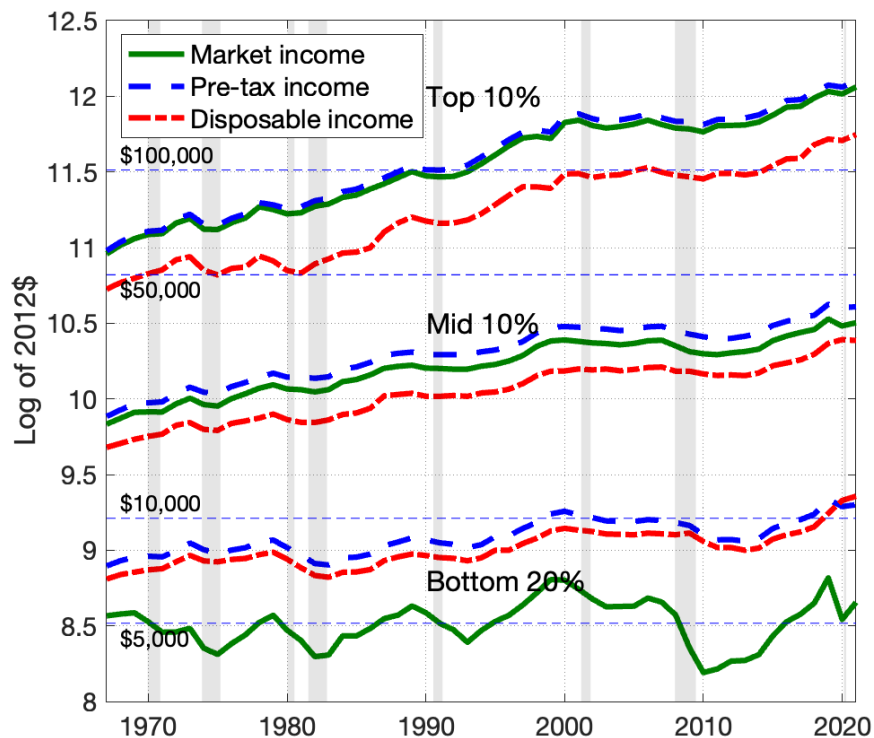


Figure 16: Different income measures across the distribution

CPS, Sample B. Households ranked by equivalized market income. Average values for equivalized market income, equivalized pre-tax income, and equivalized disposable income for households in the top 10%, middle 10%, and bottom 20% of the distribution for equivalized market income. NBER recessions are shaded in gray.

hard to gauge the extent of growth differentials. In Figure 17 we plot ratios of the top income bin to the middle bin, and the ratio of the middle bin to the bottom one. The figure indicates rising inequality at both ends of the distribution.

Inequality has risen steadily at the top, with especially rapid growth between 1980 and 1997. Taxes have somewhat blunted the rise in inequality in disposable income at the top.

The picture for inequality at the bottom is quite different. The rise in inequality in market income is heavily concentrated in recessions, and the cyclical component is so large, especially for the Great Recession and subsequent recovery, that quantifying the secular trend is not easy. Transfers massively reduce inequality at the bottom, and taxes reduce inequality further. Interestingly, the middle-to-bottom ratio for disposable income is quite flat from the mid-1980s onward. This finding is important, and we will return to

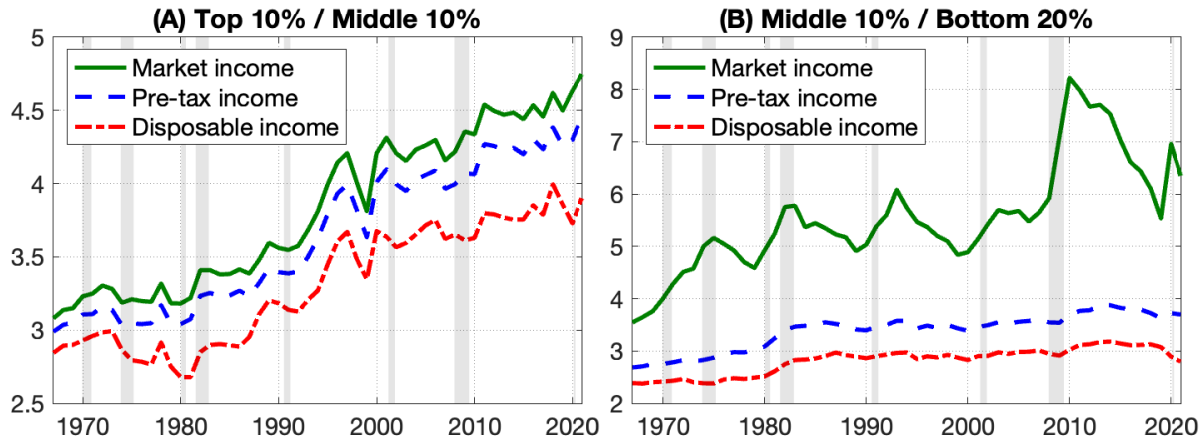


Figure 17: Inequality in different income measures

CPS, Sample B. Households ranked by equivalized market income. (A): Ratio of average equivalized income of the top 10% to average of the middle 10%. (B): Ratio of average equivalized income of the middle 10% to average of the bottom 20%. NBER recessions are shaded in gray.

it when discussing the dynamics of consumption inequality. Another striking finding is that while bottom end inequality in disposable income rose slightly during the Great Recession, it actually declined during the COVID recession. That reflects the unprecedented scale of public income support to households, and especially to low income households, in 2020 and 2021.³⁶

Figure 18 presents similar data in a slightly different way. Here, we rank households year by year three different ways: (i) by male earnings, (ii) by household earnings, and (iii) by household disposable income (all measures are equivalized). For each measure we record income at the 10th, 20th, 25th, 50th, 75th, 90th, 95th and 97th percentiles of the distribution. Then, we plot the evolution of each percentile, normalizing the 1967 value to one in each case.³⁷

These three charts show that income differentials are widening across the distribution for all of these income measures: the percentile plots all fan out over time. But the increase in inequality is smaller in household earnings than in male earnings, and smaller still in household disposable income.³⁸ Most of the widening in the gaps below the 50th

³⁶In addition to direct stimulus payments, child tax credits were expanded, and unemployment benefits and food stamp benefits were made more widely available and more generous. Obviously, disposable income is only one dimension of inequality, and some have argued (see [Stantcheva, 2022](#)) that the COVID crisis has widened inequality along other dimensions—e.g., health status.

³⁷In Panel (A) we drop households which contain no working age men.

³⁸Adding government transfers reduces somewhat the rise in inequality at the bottom, while subtracting taxes reduces the growth in inequality across the distribution.

