Unequal Growth*

Francesco Lippi
LUISS University, EIEF and CEPR

Fabrizio Perri
Federal Reserve Bank of Minneapolis and CEPR

December 2022

Abstract

Over the past 50 years, households in the United States have experienced changes in earnings dynamics that have resulted in a large increase in inequality. This paper assesses the impact of these changes on aggregate growth and welfare. We begin by inspecting a simple statistical decomposition of aggregate earnings growth using data from the Panel Study of Income Dynamics for the period 1967-2018. The decomposition expresses aggregate earnings growth as the sum of two terms. The first is the covariance between the level and growth of household earnings, which depends only on micro earnings dynamics. The second is the average growth across households, which depends on both micro and macro factors, such as a common labor productivity growth. In order to identify the impact of the changes in the micro dynamics on aggregate outcomes, we map a simple model of micro-founded growth onto the terms of the decomposition. We find that changes in households' earnings dynamics that are consistent with the micro data imply *unequal growth* across the earnings distribution that is, a transition period during which there is a change in the *shape* of the distribution of incomes that is not mean preserving, so that aggregate growth is affected. This change yields a positive effect on aggregate growth and, with incomplete markets, an ex-ante negative welfare effect.

JEL Classification Numbers: D31, O4

Key Words: Income distribution, Inequality, Growth slowdown

*We thank Giovanni Gallipoli, Neil Merhotra, Kjetil Storesletten and Ludwig Straub for insightful discussions, David Argente, Martin Garcia Vasquez, Adriana Grasso and Egor Malkov for outstanding research assistance, our editor Gianluca Clementi for constructive guidance, and seminar participants at several conferences and institutions for excellent comments. Lippi acknowledges financial support from the ERC grant: 101054421-DCS. 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

Over the past 50 years, households in the United States have experienced changes in earnings and income dynamics that have generated a large increase in earnings and income inequality (see, among others, Katz and Murphy 1992 and Heathcote et al. 2010). The objective of this paper is to measure the direct impact of these changes on aggregate growth and welfare.

Our starting point is the observation that aggregate earnings growth can be thought of as coming from two sources. The first is growth that is common (or evenly distributed) across the earnings distribution, such as aggregate productivity growth. This source has, by definition, no impact on the *shape* of earnings distribution or on earnings inequality. The second source is growth that is systematically different across the earnings distribution. This source leads to a change in the shape of the income distribution, and it can affect, at the same time, income inequality and aggregate growth. This is the source that we refer to as "unequal growth". In order to identify unequal growth, we present a statistical decomposition showing that aggregate earnings growth can be written as the sum of two terms: the first is the cross sectional (across households) covariance between earnings growth and earning levels, and the second is the (unweighted) average of household/individual earnings growth.

The key insight is that the cross sectional covariance term is connected to aggregate growth but depends only on micro earnings dynamics, so we can identify changes in these dynamics from changes in this covariance term (and underlying correlations and standard deviations). Once changes in micro dynamics are identified, we can assess their impact on aggregate growth. Moreover, by looking at the evolution of the second term of the decomposition, we can also identify changes in the growth that are common across the distribution.

We first document the evolution of the two terms of the decomposition for the United States using micro data from the Panel Study of Income Dynamics (PSID) over the period 1967-2018. The data show the well known fact that aggregate growth is slowing down (see, among others, Gordon 2012 or Summers 2015) and that inequality is increasing. More importantly for our purposes, the data show that the correlation between earnings growth and levels is negative, reflecting mean reversion, but increasing over time; low-earnings households grow faster than high-earnings households, but the growth gap between low- and high-earnings households shrinks over time.

Second, we bring these data to a simple model of micro-founded growth à la Aiyagari-Bewley-Huggett, modified to include a labor participation margin. In the model, we introduce changes in the parameters governing income dynamics and discipline these changes using standard studies on income micro dynamics and the observed aggregate moments as they appear in the statistical decomposition described above. The idea is closely linked to the analysis by Gabaix et al. (2016), who frame the evolution of income inequality as a transition from one invariant distribution to a new one, triggered by a change in the fundamentals of the household's income process. Our key contribution to the literature is the focus on the impact of these changes on aggregate growth.

The model shows that the changes in micro income dynamics that are consistent with the decomposition involve a sizeable decline in the common growth component plus a changing *unequal growth* across the income distribution that is, a distribution of growth opportunities across the earnings distribution that over the period 1978-2007 has favored high-earnings households more than in earlier periods. We then show that this changing unequal growth has resulted in an increase of aggregate growth over the whole sample period of about 0.3% per year. The intuition for this result is that since high-earnings households make up a large fraction of aggregate earnings, having them grow faster results in higher aggregate growth. Our first conclusion is that the changes in earning dynamics that drove the increase in inequality in the U.S. over the past 50 years have contributed to aggregate growth, partially offsetting the slowdown of the common growth component. We then use the model to evaluate the ex-ante welfare consequences of such a change. Our second conclusion is that in an economy with incomplete markets, the ex-ante welfare effect of the increase in *unequal growth* is negative and sizeable. The reason is that the lower growth of low-earnings households compared with their growth in earlier periods leads to prolonged income stagnation for these households, which in turn leads to non-participation, lower consumption and lower welfare. These losses are only partially offset by the gains at the top.

2 Literature review

There is a large literature that studies the impact of inequality on growth, both empirically and theoretically (see the excellent survey by Benabou 1996). Most of the literature focuses on mechanisms through which the distribution of resources (inequality) affects either endogenous factor accumulation or policies and, through those, growth. Our work shares this literature's interest in the impact of changing inequality on growth, but our focus is on the direct (or statistical) impact of changes in earnings dynamics on aggregate growth and welfare. Our work is also related to studies on the changes in earnings dynamics and their impact on inequality, such as Heathcote et al. (2010), Kopczuk et al. (2010), Atkinson et al. (2011), Guvenen et al. (2014), or Guvenen et al. (2021). On the modelling side, our work builds on the standard incomplete markets framework of Huggett (1993) and Aiyagari (1994), and the calculations of the welfare impact of increased inequality are related to those in Krueger and Perri (2004). Relative to these studies, our main innovation is the introduction of heterogeneous growth opportunities across the earnings distribution. This is inspired by recent works such as Luttmer (2011), Gabaix et al. (2016), or Benhabib and Bisin (2018) showing that heterogeneity in growth is important in explaining features of observed income, wealth, and firm distributions. Our focus on the aggregate impact of individual shocks to agents that are "large," and thus on the macro implication of micro shocks, is connected to the work of Gabaix (2011).

The papers by Jovanovic (2014), Jones and Kim (2018), and Moll et al. (2022) are closely related to the issues we explore. These papers present theoretical models to explain the *joint* evolution of income inequality and aggregate growth. They propose explicit mechanisms through which fundamental changes in the technology, or the market structure, simultaneously trigger a change of the cross sectional income inequality (for the top incomes in Kim and Jones) and a change of the aggregate growth. Jovanovic (2014) presents a model, in which an improvement in the technology for labor market matches between workers with complementary skills leads to a reshuffling of the matches in the labor market, which implies wage gains for the workers at the high end of the skill distribution and losses for the workers in the low end of the distribution.[1] Assuming a lognormal distribution of skills, the model delivers an analytic characterization of the transition dynamics following an improvement of the matching technology, which illustrates the consequences for income inequality and for aggregate growth.[2] In a nutshell, better signals about workers' skills lead to faster growth, more income inequality, and a smaller turnover in the distribution of firms' productivity.[3]

The paper by Jones and Kim (2018) presents a model of the right tail of the income distribution. Assuming an exponential income growth that is occasionally destroyed by the arrival of a new competitor, the model generates an income distribution that is Pareto.[4] Changes in top income inequality reflect changes in the power law parameter that can be triggered by shocks to information technology, taxes, and policies related to innovation blocking. The paper shares with our investigation the focus on the linkages between growth and

[1] The quantitative model by Grigsby (2021), featuring heterogeneous-skill workers, provides an insightful complement to the theory illustrating how non-uniform labor demand shocks may lead to labor relocation and negatively affect the aggregate wage. See Haskel et al. (2012) for a critical review of the hypothesis that increased globalization triggered significant effects on labor income inequality.

[2] A related model by Benabou and Tirole (2016) studies an imperfectly competitive labor market with asymmetric information about heterogeneous worker types. It shows how increased competition for the best talents leads to (a possibly inefficient) increase of income inequality. This paper, however, does not discuss the consequences for aggregate growth.

[3] See Garicano and Rossi-Hansberg (2006) for a related analysis of an economy in which agents organize production by matching with others in knowledge hierarchies. The authors discuss how changes in the cost of communication affect various dimensions of wage inequality.