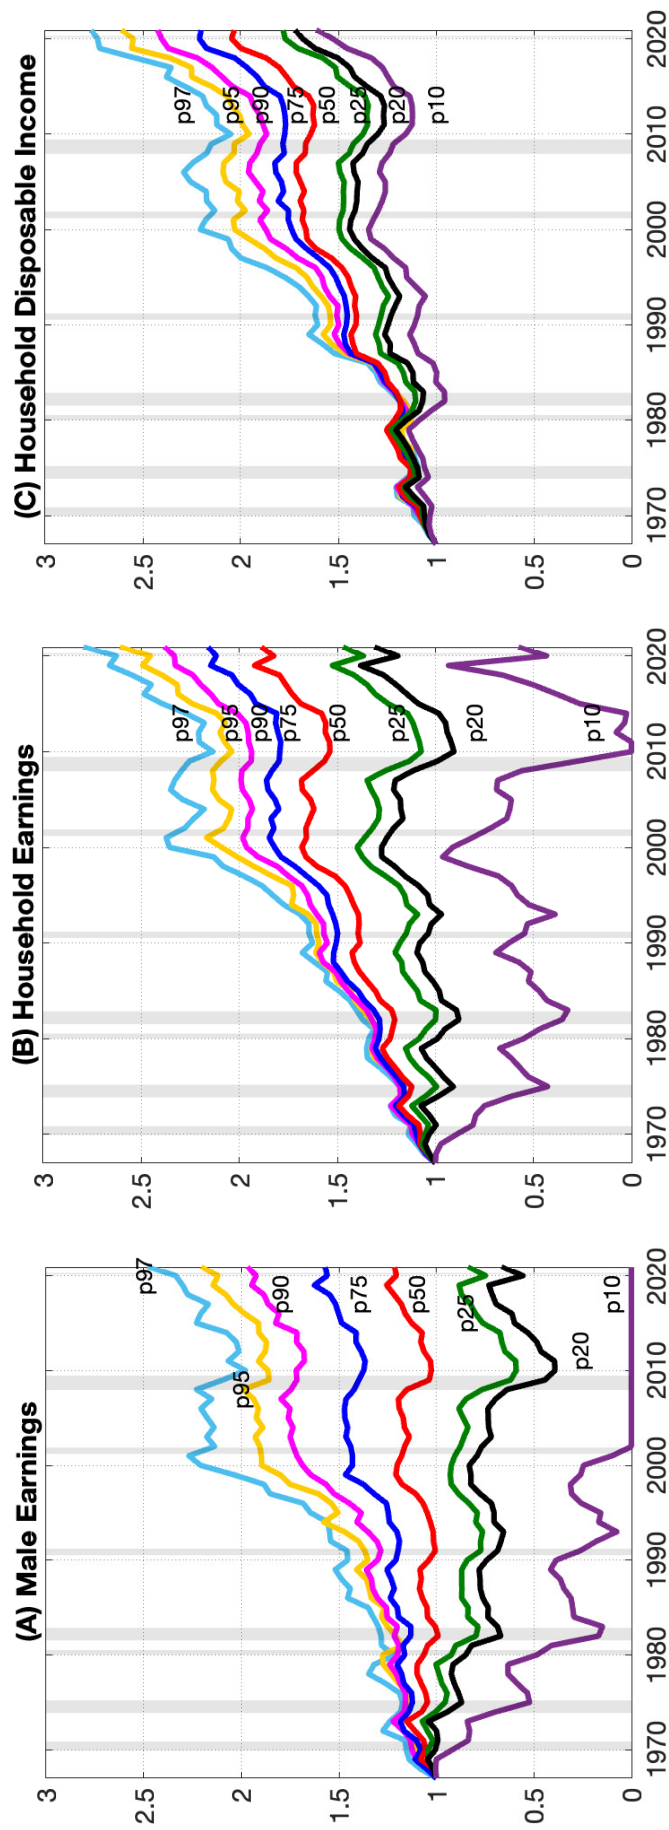


Figure 18: Fan charts for income inequality

CPS, Sample B. (A): Equivalized male earnings. (B): Equivalized household earnings. (C): Equivalized household disposable income. For each of these three real income measures, we plot the evolution of different percentiles of the corresponding cross-sectional distribution, all relative to the income value at that percentile in 1967. NBER recessions are shaded in gray.

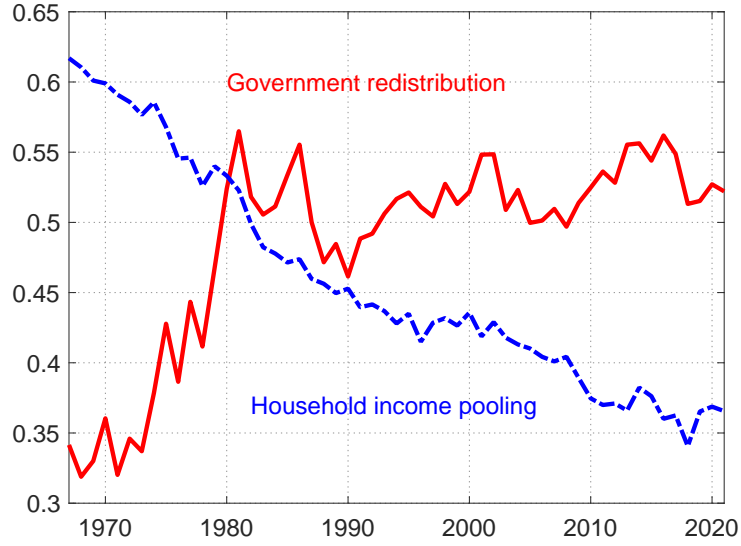


Figure 19: Inequality reduction mechanisms

CPS, Sample B. The blue line is identical to Figure 13. The red line is defined in footnote 39

percentile for disposable income occurs before the mid 1980s, while the widening further up the distribution starts later and continues through to the end of the sample period.

Another sharp message from the plot is that moving to broader income measures also changes the picture for income *growth*. There is only slight growth in real median male earnings over the sample period, but median household earnings rise strongly, and real median household disposable income more than doubles.

In Section 6.1, we documented a decline in the role of household income pooling as an inequality reducing mechanism. Here, we have described how government redistribution has tempered rising inequality. In Figure 19 we measure and contrast the importance of these two mechanisms over time. The dashed blue line is the household income pooling index from equation (2), i.e., the fraction of individual earnings inequality that is eliminated by income pooling within the household. The solid red line is a comparable index for government redistribution. It shows the fraction of inequality in equivalized household market income that is eliminated by the tax and transfer system.³⁹ This fraction grows sharply in the 1970s, and continues to grow, but at a much slower pace, later on. Overall, its increase is comparable in size to the decline in the household pooling index. Thus, over the last half century, the weakening role of the household in reducing inequality has been offset almost one for one by a rising role for the government.

³⁹Formally, the index is a ratio whose numerator is the variance of equivalized market income minus the variance of equivalized disposable income, and whose denominator is the variance of equivalized market income.

Two comments are in order here. First, our measure of government redistribution is computed using the variance of level income, to make it comparable to the household pooling index. This measure is more sensitive to inequality at the top than to inequality at the bottom. We have experimented with computing the extent of government inequality reduction using the top 10% to bottom 20% ratio as the inequality metric. We found that government inequality reduction computed this way is more counter-cyclical, but its overall trend is very similar to the red solid line in Figure 19. Second, we have also computed the fraction of household earnings inequality that is eliminated by non-labor income. We found that this fraction is negative, which means, as expected, that asset income exacerbates inequality across households because it mostly accrues to those at the top of the income distribution. With the important caveat that the CPS under-samples very wealthy households and mis-measures non labor income, we found no discernible trend over time in this fraction.⁴⁰

6.4 Accounting for missing income

In Section 4.1 we noted that households tend to under-report income in surveys, relative to the estimates that inform the national accounts, and the extent of under-reporting varies across different components of income. We now explore how this under-reporting might impact levels and trends in measured income inequality. We proceed as follows. Figure 3 reports ratios of the average values of different income components in our CPS Sample A to their NIPA counterparts. We invert these ratios (focusing on the NIPA+ measures) and use them to scale up our individual CPS income measures proportionately so that – by construction – the average values for each income component in the CPS reproduce the NIPA personal income counterpart in every survey year from 1980 onward. Note that this rescaling procedure implicitly assumes that the share of each income component that is under-reported does not vary across the household income distribution.⁴¹ We then rank households by pre-tax income (market income plus transfers) and compute ratios of top to middle and middle to bottom in Figure 20. The solid blue lines show the dynamics of inequality without any scaling up. The dashed yellow, green, purple and red lines show the impact of scaling up one income component at a time. Finally, the dotted black lines with diamond markers show income inequality when all components are scaled up.

⁴⁰See Figure B.4 in Appendix B for these additional measures of inequality reduction.

⁴¹There is one difference between how we rescale transfers here versus the procedure we applied in Section 4.1, which is that here we do not include EIP payments (stimulus checks) in either our NIPA or CPS

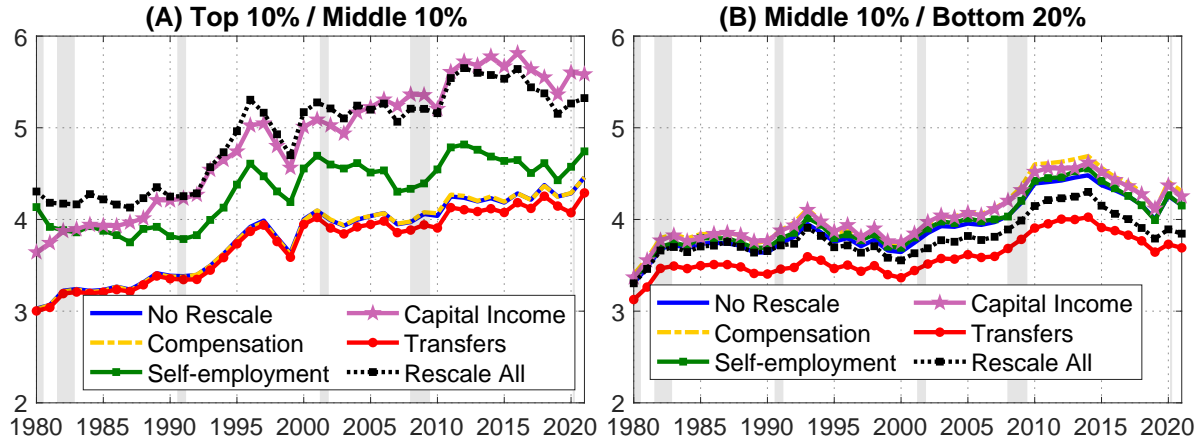


Figure 20: Impact of rescaling different income components on measured inequality CPS, Sample B. Households ranked by raw and rescaled equivalized pre-tax income. Different income components rescaled to match NIPA counterparts. (A): Ratio of average equivalized pre-tax income of the top 10% relative to the average of the middle 10%. Panel (B): Ratio of average equivalized pre-tax income of the middle 10% to the average of the bottom 20%. NBER recessions are shaded in gray.

Consider, first, the top of the income distribution. Here, rescaling compensation or transfer income has almost no impact on measured inequality. Rescaling business income, and especially capital income, however, has the effect of notably amplifying measured inequality. Mechanically, that is because households in the top 10 percent of the market income distribution receive a much larger share of their income from these sources than households in the middle of the distribution. When all income components are rescaled, the level of inequality at the top is higher throughout the sample period, but the extent of the rise in measured inequality is similar.

Panel (B) shows how rescaling affects measured inequality at the bottom. There is only one income component for which rescaling has a notable effect, and that is transfer income. Because households at the bottom derive a much larger share of their income from transfers than households in the middle, scaling up transfer income significantly reduces the level of measured inequality. And furthermore, rescaling transfer income also reduces the measured rise in income inequality at the bottom. That is a consequence of the fact that the ratio of measured transfers to NIPA transfers is declining over time (see Figure 3), so our rescaling procedure inflates income at the bottom by an increasing amount over time.

In our view this rescaling procedure for transfers presents an overly optimistic portrait of transfers, because TAXSIM treats these payments as negative taxes.

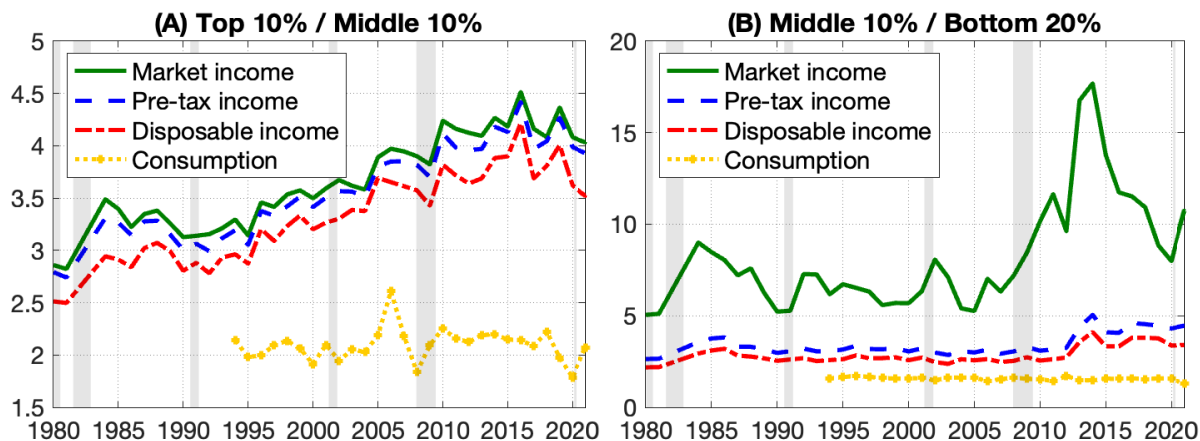


Figure 21: Income and consumption inequality

CES, Sample B. Households ranked by equivalized market income. (A): Ratios of average equivalized income / consumption of the top 10% to the average of the middle 10%. (B): Ratios of average equivalized income / consumption of the middle 10% to the average of the bottom 20%. NBER recessions are shaded in gray.

trait of inequality at the bottom. Most of the difference between CPS and NIPA transfers reflects Medicaid spending. But a large portion of Medicaid spending simply reduces uncompensated care expenses for health care providers, muddying the calculation of who benefits from this spending.⁴²

6.5 Consumption inequality

We now turn to consumption inequality and to the Consumption Expenditure Survey (CES). Figure 21 plots the dynamics of top to middle and middle to bottom inequality for various income measures as well as for consumption, as defined in Section 4.2. Note that this figure is constructed exactly the same way as Figure 17. In particular, we take the CES version of our Sample B, and rank households by equivalized household market income. This ranking identifies a top, middle and bottom group of households at each date, and we compute average income (for each of our income measures) and average consumption expenditure for each of these groups.

The dynamics of income inequality in the CES sample are quite similar to those documented in Figure 17 for the CPS (note that the sample period here is shorter). There is more inequality at the bottom of the income distribution, reflecting the fact that there are

⁴²For example, [Finkelstein et al. \(2019\)](#) estimate that 60 percent of Medicaid spending is a transfer to providers of uncompensated care for the low-income uninsured.

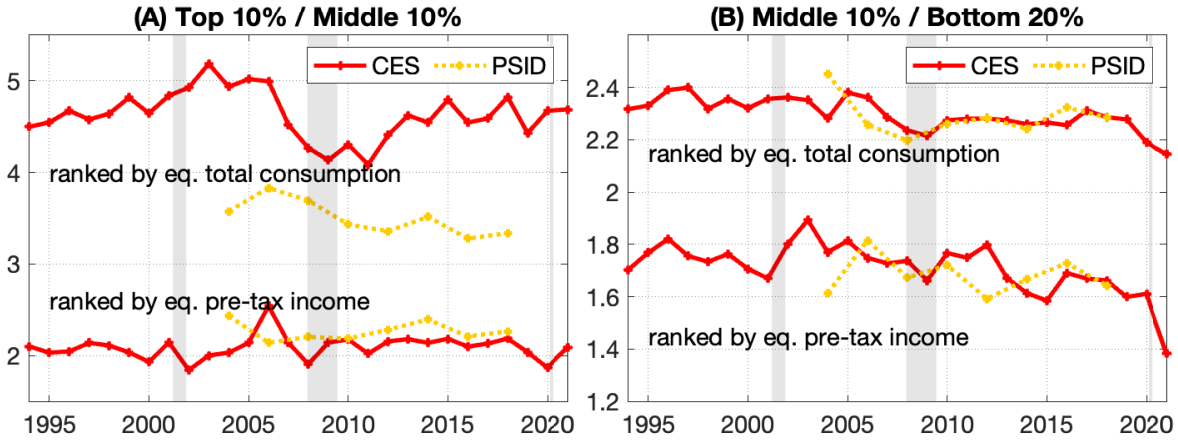


Figure 22: Consumption inequality in different data sets

CES and PSID, Sample B. Households ranked by equivalized pre-tax income and by equivalized total consumption. (A): Ratio of average equivalized consumption of the top 10% to the average of the middle 10%. (B): Ratio of average equivalized consumption of the middle 10% to the average of the bottom 20%. NBER recessions are shaded in gray.

more poor households in the CES than in the CPS (see Table 1). Another notable difference is that market income inequality at the bottom peaks in 2014 in the CES, while the peak is in 2010 in the CPS.⁴³

The striking finding for consumption is that consumption inequality is flat over time, both at the bottom and at the top. The finding that consumption inequality is flat at the bottom should perhaps come as no surprise given that we find little to no increase in inequality in disposable income at the bottom from the mid-1980s onward – a finding that applies to both the CPS and the CES. The absence of an upward trend in consumption inequality at the top is more surprising. Another notable finding is that consumption differentials are quite small relative to income differentials, both at the bottom and at the top.

The other data set in which we can measure inequality in consumption is the PSID. In Figure 22 we compare consumption inequality dynamics between the CES and the PSID in terms of the top-to-middle and middle-to-bottom ratios. The figure shows that these two data sets align well, regardless of whether households are ranked by equivalized pre-tax income or equivalized total consumption expenditures. We again observe little change in consumption inequality in the PSID over the post-2004 period, when a broad measure of consumption is available. Meyer and Sullivan (2023), who focus on well-measured

⁴³Meyer and Sullivan (2023) also find income inequality peaks in 2014 in the CES: see their Figure 1.

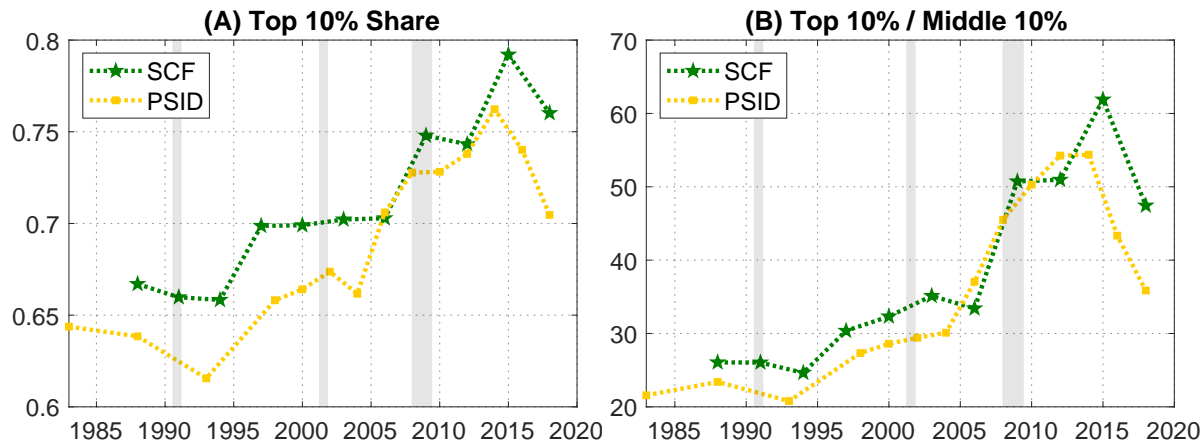


Figure 23: Wealth inequality

SCF and PSID. Sample B. Households ranked by equivalized net worth. (A): Share of total equivalized net worth held by richest 10% of households. (B): Ratio of the average equivalized net worth of the top 10% to the average of the middle 10%. NBER recessions are shaded in gray.

categories of consumption, also find essentially no increase in consumption inequality between 1983 and 2017.⁴⁴

6.6 Wealth inequality

Figure 23 plots two moments of the distribution of household net worth from the SCF and the PSID.⁴⁵ The left panel shows the share of total equivalized net worth held by the richest 10 percent of households, while the right panel plots the top to middle ratio. The two different data sets exhibit very similar trends. Both measures indicate rising wealth inequality through most of the sample period, with a peak sometime in the mid-2010s, followed by some decline. Kuhn et al. (2020) show that some of these dynamics reflect changes in the relative prices of different asset classes that account for different shares of wealth at different points in the distribution. In particular, equity is very important at the top, while housing is the key component of household wealth for those in the middle of the distribution. Thus, the collapse in home values during the Great Recession dramatically depressed median net worth, leading to a spike in the top to middle ratio.