[4] As usual, this is readily seen from the Kolmogorov forward equation for the distribution of incomes $f(y)$: assuming a growth rate γ and a killing rate δ, the invariant distribution satisfies $0 = \gamma f' + \delta f$, which gives a Pareto distribution with parameter $\alpha \equiv -\frac{\delta}{\gamma}$.

income inequality, focusing on labor and entrepreneurial income (consistent with evidence in Piketty and Saez (2003)). Jones and Kim's (2018) insightful model is designed to inspect the dynamics of the right tail of the income distribution. Our quantitative analysis focuses on the whole range of incomes, as does Jovanovic's (2014), with the aim of capturing the interactions between inequality and growth over the whole range of the income distribution.

The paper by Moll et al. (2022) deals with phenomena closely related to the ones we study. These authors develop a tractable theory that links technological innovations in particular, automation of the tasks performed by labor. They characterize how technological changes affect the returns to capital and labor, as well as inequality. Automation has two effects: it increases inequality by affecting the returns to wealth, and it leads to stagnant wages at the bottom of the distribution. An inelastic supply of capital implies that these effects are persistent, so the wages of displaced workers remain low even in the long run.

3 A micro decomposition of aggregate growth

In this section, we present a simple statistical decomposition that connects aggregate income growth to micro-level (household or individual) income growth, cross sectional income inequality, and the cross sectional correlation between income growth and income level. These types of decomposition have been widely used in industrial organization to connect sectoral productivity growth to productivity growth in individual firms (see, among others, Olley and Pakes 1996). We find it useful to apply this decomposition to household-level data (as opposed to firm-level data), because it connects aggregate growth with household income inequality, which has a more direct and relevant welfare content than firms income inequality.

Let y_{it} be the level of income of household/individual i at time t. Let Γ_{t+T} be the economy's aggregate growth over an horizon T, which is

$$\Gamma_{t+T} = \frac{E(y_{it+T})}{E(y_{it})} = E\left(\frac{y_{it+T}}{y_{it}}\frac{y_{it}}{E(y_{it})}\right)$$

where $E(.)$ is the cross sectional average. Now define

$$g_{i,t+T} \equiv \frac{y_{it+T}}{y_{it}} \quad , \quad s_{i,t} \equiv \frac{y_{it}}{E(y_{it})}$$

so that $\Gamma_{t+T} = E(g_{i,t+T} \cdot s_{i,t})$, where $g_{i,t+T}$ is income growth of unit i and $s_{i,t}$ the ratio between income of unit i and average income. Then, using the definition of covariance and the fact that $E(s_{i,t}) = 1$, we get

$$\Gamma_{t+T} = cov(g_{i,t+T}, s_{i,t}) + E(g_{i,t+T}) \ , \tag{1}$$

or equivalently,

$$\Gamma_{t+T} = corr(g_{i,t+T}, s_{i,t})\sigma(s_{i,t})\sigma(g_{i,t+T}) + E(g_{i,t+T}) \quad . \tag{2}$$

Equation (1) suggests that what matters for aggregate growth is not only the (unweighted) average individual growth $E(g_{i,t+T})$ but also the distribution of growth opportunities, as summarized by $cov(g_{i,t+T}, s_{i,t})$. The intuition for why this is the case is straightforward: the higher the covariance, the faster higher incomes grow; since the high incomes contribute more to aggregate growth, then aggregate growth is higher. Equation (2) also suggests that $cov(g_{i,t+T}, s_{i,t})$ is linked to three cross sectional moments that have an intuitive economic interpretation. The first, $corr(g_{i,t+T}, s_{i,t})$, is the correlation between level and growth at the individual level. This measure captures the degree of mean reversion (or economic rank mobility) in individual income dynamics. The second, $\sigma(s_{i,t})$ is the standard deviation of $s_{i,t}$, which is essentially a measure of cross sectional income inequality. The third, $\sigma(g_{i,t+T})$, is the standard deviation of the growth rate of individual income, which is a measure of cross sectional income volatility. The equation suggests that changes in any of these three quantities will be associated, *ceteris paribus*, with changes in aggregate growth. It is important to note that this decomposition is a statistical identity, so by itself, it cannot be used to make causal inferences on growth and inequality. Nevertheless, it provides a useful starting point for assessing the impact of changing individual income dynamics on growth. To see why, note that all the moments in the first term of equation (2) are independent from the presence of a common growth factor, call it $\bar{g}$, that equally affects the growth of all households. All the terms in the product depend only on heterogeneous individual income dynamics. The second term in equation (2) is instead potentially affected both by the common factor $\bar{g}$ and by individual income dynamics. So the evolution of the statistics in equation (2) will help us, with the aid of a simple statistical model, identify the impact on growth of the changes in income dynamics that drive in income inequality from the changes in growth that are common across all households. For this reason, the next section uses a panel of micro data to document how the terms in the decomposition have changed over time.

4 A decomposition of growth in aggregate earnings in the United States: 1967-2018

Both equation (1) and equation (2) involve cross sectional moments as well as moments related to individual income/earnings growth, so in order to bring them to the data, we need panel data on household/individual income variables. For this reason, we work with the Panel Study of Income Dynamics (PSID), which is a panel of households selected to be representative of the whole population of the United States, collected from 1967 to 1996 at an annual frequency and from 1996 to 2018 at a bi-annual frequency. Our income measure in PSID is total (for each household in the sample) wage and salary income plus farm income plus 50% of business income for each household in the sample, divided by the total number of persons in the sample. The main reason why we focus only on labor income only is that in household surveys (including the PSID), non-labor income is poorly measured, and thus including it in our measure of income makes little difference. We understand that non-labor income might actually be an important driver of income inequality, especially at the top of the distribution (see, for example, Piketty et al. 2018),and believe that an interesting extension of our research would be to use a panel of administrative data (for example, tax returns data) to assess the aggregate growth impact of changes in inequality in non-labor income. Our first check on the quality of PSID data on labor earnings is to compare the growth in aggregate labor income per capita from the full PSID sample with the corresponding measure from the National Incomes and Product Accounts (NIPA), which is wages and salaries disbursement plus 50% of proprietor income, divided by the U.S. population. Both measures are deflated using the PCE deflator.[5] Figure 1 reports aggregate growth in earnings from the PSID and from the NIPA. The solid lines report the actual annualized growth, (computed across five year averages of non-overlapping windows), while the dotted lines are polynomial trends.[6] The figure shows that aggregate growth in the PSID does not match growth in the NIPA perfectly, but the two series show a strong co-movement, suggesting that the PSID sample is a good laboratory for studying the connections between household labor income dynamics and aggregate labor income growth.[7]

[5]See appendix A for more details on data construction.

[6]Because of the bi-annual sample of the PSID, we use only (both for the PSID and the NIPA) the first, third and fifth year of each window. So, for example, the observation for 2018 measures the growth between average income in 2018, 2016, and 2014, and average income in 2012, 2010, and 2008.

[7]The figure shows that earnings in the PSID grow faster than those in the NIPA in the early part of the sample, while they grow slower than those in the NIPA in the last part of the sample. These differences have been documented in other studies, like Cynamon and Fazzari (2017) and Heathcote et al. (2022). As far as we know, a clear cause for them as not been established. We suspect that the main reason is sample selection/attrition in the PSID.