Figure 24 ranks households by equivalized disposable income, and plots the top-to-middle and middle-to-bottom ratios for consumption, disposable income and net worth.

⁴⁴For a possible explanation of the divergence between income and consumption inequality, see also Straub (2019).

⁴⁵Given our earlier finding that aggregate net worth is grossly under-counted in the CES (Figure 5), we do not analyze distributional net worth statistics for the CES.

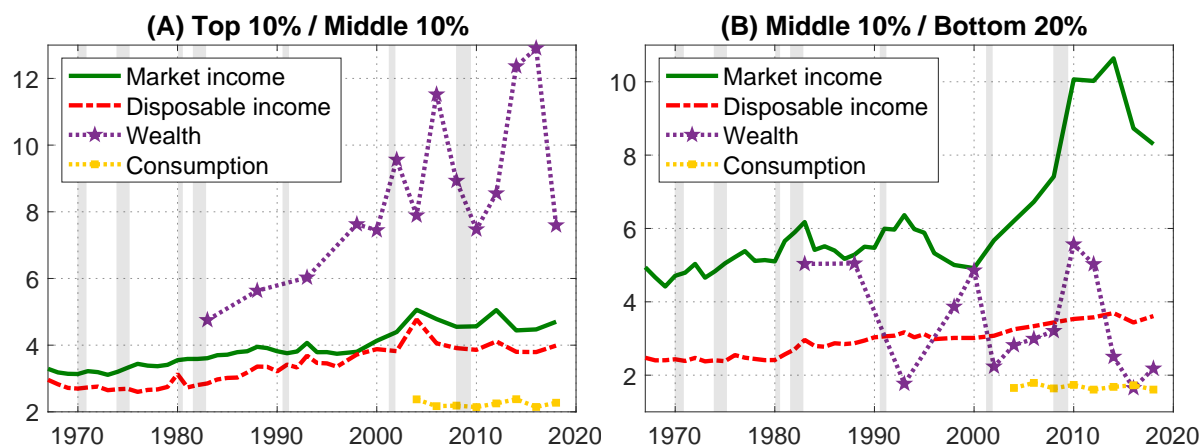


Figure 24: Contribution of wealth inequality to total household resources inequality PSID. Sample B. Households ranked by equivalized market income. (A): Ratios of average equivalized consumption, equivalized disposable income, equivalized market income, and equivalized net worth of the top 10% to the corresponding averages of the middle 10%. (B): Ratios of average equivalized consumption, equivalized disposable income, equivalized market income, and equivalized net worth of the middle 10% to the corresponding averages of the bottom 20%. NBER recessions are shaded in gray.

The main takeaway from this figure is that the wealth ratio below the median is of a similar magnitude to that of the income ratio, whereas above the median, the wealth ratio is larger and grows faster than the income ratio. Thus, the unequal wealth distribution amplifies inequality in total household resources above the median of the income distribution by much more than below the median.

7 Conclusions

In this paper we have studied the evolution of income inequality in the United States over the past half century. Our data sources are all publicly available surveys, so our analysis is easily replicable. We showed that along some dimensions—for example, labor earnings—surveys line up very well with the National Accounts, while for others—for example, capital and self-employment income—there is a larger gap.

We argued that the household budget constraint provides a useful framework for organizing the data and for parsing the various components of income, which, taken individually, sometimes paint contrasting pictures of the evolution of inequality. This approach identifies important roles for labor supply, household pooling, and the government in mediating how the dynamics of individual wage inequality translate into in-

equality in household disposable income.

A declining wage gender gap has helped moderate the overall increase in wage dispersion. At the same time, labor supply has collapsed for men at the bottom of the earnings distribution, amplifying the gap between individual labor earnings at the bottom and the middle/top of the distribution.

Historically, the household has been an important inequality-reducing force. Over time, however, this mechanism has declined quite dramatically, for two reasons. First, the gender gap in earnings has declined, so women are now much less reliant on spousal earnings. Second, the correlation between earnings within couple households has increased over time, reducing the scope for within-household insurance. The diminishing role of the household has been paralleled by an increasing role for the government. At high frequency, taxes and transfers virtually eliminate the cyclicity of inequality at the bottom of the distribution. At low-frequency, they significantly reduce the secular rise in pre-government income inequality.

Our final findings concern consumption expenditures and wealth. In the two data sets that report consumption expenditures, we find that inequality in consumption expenditures has remained remarkably stable over time. On the other hand, inequality in net worth has increased significantly across the income distribution. In particular, between the mid 1990s and the end of our sample period, households in the top 10% of the income distribution experienced significant gains in income and wealth, relative to households in the middle of the income distribution. But these gains do not appear to have translated into an increase in their relative consumption.

References

- Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Elsevier.
- Aguiar, M. and M. Bils (2015). Has consumption inequality mirrored income inequality? *American Economic Review* 105(9), 2725–2756.
- Aguiar, M. A., M. Bils, and C. Boar (2020). Who are the hand-to-mouth? National Bureau of Economic Research.
- Altonji, J. G., D. M. Hynsjö, and I. Vidangos (2022). Individual earnings and family income: Dynamics and distribution. *Review of Economic Dynamics*.
- Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica* 85(3), 693–734.
- Attanasio, O., E. Battistin, and H. Ichimura (2007). What really happened to consumption inequality in the United States? In E. Berndt and C. Hulten (Eds.), *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, pp. 515–543. University of Chicago Press.
- Attanasio, O. P. and L. Pistaferri (2016). Consumption inequality. *Journal of Economic Perspectives* 30(2), 3–28.
- Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in US wage inequality: Revising the revisionists. *The Review of Economics and Statistics* 90(2), 300–323.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118(4), 1279–1333.
- Ben-Shalom, Y., R. Moffitt, and J. K. Scholz (2012). An assessment of the effect of anti-poverty programs in the United States. In P. N. Jefferson (Ed.), *The Oxford Handbook of the Economics of Poverty*, pp. 709–749. Oxford University Press.
- Benhabib, J. and A. Bisin (2018). Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature* 56(4), 1261–1291.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Borella, M., M. De Nardi, M. Pak, N. Russo, and F. Yang (2022). The importance of modeling income taxes over time. US reforms and outcomes. National Bureau of Economic Research.

- Bricker, J., A. Henriques, J. Krimmel, and J. Sabelhaus (2016). Measuring income and wealth at the top using administrative and survey data. *Brookings Papers on Economic Activity* 2016(1), 261–331.
- Castaneda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (1998). Exploring the income distribution business cycle dynamics. *Journal of Monetary Economics* 42(1), 93–130.
- Cortes, G. M., N. Jaimovich, and H. E. Siu (2017). Disappearing routine jobs: Who, how, and why? *Journal of Monetary Economics* 91, 69–87.
- Cynamon, B. Z. and S. M. Fazzari (2017). Household income, demand, and saving: deriving macro data with micro data concepts. *Review of Income and Wealth* 63(1), 53–69.
- De Nardi, M. and G. Fella (2017). Saving and wealth inequality. *Review of Economic Dynamics* 26, 280–300.
- Diaz-Gimenez, J., V. Quadrini, and V. Rios-Rull (1997). Dimensions of inequality: Facts on the U.S. distribution of earnings, income and wealth. *Federal Reserve Bank of Minneapolis Quarterly Review*, 21, 3–21.
- Feenberg, D. and E. Coutts (1993). An introduction to the TAXSIM model. *Journal of Policy Analysis and Management* 12(1), 189–194.
- Finkelstein, A., N. Hendren, and E. F. Luttmer (2019). The value of Medicaid: Interpreting results from the Oregon Health Insurance Experiment. *Journal of Political Economy* 127(6), 2836–2874.
- Fisher, J., D. Johnson, T. Smeeding, and J. Thompson (2018). Inequality in 3-D: Income, consumption, and wealth. Finance and Economics Discussion Series (FEDS), Board of Governors of the Federal Reserve System.
- Fortin, N. M. and T. Lemieux (1997). Institutional changes and rising wage inequality: Is there a linkage? *Journal of Economic Perspectives* 11(2), 75–96.
- Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Gramm, P., R. Ekelund, and J. Early (2022). *The Myth of American Inequality: How Government Biases Policy Debate*. Rowman & Littlefield Publishers.
- Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2014). Marry your like: Assortative mating and income inequality. *American Economic Review* 104(5), 348–353.
- Guner, N., C. Rauh, and G. Ventura (2023). Means tested transfers in the U.S.: Facts and parametric estimates. University of Cambridge.
- Guvenen, F., L. Pistaferri, and G. L. Violante (2022). Global trends in income inequality and income dynamics: New insights from GRID. *Quantitative Economics* 13(4), 1321–1360.

- Heathcote, J. and F. Perri (2018). Wealth and Volatility. *The Review of Economic Studies* 85(4), 2173–2213.
- Heathcote, J., F. Perri, and G. L. Violante (2010). Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006. *Review of Economic Dynamics* 13(1), 15–51.
- Heathcote, J., F. Perri, and G. L. Violante (2020). The rise of U.S. earnings inequality: Does the cycle drive the trend? *Review of Economic Dynamics* 37, S181–S204.
- Heathcote, J., K. Storesletten, and G. L. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review* 104(7), 2075–2126.
- Helpman, E. (2018). *Globalization and Inequality*. Harvard University Press.
- Hoffmann, F., D. S. Lee, and T. Lemieux (2020). Growing income inequality in the United States and other advanced economies. *Journal of Economic Perspectives* 34(4), 52–78.
- Hyslop, D. R. (2001). Rising U.S. earnings inequality and family labor supply: The covariance structure of intrafamily earnings. *American Economic Review* 91(4), 755–777.
- Johnson, D. and S. Shipp (1997). Trends in inequality using consumption-expenditures: The US from 1960 to 1993. *Review of Income and Wealth* 43(2), 133–152.
- Juhn, C. and K. M. Murphy (1997). Wage inequality and family labor supply. *Journal of Labor Economics* 15(1, Part 1), 72–97.
- Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? Evidence and theory. *The Review of Economic Studies* 73(1), 163–193.
- Kuhn, M. and J.-V. Ríos-Rull (2016). 2013 update on the U.S. earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review* 37(1), 2–73.
- Kuhn, M., M. Schularick, and U. I. Steins (2020). Income and wealth inequality in America, 1949–2016. *Journal of Political Economy* 128(9), 3469–3519.
- Lippi, F. and F. Perri (2023). Unequal growth. *Journal of Monetary Economics* 133, 1–18.
- Mare, R. (1991). Five decades of educational assortative mating. *American Sociological Review* 56, 15–32.
- McKinney, K. L., J. M. Abowd, and H. P. Janicki (2022). US long-term earnings outcomes by sex, race, ethnicity, and place of birth. *Quantitative Economics* 13(4), 1879–1945.
- Meyer, B. D. and J. X. Sullivan (2023). Consumption and income inequality in the United States since the 1960s. *Journal of Political Economy* 131(2), 247–284.
- Nelson, J. A. (1994). On testing for full insurance using Consumer Expenditure Survey data. *Journal of Political Economy* 102(2), 384–394.

- Piketty, T., E. Saez, and G. Zucman (2017). Distributional national accounts: Methods and estimates for the United States. *The Quarterly Journal of Economics* 133(2), 553–609.
- Santos, C. and M. Tertilt (2023). How families matter for income inequality. *IFS Deaton Review of Inequalities*.
- Slesnick, D. T. (2001). *Consumption and social welfare: Living standards and their distribution in the United States*. Cambridge University Press.
- Smith, M., D. Yagan, O. Zidar, and E. Zwick (2019). Capitalists in the twenty-first century. *The Quarterly Journal of Economics* 134(4), 1675–1745.
- Stantcheva, S. (2022). Inequalities in the times of a pandemic. *Economic Policy* 37(109), 5–41.
- Straub, L. (2019). Consumption, savings and the distribution of permanent income. Harvard University.
- Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy* 110(4), 825–853.

Appendix

A Details of data construction

In this appendix we discuss in detail how we constructed the data in each figure.

A.1 Definition of income measures

We first define our CPS income measures and the exact IPUMS CPS income components in each, since the CPS is the main data source we use in the paper.

Labor Income 1967-1986 = incwage, 1987-2016 = inclongj (if srcearn=1 (wage and salary)) + oincwage, where incwage = income from wage and salary; inclongj = earnings from longest job before deductions; oincwage = income from other wage and salary.

Self Employment Income 1967-1986 = incbus + incfarm, 1987-2016 = inclongj (if srcearn=2 or 3 (farm or non-farm self-employment)) + oincbus + oincfarm, where incbus = income from non-farm self-employment; incfarm = income from farm or nonincorporated self-employment; oincbus = income from other work (own business self-employment); oincfarm = income from other work (farm self-employment).

Earnings = labor income + self-employment income.

Private Transfers 1967-1974 = incother, 1975-1986 = incretir + incaloth + incother, 1987-2016 = incother + incalim + incchild + incdisa1 + incdisa2 + incasist + increti1 + increti12 + incsurv1 + incsurv2 where incother = income from other sources; incretir = income from retirement funds; incaloth = income from alimony and child support; incalim = income from alimony; incchild = income from child support; incdisa1 = income from disability income – primary source; incdisa2 = income from disability income – secondary source; incasist = income from contributions, assistance from friends; increti1 = income from retirement income – primary source; increti12 = income from retirement income – secondary source; incsurv1 = income from survivors income – primary source; incsurv2 = income from survivors income - secondary source.

Net Asset Income 1967-1974 = incidr, 1975-1986 = incint + incdrt, 1987-2016 = incint + incdivid + incrent, where incidr = income from interest, dividends and net rentals; incint = income from interest; incdrt = income from dividends, rents and trusts; incdivid = income from dividends; incrent = income from rent.

Market Income (=pre-government income) = earnings + private transfers + net asset income.

Public Transfers 1967-1974 = incwelfr + incgov + incss, 1975-1986 = incwelfr + incgov + incss + incssi, 1987-2016 = incwelfr + incss + incssi + inceduc + incvet + incwkcom + incunemp, where incwelfr = income from public assistance or welfare; incgov = income from unemployment compensation, workmen's compensation, government employee pensions, and veterans' benefits.; incss = income from Social Security or railroad retirement – from US govt; incssi = income from supplemental security; inceduc = income from educational assistance; incvet = income from veterans payments; incwkcom = income from worker's compensation; incunemp = income from unemployment compensation.

Pre-tax Income = market income + public transfers + 0.5*FICA.

Disposable Income = pre-tax income - federal-level taxes - state-level taxes - FICA.

Note that both tax liabilities and FICA are estimated at the household-level using the NBER's TAXSIM program.

A.2 Top-coding

Top-coding is an important issue to address in survey data, both for computing means, and for measuring the evolution of inequality at the top of the income distribution.

We first discuss how we deal with the top-coding issue in the CPS.

Public top-code thresholds vary widely across income categories, and across time. An additional problem is that the Census Bureau's internal data are also subject to censoring (to economize on computer tape, and to protect against gross errors). For example, the public use censoring point for the variable incwage (income from wages and salaries) was \$50,000 for the income years 1975–1980, \$75,000 for 1981–1983, and \$99,999 for 1984–1986. For the same variable, the internal CPS censoring points were \$99,999 for the period 1975–1984, and \$250,000 for 1985–1986.

The Census Bureau also changed the procedures to report the topcode and replacement values twice over the income years 1967-2021. In income year 1995 (i.e. survey year 1996), the CPS started reporting cell means for top-coded observations, with cells identified by gender, race, and work experience. From 1997 onward (i.e., survey year 1998), more income variables were included in this procedure.

In 2010 (i.e. survey year 2011), the Census Bureau shifted from the cell-means replacement value system to a rank proximity swapping procedure. In this technique, all values greater than or equal to the income topcode are ranked from lowest to highest and systematically swapped with other values within a bounded interval. All swapped values are also rounded to two significant digits.

The procedure that we use to address top-coding is as follows:

- Before 1997 (or 1995 for income variables `oincwage`, `oincbus`, `oincfarm`), we use a regression approach, described below.
- For 1997 (or 1995 for income variables `oincwage`, `oincbus`, `oincfarm`) to 2009, we use the cell-means replacement value system for all income categories, except for the variable `inlongj` (i.e. income from primary source).
- After 2009, we use the rank proximity swapping procedure for all income categories, except for the variable `inlongj`.
- In order to measure changes in inequality consistently over time, we use the regression approach for `inlongj` for all income years 1967 - 2016.

Here are more details on the different approaches:

Regression Approach: In the regression approach, we deal with top-coded observations by assuming the underlying distribution for each component of income is Pareto, and follow a suggestion from David Domeij by forecasting the mean value for top-coded observations by extrapolating a Pareto density fitted to the non-top-coded upper end of the observed distribution. This procedure automatically takes care of the internal censoring problem, since the internal threshold always exceeds the public use limit. It also has the advantage that in principle it adjusts appropriately to changes in top-code thresholds. We apply this procedure at the most disaggregated decomposition of income possible. Thus, for example, for each year we divide the set of observations for the variable `inlongj` (income from primary source) according to whether or not they are flagged as wage and salary or self-employment, and run separate regressions on the two sets of observations. This is important for two reasons. First, for any given individual, while one type of income may be top-coded, others will not be. Second, there is more upper tail concentration in some types of income than others.