Figure 1: Growth in labor income: NIPA and PSID

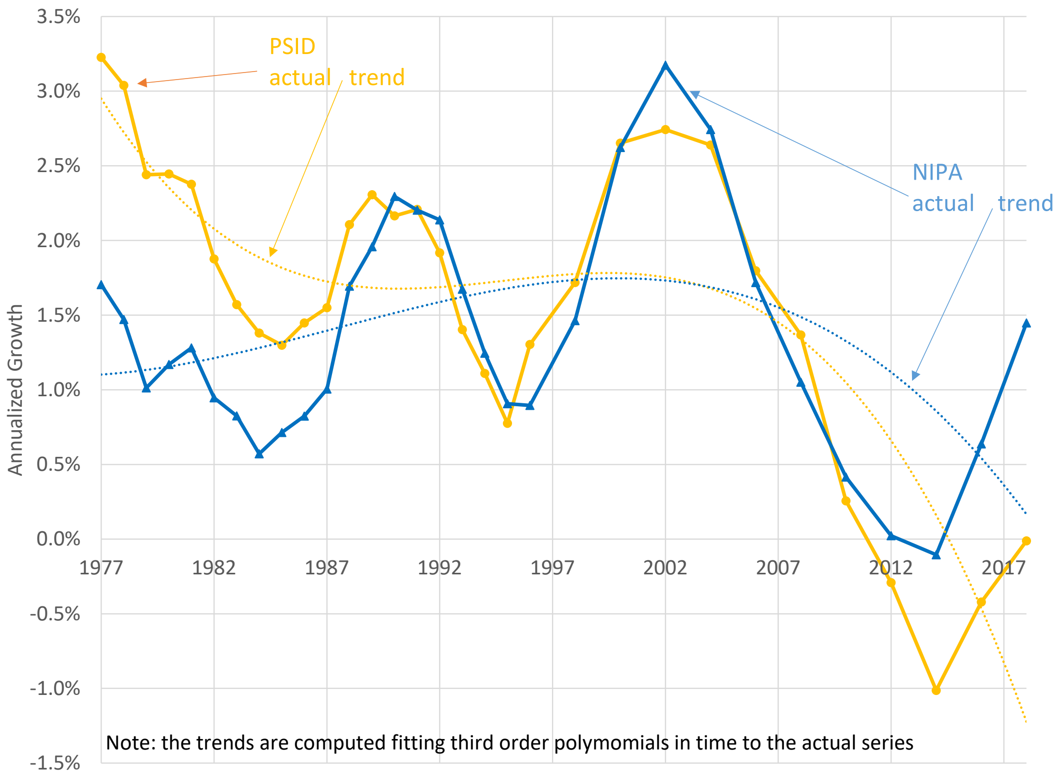

Figure 2: Inequality in household labor income: 90/20 ratio in PSID and CPS

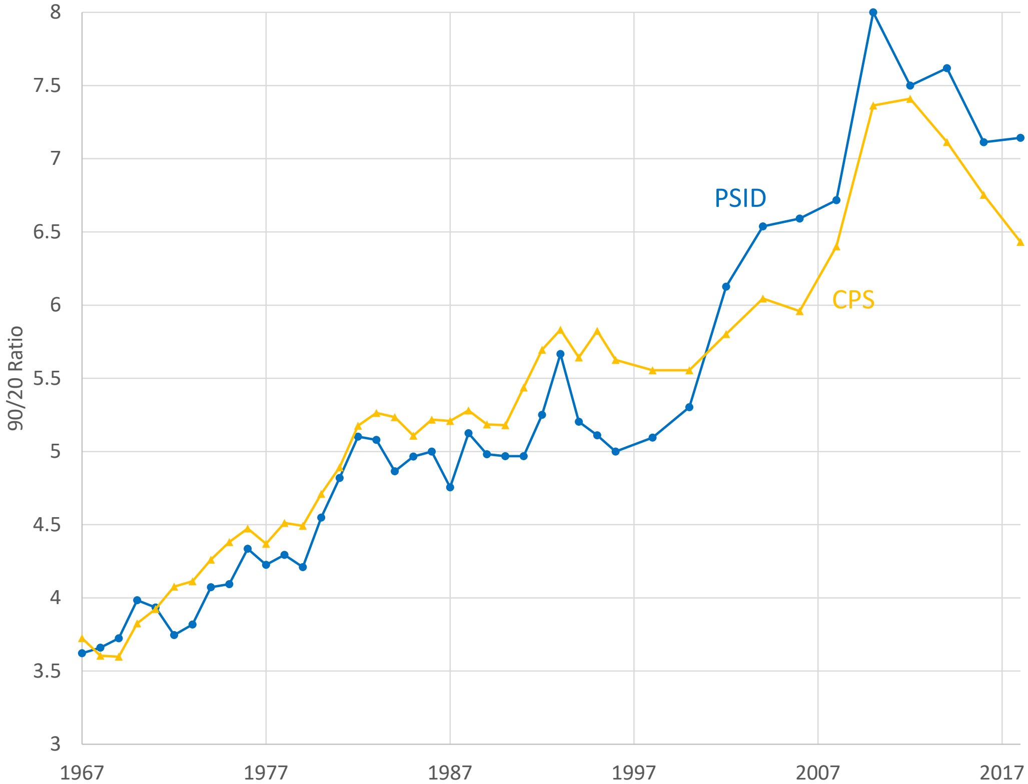

Figure 2 also shows that the PSID captures well the patterns of household income inequality in the United States documented in a much larger cross sectional survey, the March Current Population Survey (CPS). The figure plots a commonly used measure of inequality, the ratio of 90th percentile of the household earnings distribution to the 20th.[8] The figure shows that both surveys are aligned both in terms of level and in terms of the secular increase of earnings inequality in the United States.

Since Figure 1 and Figure 2 show that the data in the PSID capture well the evolution of aggregate growth and inequality, we now proceed to compute the data equivalent in the PSID of s_i and g_i, which are the basic elements of the decomposition in equation (1) and equation (2). In order to reduce measurement error, we aggregate individual PSID data along two dimensions.[9] First, instead of using current labor earnings, y_{it}, we use an average of real (PCE deflated) labor earnings over a five year window, so $\bar{y}_{it} = (y_{it} + y_{it-2} + y_{it-4})/3$

[8]The earnings measure in the CPS and the PSID is total wage and salary income plus 50% of household business and farm income. Inequality measures are computed for households with heads between age 25 and 60. The average sample size in the PSID is around 4,000 households per year; the size in the CPS is 10 times larger.

[9]Guvenen et al. (2014), who also analyze the relation between level and growth in individual earnings data, use a similar aggregation.

is our measure of earnings.[10] Second, in each year, we aggregate households in 10 deciles of $\bar{y}_{it}$. Formally, let I_t be the group of households that is in the i_{th} decile of the $\bar{y}_{it}$ distribution in period t. We define

$$g_{i,t} = \frac{\sum_{j\in I_t} \bar{y}_{j,t+6}}{\sum_{j\in I_t} \bar{y}_{j,t}} \frac{\bar{P}_t}{\bar{P}_{t+6}} \qquad \text{and} \qquad s_{i,t} = \frac{\sum_{j\in I_t} \bar{y}_{j,t}}{\sum_{I_t} \sum_{j\in I_t} \bar{y}_{j,t}}\ , \tag{3}$$

where $\bar{P}_t$ is average population in the PSID sample in periods t, $t-2$, and $t-4$. Our sample includes all households with a head between age 25 and 60 that are in the sample for at least 11 years (from $t+6$ to $t-4$). Note finally that the growth rate of earnings in a given decile is computed using the same group of households in t and $t+6$. In Figure 3, we show the s_{it} and g_{it} for four points in our sample: the beginning (final years of the growth window are 1977-78), two mid-points (final years are 1986-87 and 2006-08), and the end (final years are 2016-18). Starting with the curve in panel (a), we want to highlight three features. The first is that earnings growth is unequal across the earnings distribution, with households at the bottom of the distribution experiencing faster growth. The negative slope of the curve is consistent with a mean reverting income process. The second is that the curve is L-shaped that is, quite steep at the bottom end of the distribution (for $s_i < 1$) and fairly flat at the top of the earnings distribution (for $s_i > 1$). The third feature is that the support of the curve is fairly concentrated, with income of the top decile being only twice the income of the middle decile. Moving to the middle periods (panels b and c), first notable change is that the curve becomes U-shaped, with growth of the top decile being faster than the growth of the middle deciles. This faster growth at the top results in a widening of the support of the earning distribution (i.e. ,increasing inequality). Finally, notice that the curve shifts down over time suggesting a reduction in growth for most deciles. Panel d shows that in the last years of the sample (growth ending in 2016-18), the curve returns L-shaped, with a more noticeable spike of growth at the bottom.

After showing the evolution of unequal growth in the United States, we provide some evidence on the connection between unequal growth, inequality and aggregate growth. The top panel of Figure 4 shows two components of the decomposition in equation (2). The solid line depicts the standard deviation of s_i, a measure of income inequality, while the dashed line depicts $corr(s, g)$, the correlation between income level and income growth. The panel shows that there is co-movement between earnings inequality $\sigma(s)$ and the correlation between earnings level and earnings growth. There are two periods in the sample (highlighted by the shaded areas, which correspond to panels b and c in Figure 3) when the correlation

[10] The reason why we don't use all years in the window is that PSID data are bi-annual after 1996, so this is the only way to obtain a consistent measure of $\bar{y}_{it}$ throughout the sample.

Figure 3: Evolution of unequal growth in the United States: 1977-2018

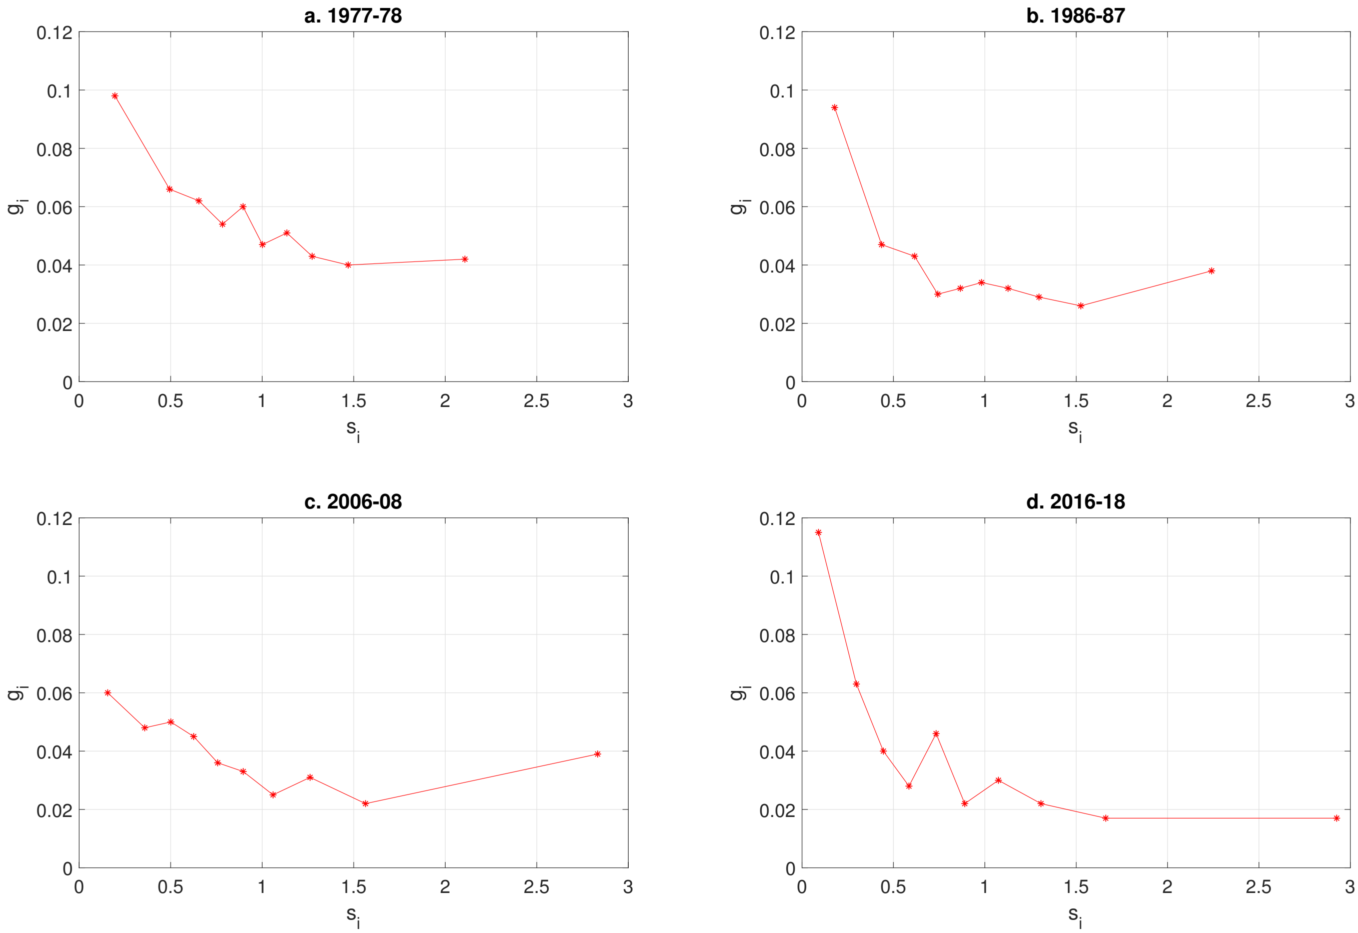

between level and growth peaks. These periods are associated with large increases in income inequality. As discussed above, an increase in the correlation between earnings level and growth implies that high-earnings households tend to grow faster and hence increase earnings inequality.

The most interesting insight for our purpose comes from comparing the top with the bottom panel of Figure 4, which reports the term Γ_t in equation (1), measured aggregating all households in our PSID sample.[11] Comparing the two panels shows that during both

[11]It is important to note that the aggregate growth rate in the bottom panel of Figure 4 is higher than the one reported in Figure 1 (see Figure B.3 in Appendix B). The reason for the difference is the selection of the sample. In Figure 1 aggregate growth is computed using all households in the PSID, as we want to compare the PSID with the NIPA. In Figure 4 aggregate growth is computed only from households that are in the panel for at least 11 years, and in each of those years, the head is between 25 and 60 years old. This implies that households in the sample used in Figure 4 are younger than the full sample and thus tend to have faster earnings growth. However Figure B.3 shows that the growth decline and the medium run fluctuations in growth are similar across the two samples.

episodes, when the correlation between income and growth peaks, we also observe higher aggregate growth (or a reduction in the growth decline). This evidence suggests that changes in unequal growth can drive, at the same time, increases in earnings inequality and changes in aggregate growth. We want to stress here that aggregate growth can obviously be affected by other factors, which will show up in the term $E(g_i)$ in equation (2), so we should not expect a perfect correspondence between aggregate growth and $corr(s, g)$.[12]

So far, we have documented a series of facts relating growth and inequality in the United States over the past 50 years. Aggregate growth has declined and inequality has increased. The decline in growth has not been uniform across the income distribution, and in the middle years of our sample, we have documented faster earnings growth of households at the very top of the distribution (relative to that of households in the middle). Towards the end of our sample, on the other hand, we observe faster earnings growth of households at the very bottom of the distribution. In the following section, to identify the main changes in the process for earnings, we present a simple model of household earnings formation and use it to match the facts documented above. The model allows us to measure the impact of changes in earning dynamics on aggregate growth and to quantify its effect on the welfare of households in the United States.

5 A Bewley-Aiyagari-Huggett model

We consider a standard Bewley-Aiyagari-Hugget small open economy, with a few simple modifications to the household income process introduced to capture the features and the changes in the earnings distribution documented above. We then explore the effect of these changes on aggregate growth and on welfare. Time is discrete, and we take a period in the model to be a quarter. The reason why we choose a quarterly frequency is that an important component of household earnings risk is the possibility of brief (shorter than one year) zero-earnings spells. These spells can be captured only by a model with a frequency that is higher than a yearly one. The economy is inhabited by a continuum of infinitely lived households with standard preferences over consumption flows, denoted by

$$E_t \sum_{j=0}^{\infty} \beta^j u(c_{t+j}) ,$$

where $\beta > 0$ is the discount factor and $u(.)$ is a standard utility function, which is assumed to be $CRRA$, that is $u(c) = \frac{c^{1-\theta}}{1-\theta}$ with $\theta > 0$.

[12] Figure B.1 and B.2 in Appendix B report the evolution of all the components of the decomposition in equations 1 and 2.