Cell-Means Replacement: When the Census Bureau began reporting cell means, they drastically reduced public use censoring points for many income categories: the threshold for interest income declined from \$99,999 to \$35,000 between income years 1997 and 1998 and to \$25,000 in 2002, while the threshold for dividend income declined from \$99,999 to \$15,000. We found that when the distribution is truncated too far to the left, the Pareto-extrapolation procedure does not always perform well. Thus, for income years 1998 to 2009, we use cell means for all income categories, except income from primary source.

Rank Proximity Swapping In 2012, the Census Bureau released a series of revised income topcodes files. Each file contains income values to replace topcoded income components for every ASEC sample from 1975 to 2009 (or survey years 1976-2010). IPUMS CPS provides the Census revisions with IPUMS identifiers and income variable names on their [website](#). However, not all the top-coded values have a corresponding swap value in the file provided by IPUMS CPS. We check all the variables for computing earnings and asset income and find that except for incfarm and incdrt, all the other variables have some missing swap values for at least one year. The coverages are 66.4%, 91.7%, 57.5%, 94.8%, 74.5%, 100%, 60.9%, 51%, 74.5%, 99.6%, and 69.1% for variables inclongj, incwage, oincwage, incbus, oincbus, incfarm, oincfarm, incint, incrent, incdrt, and incdivid respectively. That means, for example, 66.4% of the topcoded values for variable inclongj have swap values in the file.

Top Coding in Other Data Sets: Since top-coding affects very few observations in the PSID, we do not make any adjustments for top-coding in the PSID.

For the CES, only a very limited number of consumption categories are subject to top-coding, i.e. only some categories of medical spending (such as hospital services). Since the non-durable consumption expenditures considered in our paper exclude health expenditures, we do not need to correct for it. Topcoding in earnings is potentially more important as the fraction of topcoded earnings observations in some years reaches 2% of the sample. Also public topcoding thresholds vary across income categories, and across time. We deal with top-coded observations in the CES following the same regression approach as we use in the CPS. That is, we assume that the underlying distribution for each component of income is Pareto, and we forecast the mean value for top-coded observations by extrapolating a Pareto density fitted to the non- top-coded upper end of the observed distribution. This procedure automatically adjusts appropriately to changes in top code thresholds.

We apply the same regression-based procedure in the ACS as we do in the CPS and the CES.

A.3 Details of data construction for each figure

Figure 1: Wages and Salaries per capita NIPA: Wage and salaries (BEA-NIPA Table 2.1 line 3) divided by US population (BEA-NIPA Table 7.1 line 18), adjusted by PCE deflator

CPS: Sample A: individual level earnings = wages/salaries + self-employment income, divided by total number of individuals in Sample A, adjusted by PCE deflator

PSID: Sample A: household level earnings = wages/salaries + self-employment in-

come, divided by the average household size in Sample A, adjusted by PCE deflator

ACS: Sample A: individual level earnings = wages/salaries + self-employment income, divided by total number of individuals in Sample A, adjusted by PCE deflator

CES: Sample A: household level earnings = wages/salaries + self-employment income, divided by the average household size in Sample A, adjusted by PCE deflator

SCF: Sample A: household level earnings = wages/salaries + self-employment income, divided by the average household size in Sample A, adjusted by PCE deflator

Figure 2 and Figure B.2: Pretax income per capita NIPA+ is personal income that we draw directly from BEA Table 2.1 line 1. NIPA- is a pre-tax measure we construct to be as consistent as the definition used in the CPS as possible.

To construct NIPA- we start from NIPA+ and exclude three major categories of income: (1) line 7 in BEA Table 2.1: employer contributions for employee pension and insurance funds; (2) imputed rents of owner-occupied housing; and (3) Medicare and Medicaid benefits.

Both variables are divided by US population (BEA-NIPA Table 7.1 line 18) and adjusted by the PCE deflator.

CPS: Sample A: individual level pre-tax income = earnings + private non-earnings (interest, dividend, child, retirement, rent) + public transfer (social security income, UI, veteran's benefits, income from SSI, income from educational assistance, income from worker's compensation, welfare (public assistance) income (INCWELFR), EIP payments (stimulus checks) + 0.5*FICA estimated using NBER TAXSIM, divided by total number of individuals in Sample A, and adjusted by the PCE deflator.

Note that in order to be consistent with NIPA- and NIPA+ definitions of pre-tax income, the CPS pre-tax income measure here does not include income from assistance (incasist), but it includes EIP payments (stimulus checks). In order to use the CPS series to estimate the EIP payments, we take federal tax credits (fedtax-fedtaxac in IPUMS CPS) and subtract both the earned income tax credit (eitcred in IPUMS CPS) and the additional child tax credit (actccrd in IPUMS CPS).

In Figure B.2, we plot series of pre-tax income for the other surveys. The definitions for pre-tax income in the other surveys are as follows:

PSID: Sample A: household level pre-tax income = total family money income (e.g., V81 for 1968 in PSID) + 0.5*FICA estimated using NBER TAXSIM, divided by the average household size in Sample A, adjusted by the PCE deflator. The income reported here was collected in t about tax year $t - 1$. Please note that this variable can contain negative

values. Negative values indicate a net loss, which in waves prior to 1994 were bottom-coded at \$1, as were zero amounts. These losses occur as a result of business or farm losses. This variable is the sum of these seven variables: (1) ER77352 Reference Person and Spouse/Partner Taxable Income, (2) ER77413 Reference Person and Spouse/Partner Transfer Income (Except Social Security), (3) ER77420 Taxable Income of Other FU Members, (4) ER77441 Transfer Income of OFUMS (Except Social Security), (5) ER77442 Reference Person Social Security Income, (6) ER77444 Spouse/Partner Social Security Income, and (7) ER77446 OFUM Social Security Income.

ACS: Sample A: individual level pre-tax income = inctot (total pre-tax personal income, using variable directly from IPUMS USA) + $0.5 \times \text{FICA}$ estimated using NBER TAXSIM, divided by total number of individuals in Sample A, adjusted by the PCE deflator

CES: Sample A: household level pre-tax income = family total income before tax (using variable finchtax directly from CES) + $0.5 \times \text{FICA}$ estimated using NBER TAXSIM, divided by the average household size in Sample A, adjusted by the PCE deflator

SCF: Sample A: household level pre-tax income = family total wage/salary income + income from a professional practice, business, or farm + private transfer (child support or alimony + income from any other sources) + asset income (income from non-taxable investments such as municipal bonds + interest income + dividends + income from net gains or losses from the sale of stocks, bonds, or real estate + income from net rent, trusts, or royalties from any other investment or business) + public transfer (income from unemployment or worker's compensation + income from ADC, AFDC, food stamps, or other forms of welfare or assistance such as SSI + income from Social Security or other pensions, annuities, or other disability or retirement programs) + $0.5 \times \text{FICA}$ estimated using NBER TAXSIM, divided by the average household size in Sample A, adjusted by the PCE deflator

Figure 3: Gaps between NIPA and CPS (by component) Transfer, NIPA+ = Table 2.1: line 17 (Government social benefits to persons)

Transfer, NIPA- = Table 2.1: line 18 (Social Security) + line 21 (Unemployment insurance) + line 22 (Veterans' benefits) + line 23 (Other)

Capital income, NIPA+ = Table 2.1: line 12 (Rental income of persons with capital consumption adjustment) + line 13 (Personal income receipts on assets: personal interest income + personal dividend income)

Capital income, NIPA- = Table 2.1: line 12 (Rental income of persons with capital consumption adjustment) + line 13 (Personal income receipts on assets: personal interest

income + personal dividend income) - line 8 of NIPA Table 7.9 (imputed rents of owner-occupied housing)

Figure 4: Consumption Expenditures

Non-durable consumption

- NIPA: Table 2.4.5: Food and beverages purchased for off-premises consumption + Clothing and footwear + Transportation services + Food services + Education services + Household utilities + Motor vehicle fuels, lubricants, and fluids + Recreational items + Recreation services + Household maintenance + Telecommunication services + Net motor vehicle and other transportation insurance
- CES. The CES definition is consistent with the NIPA measure. Sample A. Non-durables = transportation (car insurance, repair, gas, parking, bus, taxi, other) + education + utilities (heat, elec, water, other) + clothing + leisure (trips + other recreation activities) + childcare + phone, divided by the average household size in Sample A, adjusted by the PCE deflator. Note that health expenditures are not included.

Note that for some types of non-durable goods (especially apparel and footwear), there is a big gap between the expenditure per capita we get from the CES interview survey and the counterpart time-series reported by the BLS. In general, our results calculated using the PUMD from the BLS differ from the data in the published tables for two reasons: (1) the published tables use integrated data from both the Interview and Diary Surveys; and (2) the published tables use confidential data, and are not adjusted to comply with non-disclosure requirements.

Durable consumption

- NIPA: Table 2.4.5: motor vehicles and parts + furnishings and durable household equipment
- CES. The CES definition is consistent with the NIPA measure. Sample A. Durables = cars (expenses) + furniture/equipment, divided by the average household size in Sample A, adjusted by the PCE deflator.