Figure 4: Level-growth correlation, inequality and aggregate growth

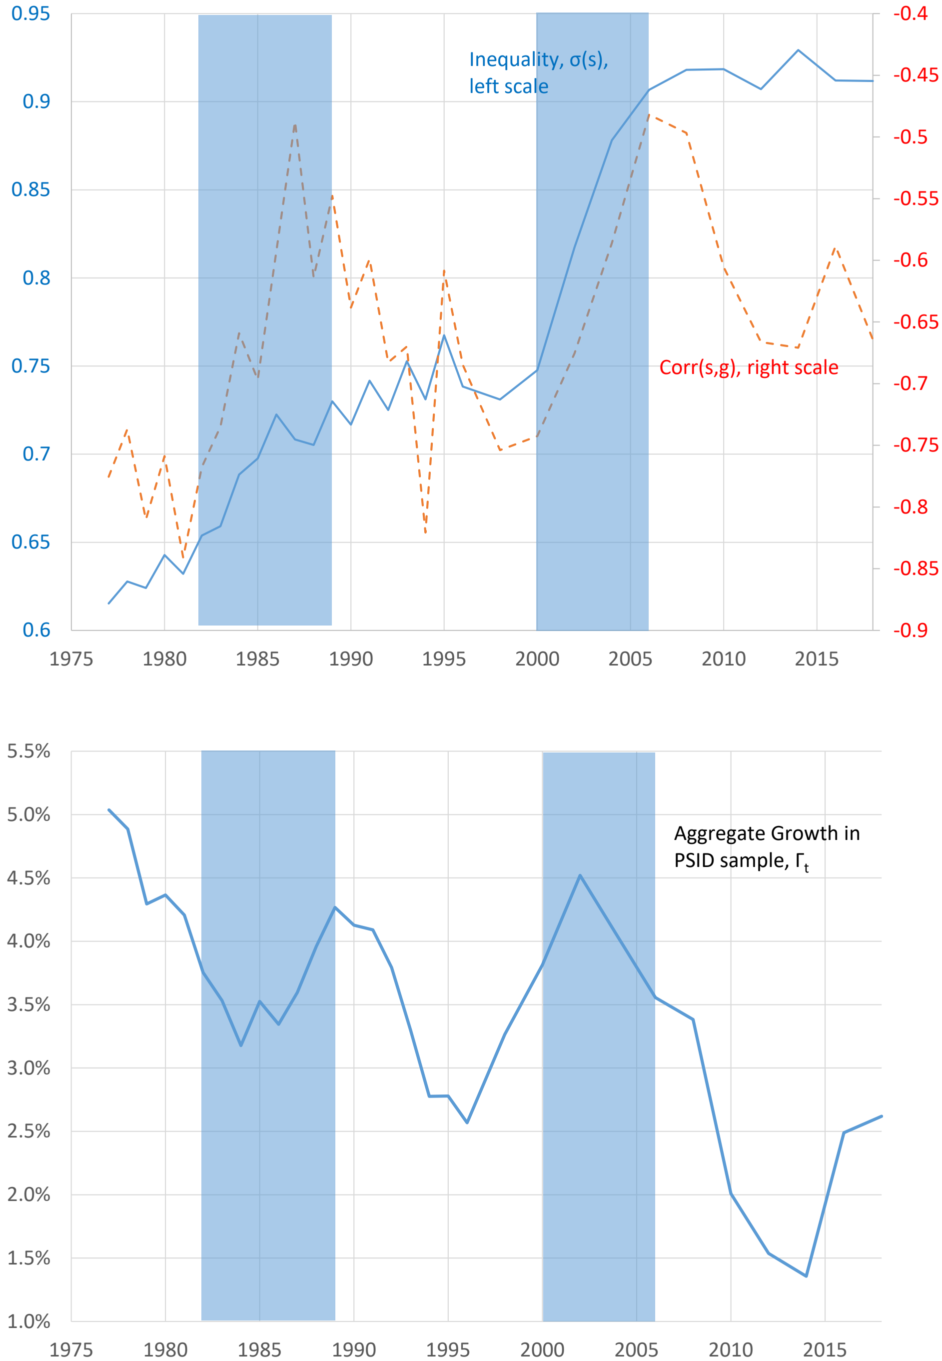

5.1 Earning potential

Each household in each period receives an idiosyncratic realization of its earning potential Y_{it}. We model earning potential as

$$\log Y_{it} \equiv y_{it} = \alpha_i + f_{it} + e_{it}. \tag{4}$$

The first two components are meant to capture permanent differences in earnings potential across households, so we define $p_{it} \equiv \alpha_i + f_{it}$ and $\tilde{s}_{it} \equiv \frac{e^{p_{it}}}{E_i(e^{p_{it}})}$ to be the relative position in terms of permanent earnings potential of household i. The third term, e_{it}, denotes an autoregressive income component described below.

The first component, α_i, is a standard fixed effect meant to capture initial permanent differences in earnings potential across households. We assume

$$\alpha_i \sim N(0, \sigma_\alpha) \ .$$

The second component of the earnings potential process, f_{it}, which we name the growth factor, is going to be the driver of the increase in income inequality. It evolves according to

$$f_{it} = f_{it-1} + \bar{g}_t \ + \delta_t \frac{(\tilde{s}_{it-1} - 1)}{(1 + \tilde{s}_{it-1})} \ . \tag{5}$$

The important element in equation (5) is that the earnings growth of household i depends on $\tilde{s}_{it-1}$. First, consider the case in which $\delta_t = 0$. In this case, each household experiences a common earnings growth rate $\bar{g}_t$. In our experiments, this is going to be the relevant case in the initial and final steady state. During a transition period, however, we will allow the parameter δ_t to be different from 0 and to be positive in particular, so that households with permanent earnings above the mean ($\tilde{s}_{it-1} > 1$) can have faster growth than households with permanent earnings below the mean. As we will show below, when $\delta_t > 0$, inequality is increasing, so this will be our modelling device to obtain the observed increase in earnings inequality. One possible structural interpretation of the two components α_i and f_{it} is that α_i captures the value of the initial skill endowment of household i and f_i captures the changing value of this skill (see, for example, Lochner and Shin 2014).

The final component e_{it} is a standard autoregressive process, which we model as

$$\begin{aligned} e_{it} &= \rho e_{it-1} + \varepsilon_{it} \quad , \quad \varepsilon_{it} \sim N(0, \sigma^2_\varepsilon(\tilde{s}_{it})) && (6) \\ \sigma^2_\varepsilon(\tilde{s}_{it}) &= \frac{\sigma^2_\varepsilon}{\tilde{s}^\chi_{it}} \ . && (7) \end{aligned}$$

Note that the parameter χ links the volatility of shocks of the income process, $\sigma_\varepsilon^2(\tilde{s}_{it})$, to $\tilde{s}_{it}$, the position of household i in the permanent earnings distribution. When $\chi = 0$, shocks have the same variance across households; when $\chi > 0$, poorer households have higher volatility of earnings shocks. This modelling choice is motivated by a large body of research that has documented that households at the bottom of the income distribution face higher volatility in their earnings shocks (see, among others, Meghir and Pistaferri 2004).

5.2 Work choices and earnings

In each period, each household with earning potential Y_{it} has the option to work and earn its potential minus taxes, or not work (and thus have 0 earnings) and receive a transfer income ϕ_t, which changes over time.

When households work on the market, they pay a flat tax that the government uses to finance the transfer income. The process for earnings (before transfer and taxes) of household i, which we denote by $h(Y_{it})$, is thus given by

$$h(Y_{it}) = \begin{cases} Y_{it} & \text{if } Y_{it}(1-\tau) \geqslant \phi_t \\ 0 & \text{if } Y_{it}(1-\tau) < \phi_t \end{cases} .$$

In a given quarter, this feature of the model will generate households with positive earnings as well as households with no earnings.

5.3 The household problem

The household consumption saving problem is standard. In the baseline case, we assume incomplete markets so that each household can borrow and save using an uncontingent bond, which pays an exogenously given, potentially time varying interest rate r_t. Bond holdings have to be above a borrowing constraint $\bar{b} \leq 0$. The problem can then be written as

$$\begin{aligned} & \max_{c_{t+j}, b_{t+j}} E_t \sum_{j=0}^{\infty} \beta^j u(c_{t+j}) \qquad (8) \\ & s.t. \\ c_{t+j} \;=\; & b_{t+j-1}(1+r_t) + \max((1-\tau)h(Y_{it+j}), \phi_{t+j})) - b_{t+j}, \\ & b_{t+j} \geq \bar{b} \qquad b_{t-1} \text{ given.} \end{aligned}$$

5.4 Equal growth stationary equilibria

We first analyze stationary equilibria in which there is no unequal growth ($\delta = 0$) and in which all parameters, including the aggregate growth rate of the economy $\bar{g}_t$, the interest rate r_t, and the transfer income ϕ_t, are constant. An *equal growth* equilibrium is a distribution of households over earning potential and assets, $\mu(Y, b)$, plus a household decision rule $b'(b, Y)$ and a tax rate τ satisfying the following conditions:

1. The decision rules solve the household decision problem in equation (8).
2. The distribution $\mu(Y, b)$ induced by the decision rules of the households and law of motion Y is time invariant.
3. The government budget constraint is satisfied:

$$\int \tau h(Y) d\mu(Y, b) = \int \phi I(h(Y) = 0) d\mu(Y, b) \ ,$$

   where $I(.)$ is the indicator function.

Note that in an *equal growth* equilibrium, all individual and aggregate variables grow at the constant rate of $\bar{g}$; hence, when we solve for it, we solve for the equilibrium in an economy in which all variables are detrended by the growth factor f_t and the discount factor β and the interest rate on bonds $1 + r$ are suitably rescaled.[13]

5.5 Unequal growth equilibria

We label *unequal growth equilibria* those that arise during a transition from one equal growth stationary equilibrium to another. We assume the economy starts in a stationary equilibrium with aggregate growth rate $\bar{g}$ and at time t_0, it experiences a sequence of changes in parameters for N periods. Specifically, we will consider the case in which $\delta_t > 0$ and in which $\bar{g}_t$, r_t, and ϕ_t change linearly for $t \in [t_0, t_0 + N]$. Changes in these parameters are unexpected by the agents in each period, and we assume that in each $t_i \in [t_0, t_0 + N]$, when agents look forward they expect $\delta_t = 0$ and $\bar{g}_t = \bar{g}_{t_i}, r_t = r_{t_i}$, and $\phi_t = \phi_{t_i}$ for $t > t_i$. After period $t_0 + N + 1$, we assume that the economy settles to a final constant growth rate $\bar{g}_F$, interest rate r_F and transfer income ϕ_F and that $\delta_t = 0$. An unequal growth equilibrium is a sequence of distributions $\mu_t(Y, b)$ and a sequence of decision rules $b'_t(b, Y)$ and taxes τ_t for $t \in [t_0, \infty]$, satisfying the following conditions:

[13] In particular, the interest rate in the detrended economy is equal to $\frac{1+r}{1+\bar{g}}$, and the discount factor is equal to $\beta(1 + \bar{g})^{(1-\theta)}$

1. Given expectations on future parameter changes, the decision rules solve the household decision problem (8).

2. The sequence of distributions are consistent with the decision rules and law of motion for Y

3. The government budget constraint is satisfied in every period:

$$\int \tau_t h(Y) d\mu_t(Y,b) = \int \phi_t I(h(Y)=0) d\mu_t(Y,b) \ .$$

4. As t goes to ∞ the distribution $\mu_t(Y,b)$ converges to a time invariant distribution $\mu_F(Y,b)$

5.6 Calibration

In this section we first describe how we parameterize the economy in the initial equilibrium with equal growth. We then discuss how we modify the parameters during the unequal growth transition.

Equal growth equilibrium

Table 1 summarizes our parameter values for the initial equal growth equilibrium, which we calibrate to match features of the earnings distribution in the PSID in the late 1960s and mid-1970s, before the increase in inequality started. The first six parameters of the table are chosen so that the relation between earnings level and growth (g_i and s_i) in the model matches the one in the PSID data averaged for the first two years of our sample (1977 and 1978) for which we can construct measures of g_i and s_i. Figure 5 illustrates the matching between data and model.[14] Note that a crucial element of the earnings process that allows the model to match the data is the variance of autoregressive shocks to be declining with the level of earnings. To understand why this is the case, note that the curve in the data in Figure 5 is fairly flat at the top of the distribution (for $s > 1$), while steep and downward sloping at the bottom of the distribution (for $s < 1$). The autoregressive component (e_{it}) of the earnings process with a variance that is constant across the distribution (i.e. $\chi = 0$ in equation 7) generates a downward sloping line (because of mean reversion), while the fixed earnings level effects (α_i) generate a flat line along the whole distribution (because growth

[14] In order to produce Figure 5, in the model we simulate quarterly earnings paths for a large number of households. Then, we aggregate quarterly earnings into annual figures, and finally we aggregate across years and across deciles, exactly as we do for the data in the PSID

Table 1: Parameters in the initial stationary equilibrium

Income Process Parameters		
Name	Symbol	Value
Variance of fixed effects	σ_α	0.45
Annualized persistence of AR shocks	ρ	0.92
Baseline st.dev. of AR shocks	σ_ε	5.5%
Standard deviation gradient	χ	1
Annualized common growth	$\bar{g}$	4.6%
Transfer income (% of average Y)	ϕ	26%
Tax rate	τ	1.0%
Unequal growth	δ	0
Preference Parameters		
Annualized discount factor	β	0.98
Risk Aversion	θ	1
Other Parameters		
Borrowing constraint	$\bar{b}$	0
Annualized interest rate	r	5%

is independent of levels). A process that is the sum of two processes generates a line with a constant negative slope across the distribution, that cannot match the shape of the data curve at both the bottom and top of the distribution. However, when the variance of the shocks of the autoregressive component declines with the level of earnings (i.e. $\chi > 0$ in equation 7), the fixed effects are the main drivers of earning dynamics at the top of the distribution, while the autoregressive component drives the earnings dynamics at the bottom of the distribution. In this case, the model can generate a g_i/s_i curve that is flat at the top and steep at the bottom, thereby matching the data, as shown in Figure 5.

Another important parameter that allows the model to match the data well is ϕ, which determines the fraction of households with 0 earnings. The parameter is what allows the curve in the model in Figure 5 to match the spike in growth for the bottom decile of s_i in the data. The value of ϕ that we choose also implies a share of households with 0 earnings in a given quarter of approximately 4%, which is plausible for our sample of prime-age households.

Finally, we set the (annualized) interest rate on bonds r to 5%, the annualized discount factor β to 0.98, the borrowing constraint $\bar{b}$ to 0, and the risk aversion θ to 1 (log preferences). These parameter choices jointly imply that in the initial steady state, the aggregate wealth to income ratio in the model economy is around 2.

Figure 5: Initial steady state: data and model

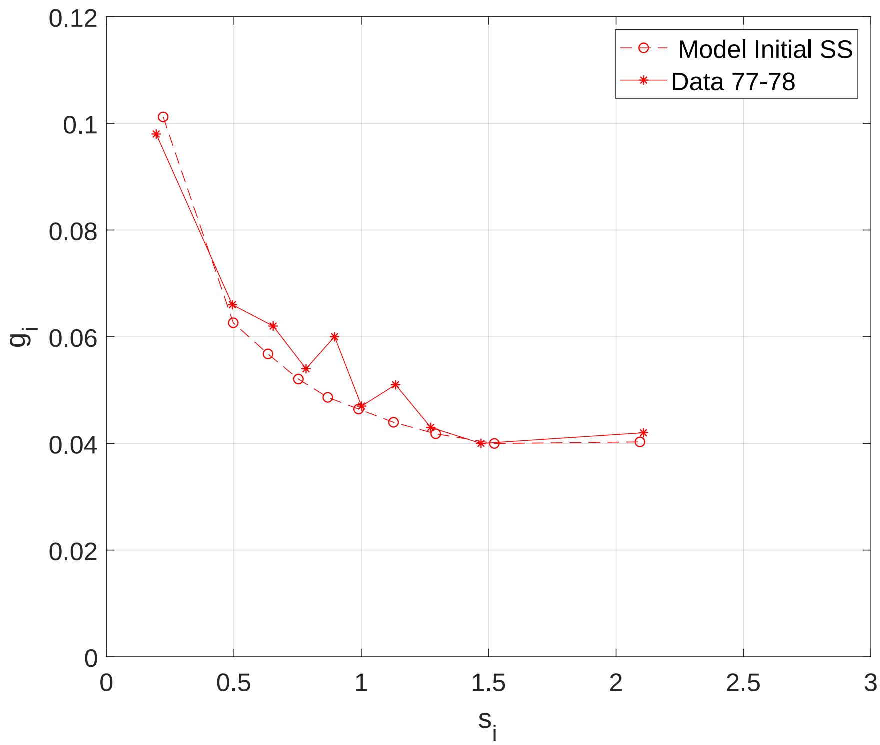

Unequal growth equilibrium

Once we have calibrated the model to the initial steady state, we consider a transition period. We assume the parameter δ increases from 0, its value in the initial steady state, to 0.009, for a period of 30 years. This choice, together with equation 5, implies, for example, that during that period, a household with permanent earnings that are twice the mean ($\tilde{s}_i = 2$) experiences earning growth 1.2% per year faster than a household with earnings at the mean ($\tilde{s}_i = 1$). After 30 years, the parameter δ reverts to 0. This parameter change is chosen so that the model exactly replicates the increase in earnings inequality (the increase in standard deviation of the s_i) documented in the top panel of Figure 4. This parameter change is the only device that will generate changes in the cross sectional distribution of potential earnings in our model, and thus it is the main driver of our results.

Before comparing the outcome of the model with the data, we also discuss the other parameter changes we implement during the unequal growth period. The first change is a linear reduction in the transfer income parameter ϕ from 26% of average income in the initial steady state to about 10% of average income in the final steady state. With constant transfer income, in response to the increase in unequal growth, the model would generate an increase in non-participation so large that the first decile of earnings would go to 0. The reason is that with unequal growth, low earnings households are more likely to experience low growth in potential earnings, which leads to non-participation and 0 earnings. For this reason, we set a reduction in transfer income, so the fraction of households with 0 earnings remains constant across time and equal to 4%, which is the participation rate in the initial steady state.[15] We view this decline in transfer income as a reduced form way to capture an increasing incentives for working, which is particularly relevant for women in the first part of our sample, in face of declining earning prospects at the bottom of the distribution.

The second change is a reduction in the common growth factor $\bar{g}$ from 4.6% to 1.7%, chosen so that the model matches the reduction in aggregate growth, as displayed in the bottom panel of Figure 4. Our estimate of the reduction in the common factor of earnings growth is large, suggesting that changes such as technological slowdown (see, for example Gordon 2012), or the decline in labor share (see, for example, Elsby et al. 2013) have had an important effect on the evolution of the growth in labor earnings in the United States. Our results below suggest that this effect has been partly muted by the unequal growth in earning dynamics.

The final change is connected to the second, and it involves a decline in the interest rate. Note the substantial reduction in aggregate growth implies that agents face a stronger desire to save (or less desire to borrow) which, with a constant interest rate, would generate an implausibly large increase in the wealth to income ratio. For this reason we assume that the interest rate falls from 5% in the initial steady state to a level of 2% in the final steady state. The final level of the interest rate is chosen so that, in the absence of unequal growth, the aggregate wealth to income ratio stays constant around 2.