Note that consumption expenditures are on a quarterly basis. We compute quarterly averages first, then multiply by four to get annual expenditures.

Figure 5 Aggregate Wealth

- Flow of Funds data is for the household sector only (excluding non-profit institutions)
- SCF: household level net worth, divided by the average household size in Sample A, adjusted by the PCE deflator
- PSID: total family wealth including home equity (WEALTH2), divided by the average household size in Sample A, adjusted by the PCE deflator
- CES: total family wealth = financial wealth (finwea) + approximate value of property that would sell for on today's market (propvalx) - home equity loans (qblncm3g) - mortgages (qblncm3x) - amount owed prior to last payment (totowed), divided by the average household size in Sample A, adjusted by the PCE deflator. Note that wealth data are on a quarterly basis. We compute quarterly averages first, then multiply by four to get annual data.

Figures 6 and 7 Wage inequality We construct these figures using Sample C, where all the individuals are between 25 and 60 years old. We mark the wage observation as missing if annual hours are less than 260. There are two measures of hours in the CPS: (1) "new hours": hours worked in an average week last year; (2) hours worked last week. The new hours measure is superior, but is available only since 1976. We therefore generate two measures of wages, using both measures of hours. For any statistic related to wages (e.g., the variance of log wage), we splice together the series computed both ways, statistic by statistic. This procedure also applies to Figure 8, 9, and 10.

Figure 9 Observables and residuals The variance of residuals are generated from a year-by-year regression of log wages on standard demographics including an education dummy (= 1 for college degree, and 0 otherwise), a race dummy (= 0 for white, = 1 for black, = 2 for other races), a task dummy (= 0 for routine cognitive and manual (RCM) tasks, = 1 for non routine manual (NRM) tasks, = 2 for abstract cognitive (AC) tasks), age, and age squared.

Figure 11 Earnings, wages, and weeks worked in three slices of the earnings distribution We start with the Sample C in which all individuals are between 25 and 60 years

old, and select the sub-sample of all men. The number of weeks measure we use to construct the weekly wage is an intervalled variable *weeks worked last year* (the non-intervalled variable is available starting only in 1976) and can take values in the set {0, 6, 20, 33, 43, 48.5, 51} as in [Heathcote et al. \(2020\)](#).

Figure 16 and 17 Household income inequality We use Sample B where:

- Market income = earnings + private non-earnings (interest, dividend, child, retirement, rent, assistance)
- Pre-tax income = market income + public transfer (social security income, UI, veteran's benefits, income from SSI, income from educational assistance, Income from worker's compensation, Welfare (public assistance) income) + 0.5*FICA estimated using NBER TAXSIM. Note that here, we do not include EIP (stimulus checks), because they are considered as tax credits and included in disposable income.
- Disposable income = pre-tax income - federal level tax liabilities - state level tax liabilities - FICA (all estimated using NBER TAXSIM). Federal taxes are net of tax credits, including EIP. Note that state level tax liabilities are available only since 1977, so we splice together two series for statistics involving disposable income computed (1) without state taxes and (2) with state taxes, statistic by statistic.

Figure 21 Household consumption inequality

- nondurables = transportation (car insurance, repair, gas, parking, bus, taxi, other) + education + utilities (heat, electricity, water, other) + clothing + leisure (trips + other recreation activities) + childcare + phone
- durables = cars (expenses) + furniture/equipment
- expenditures in the figure = nondurables + durables

Note health expenditures and housing are not included.

B Additional Figures

This Appendix contain additional figures that complement those in the main text of the paper.

Household size and wages and salaries per household In Figure B.1 we plot household size and wages and salaries per household in the PSID, the CPS and the ACS, Sample A. The main point is to show that the reason why wages and salaries per capita in PSID are higher than than in the two other surveys for most of the period (as shown in Figure 1) is that in PSID household size is significantly lower than household size in the other two surveys.

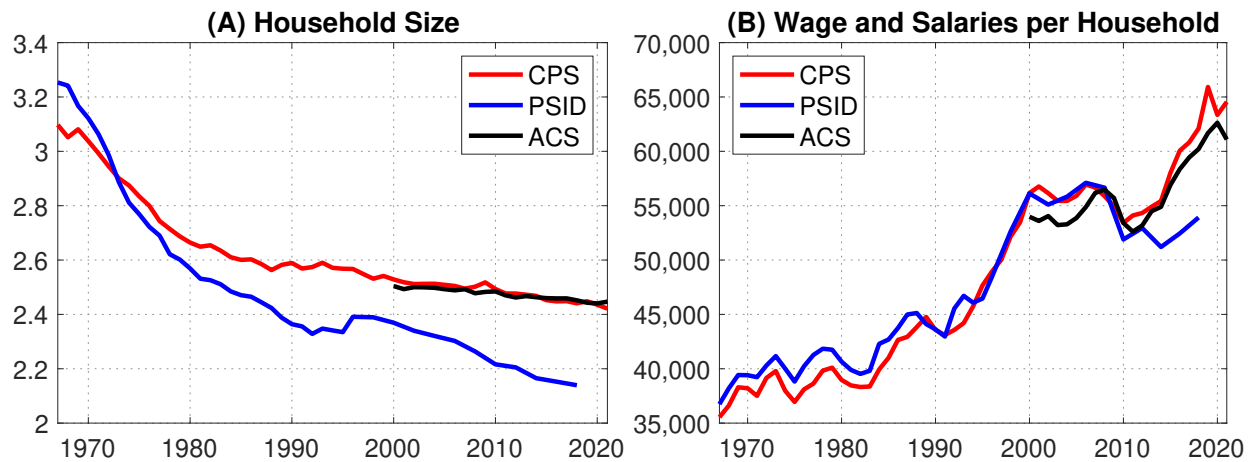


Figure B.1: Panel (A): Household size. Panel (B): Wages/salaries per household. PSID, CPS and ACS, Sample A

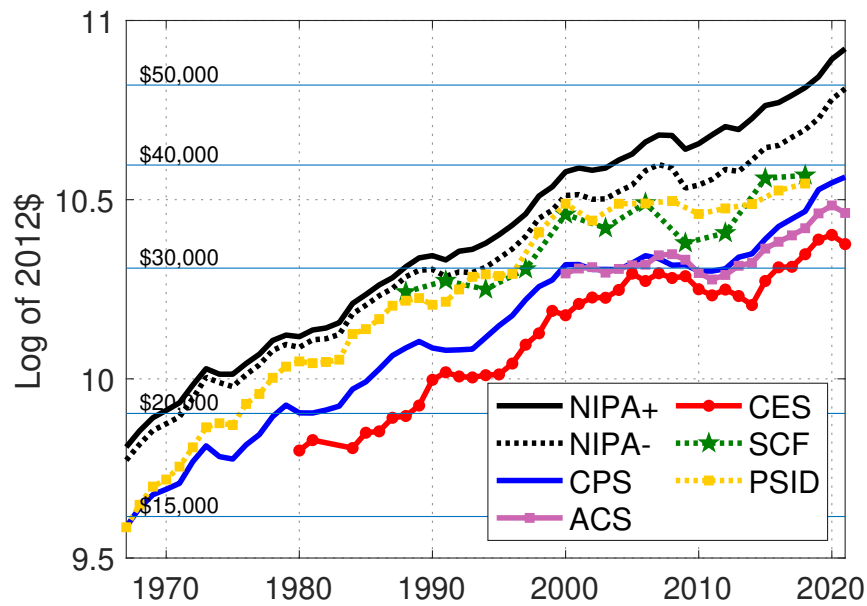


Figure B.2: Pretax Income in NIPA and 5 surveys, Sample A

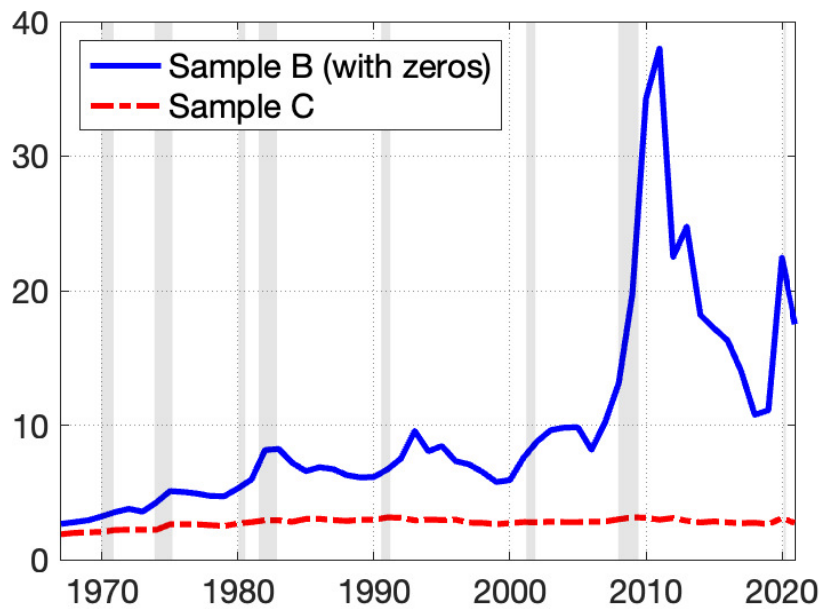


Figure B.3: The ratio of earnings of the middle 10% to the bottom 20% for two samples of men. CPS, Sample B and C.