5.7 Results

Figure 6 and 7 show several statistics implied by the model and contrast them with the same statistics in the PSID data. Figure 6 shows that the earnings process with unequal growth captures the change in shape of the level-growth curve, from an L-shape in the initial period (panel a), to a U-shape in the middle periods (panels b and c) and back to a L-shape in the final period (panel d), with a spike in growth for low levels of earnings. Figure 7 shows also that the model produces a path for the correlation (panel b) and the covariance (panel c) between level and growth which is initially increasing and, towards the end of the sample, declining. To understand why the model generates these patterns, it is useful to divide our sample in two sub-periods. The first is 1978-2007, which is when the earnings process involves unequal growth (i.e. $\delta > 0$), and the second one is 2008-2018, when there is no

[15] Since the transfer income declines and the fraction of recipients stays constant, the tax rate needed to balance the budget falls from 1% in the initial steady state to 0.4% in the final steady state.

longer unequal growth (i.e. $\delta = 0$). In the first sub-period unequal growth favors the growth of high income households, it increases inequality (panel a), as well as correlation (panel b) and covariance (panel c) between earnings levels and growth. Note that the increase in the correlation predicted by the model is more pronounced than the increase in the covariance, as the increase in $\sigma(s)$ and the fact that $corr(s, g) < 0$ mute the increase in the covariance. In the second subperiod there is no longer unequal growth, yet panel (a) in Figure 7 shows that inequality keeps increasing. The reason is that households at the bottom of the distribution now have a lower $\tilde{s}_i$ and because of this they face more volatile shocks (see equation 7) to the autoregressive component of earnings e_{it}. Since e_{it} are persistent, it takes time for their initial distribution (which is realized at the end of the unequal growth period) to converge to the more dispersed stationary distribution implied by the more volatile shocks. During this transition earnings inequality keeps increasing. This increase in dispersion at the bottom of the distribution yields a larger number of households which transit between non participation (and thus 0 earnings) and participation, and that generates the spike in earnings growth rate, observed in panel (d) of figure 6, for the households at low earnings levels. The spike in growth rate at the bottom of the earnings distribution in turn explains the increase in the dispersion of earnings growth rates (observed toward the end of the sample in panel (d) of Figure 7), and the decline in the correlation and the covariance between level and growth (Figure 7, panels b and c), which are no longer pushed up by the presence of unequal growth. The conclusion from this comparison between model and data is that an increase in dispersion of the permanent component of household earnings, increase which we label "unequal growth", captures well many features of earnings dynamics observed in PSID. Alternative choices for modelling the increase in earnings inequality, such as increase in the volatility or persistence of the autoregressive component of earnings, are not able to replicate equally well the features of the PSID data we have documented.

Our final result involves assessing the aggregate impact of unequal growth. Panel (f) in Figure 7 shows that along the unequal growth equilibrium aggregate growth declines. However along the unequal growth equilibrium the common growth factor $\bar{g}$ declines from 4.6% to 1.7%. To isolate the impact of unequal growth (that is $\delta_t > 0$) we simply compute aggregate growth in an unequal growth equilibrium where $\delta_t = 0$ for all t and subtract it from the aggregate growth in the baseline unequal growth equilibrium. Figure 8 plots this difference. The figure shows that unequal growth increase aggregate growth. Overall it can account for an increase in aggregate earnings growth along the transition of an average 0.3% per year. The logic of this result is that unequal growth, by increasing the growth of high earnings households, increases the covariance between level and growth, which results in higher aggregate growth.

Figure 6: Relation between earnings level and growth: Data and model

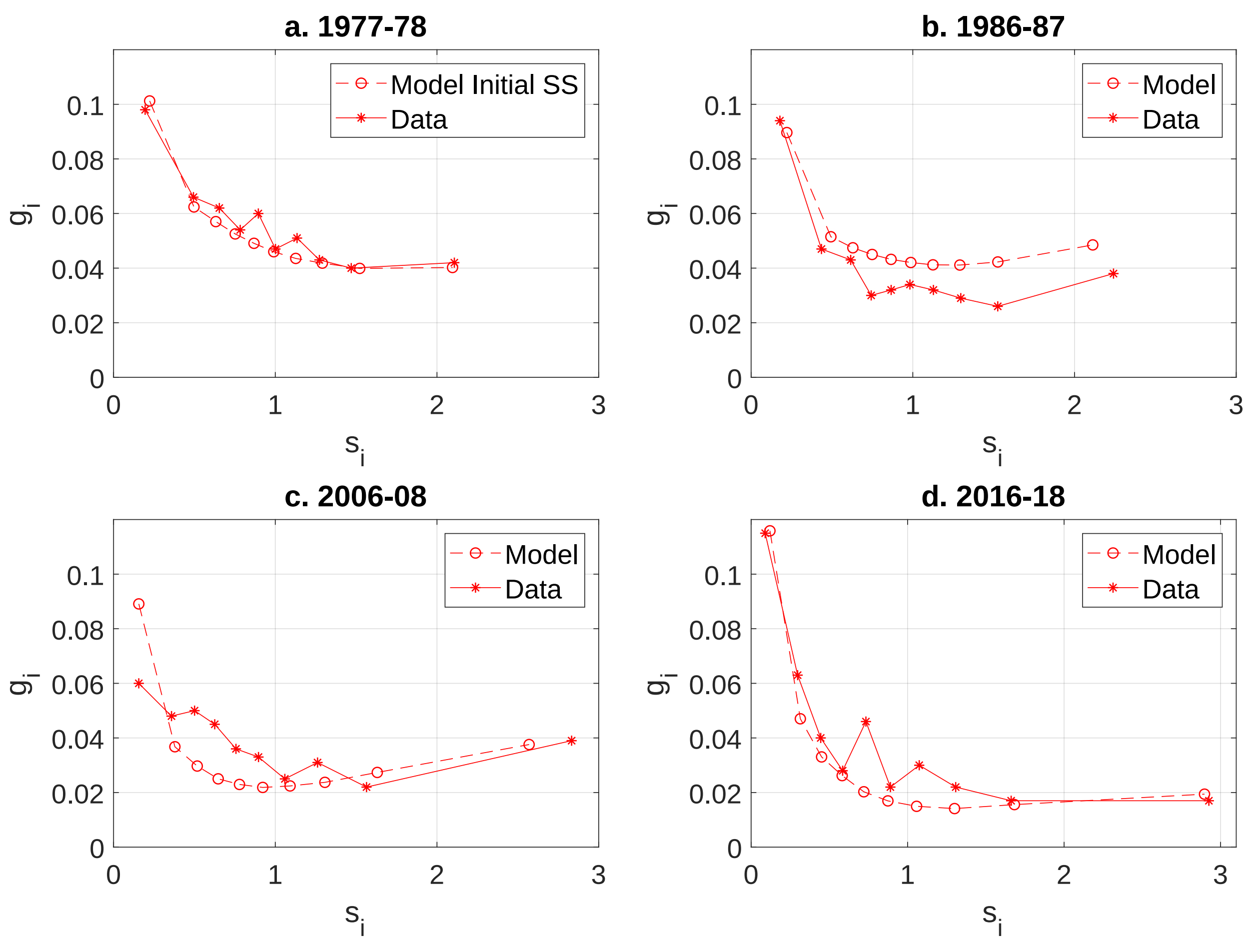

5.7.1 Welfare

Even though "unequal growth" increases aggregate growth, its welfare impact on households is not obvious. We now proceed to analyze of the welfare impact of "unequal growth". As expected, the impact depends crucially on two factors: the curvature of the utility function, which in this class of models captures the social cost of consumption inequality, and the degree of market incompleteness. In Table 2 we measure the welfare cost (in lifetime consumption equivalent units) of moving from the baseline unequal growth equilibrium to an unequal growth equilibrium where $\delta_t = 0$ for all t. In other words, the entries in the table measure the percentage of lifetime consumption a household under the veil of ignorance is willing to give up to avoid the changes in δ_t that generate unequal growth. We consider two values of the risk aversion ($\theta = 1$, which is log utility, and $\theta = 2$) and three market structures: the bond economy, (BE, the economy described above), a complete markets economy (CM), and financial autarky (FA). In financial autarky, households simply consume their (after tax