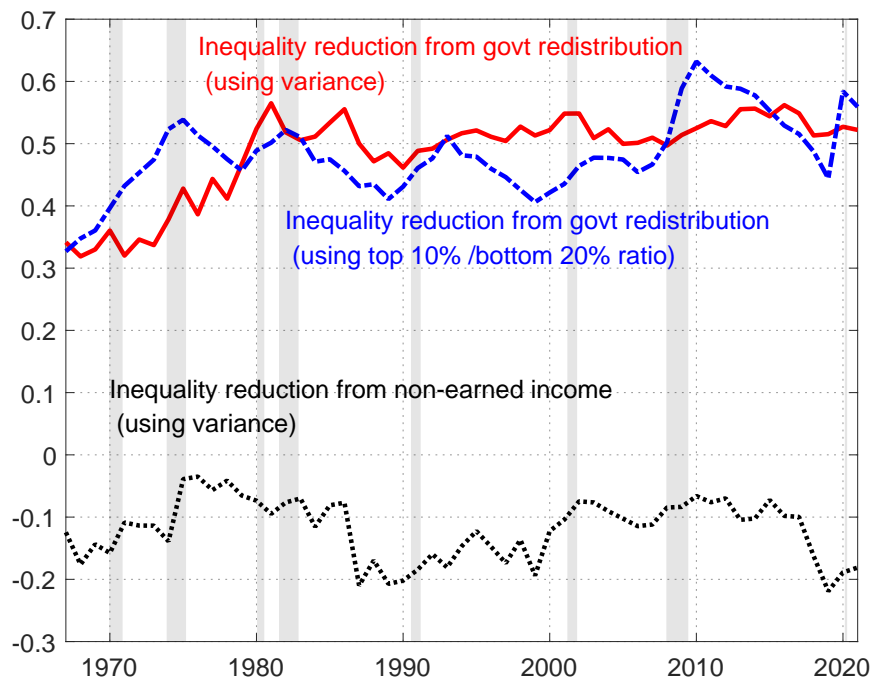


Figure B.4: Alternative measures of inequality reduction. CPS, Sample B. NBER recessions are shaded in gray

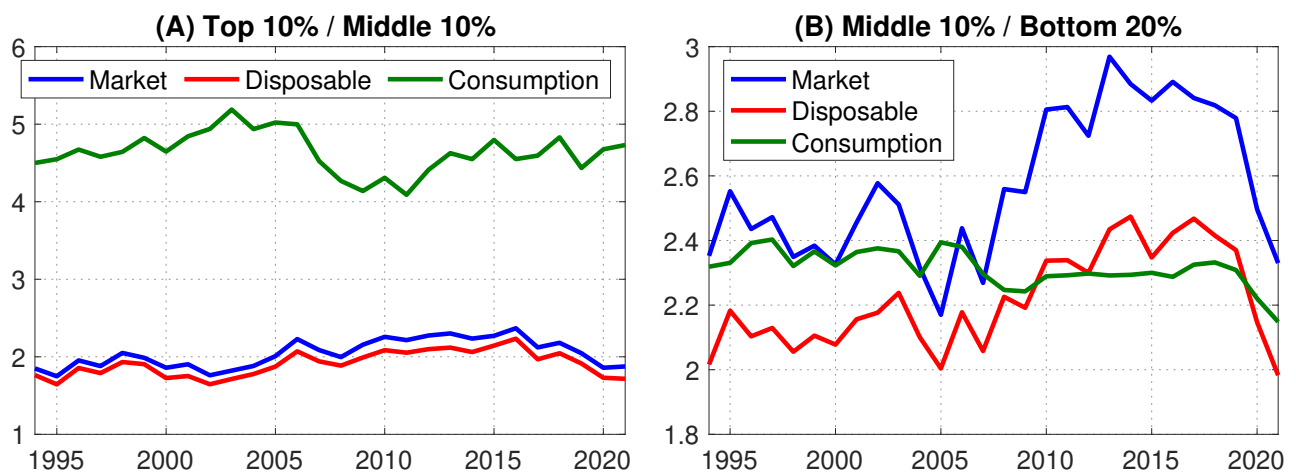


Figure B.5: Panel (A): Inequality in income and consumption at the top. Panel (B): Inequality in income and consumption at the bottom. Households ranked by equivalized total consumption. CES, Sample B.

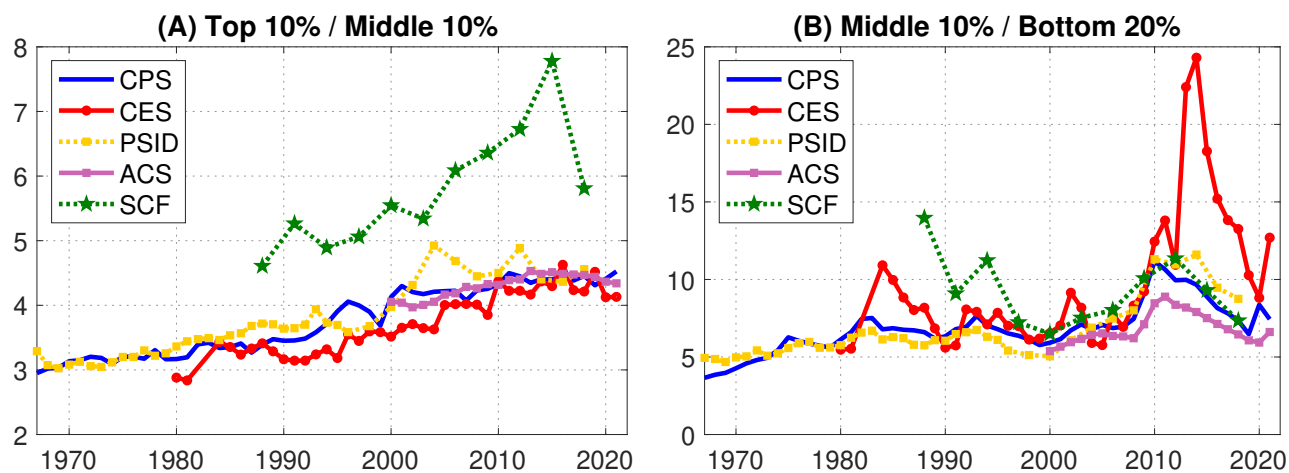


Figure B.6: Cross-datasets Comparison (Earnings). Panel (A): Earnings Inequality at the top. Panel (B): Earnings Inequality at the bottom. CPS, PSID, CES, ACS, and SCF, Sample B.

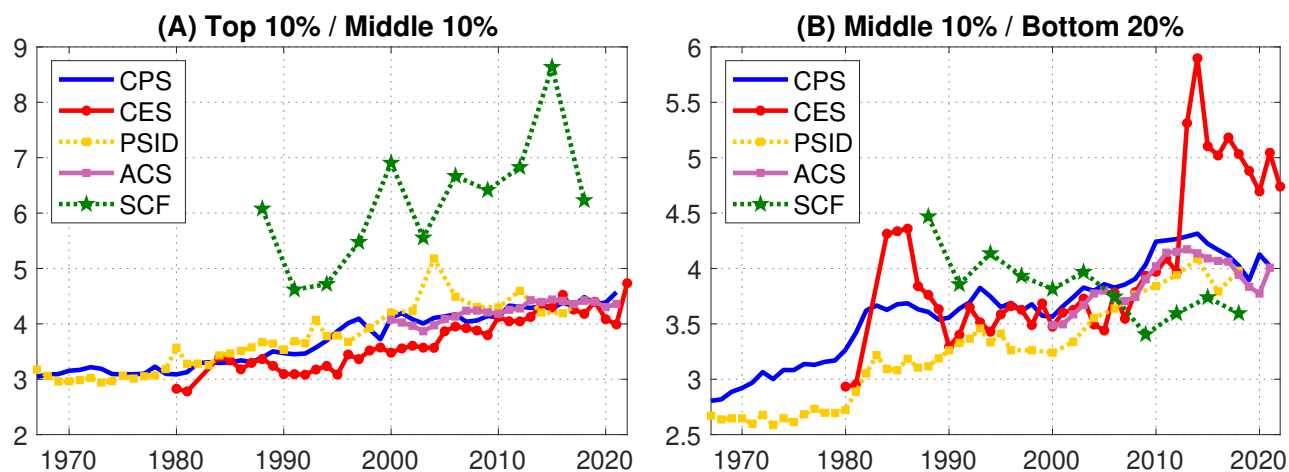


Figure B.7: Cross-datasets Comparison (Pre-tax Income). Panel (A): Pretax Income Inequality at the top. Panel (B): Pretax Income Inequality at the bottom. CPS, PSID, CES, ACS, and SCF, Sample B.