Figure 7: Time paths: data and model

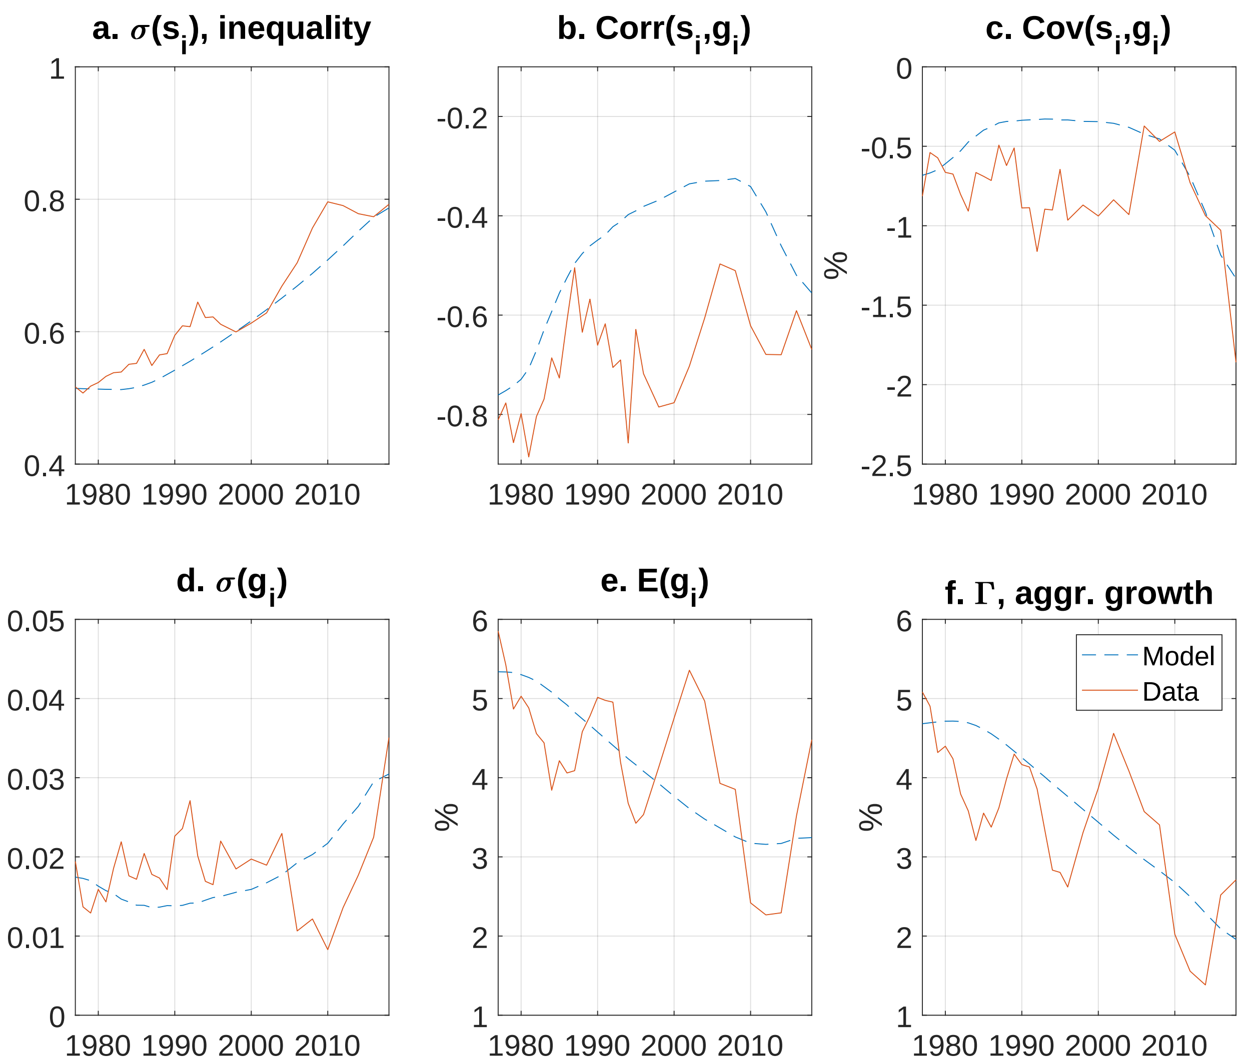

and transfer) earnings. In the complete markets economy, households still face a borrowing constraint but can sign contracts that fully insure them against their idiosyncratic earnings risk (including changes in unequal growth) and allow them not to work when their potential earnings are below ϕ_t (so that aggregate earnings under complete markets are the same as in the other market structures). Since parameters are such that $\beta\frac{1+r_t}{1+g_t} < 1$ for each t, and all the risk is insured, under complete markets all households prefer current consumption to future consumption. This implies that in each period, each household consume all its income, which is, because of insurance, a fixed fraction of aggregate earnings.

The first column of the table shows that under complete markets, unequal growth produces welfare gains. The logic is that the benefits of the higher aggregate growth displayed in Figure 8 are shared among all households. Welfare gains are declining with more curvature in

Figure 8: Aggregate impact of unequal growth

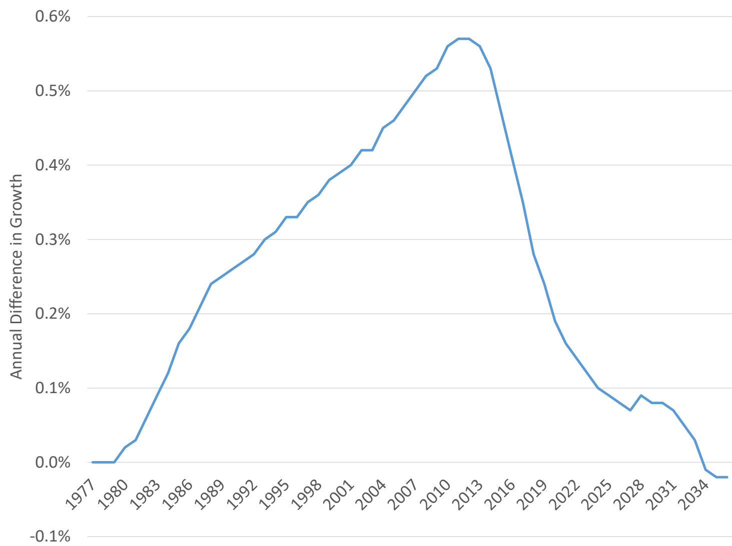

Table 2: Welfare costs of unequal growth

	Market Structure		
Risk aversion (θ)	CM	BE	FA
$\theta = 1$	-9.9%	+2.7%	+6.2%
$\theta = 2$	-5.3%	+7.4%	+15.2%

utility, as the higher curvature implies the extra resources brought in by unequal growth are valued less. When markets are incomplete, however, high-earnings households benefit, and low-earnings households lose. This results in ex-ante welfare losses, which, with curvature θ equal to 2, can be very substantial. It is useful to think of the losses in incomplete markets as arising from two features. The first is that poor agents experience lower growth and thus are stuck with permanently lower component of their earnings. The other is that with lower income, they also experience more volatile shocks. In financial autarky, both of these features affect welfare negatively, hence the large welfare losses. In the bond economy, agents can (partly) insure against the more volatile shocks but still suffer the adverse consequences of the permanently lower component of income, which explain why the welfare losses in the bond economy are also large. Another way to understand these large welfare losses in the bond economy is that the process of unequal growth causes increases in the dispersion of "permanent income" (see Bowlus and Robin 2004, Abbott and Gallipoli 2022 and Straub 2019), which translate into dispersion in consumption and welfare losses. Note also that when the curvature is high ($\theta = 2$), the gap in the welfare impact of unequal growth between complete markets and incomplete markets exceeds 20% of lifetime consumption. This is not surprising, but it highlights that a period of unequal growth increases the social value of better risk sharing or social insurance mechanisms. This could be done by introducing, in response to unequal growth, better credit/insurance markets and/or an increase in the transfer income ϕ_t. We view these as promising avenues for future research, but in order to explore them, one would need a more careful model of credit markets and a model of the incentive effects of transfers and taxes needed to finance them. In the context of our economy, we have considered a very simple experiment, which is to introduce unequal growth in an economy where households have a more slack borrowing constraint, set to 50% of average income, and all other parameters are left unchanged. We find that in this economy, the welfare effects of unequal growth are virtually indistinguishable from the effects in the benchmark economy with the borrowing constraint set to 0. We conjecture that this is because in the initial steady state of the economy with a more slack borrowing constraint, low earnings agents, which are the most negatively affected by the unequal growth, have, on average, less assets than in the economy with a tighter constraint, so the distance between their assets and the borrowing constraint is similar across the two economies. This implies that they are impacted by unequal growth in similar way, suggesting that with respect to mitigating the adverse impact of unequal growth, it is not the absolute availability of credit that matters but rather the emergence of better credit and insurance mechanisms as the unequal growth materializes.

5.7.2 Unequal growth and the wealth distribution

In this section, we briefly discuss the impact of unequal growth on the wealth distribution. In our benchmark calibration, the decline in aggregate growth faced by the small open economy is matched by a decline in the interest rate; this assumption, together with log utility, implies that in the absence of changes in unequal growth, the wealth distribution stays unchanged during the transition, as well as in the final steady state. This allow us to attribute all the changes in wealth to changes in unequal growth. Our main numerical result is that after the period of unequal growth has finished, the wealth to income ratio in the final steady state is lower relative to the one in the initial state, and wealth inequality is higher. In particular the wealth to income ratio declines from 2 to 1.7, and the share of wealth held by the top 10% increases from 31% to 34%. The result implies that in a closed economy unequal growth would cause an increase in interest rates. To understand the result it is useful to first note that the earnings process in the final steady state is the same process as in the initial steady state, with more dispersed fixed effects α_i. The effect of unequal growth can indeed be described as a spreading of the fixed effects around their mean (see equation 5). In an economy with more dispersed earnings fixed effects obviously earnings/income inequality is higher, and since household choices imply a positive association between earnings and wealth, wealth inequality is also higher. However the impact on the wealth to income ratio is not obvious. We conjecture that the feature that makes the wealth to income ratio decline is that earnings risk declines with the relative position in the distribution of fixed effects (equation 7). As we have discussed above a key driver of the increase in aggregate income/earnings is the faster (over time) earnings growth of households at the top the earnings distribution. However these high earnings households also experience a decline in earnings risk, hence they reduce their saving rate and the growth in their wealth can be lower than their income growth. For this reason, households at the top of the distribution can be responsible for a reduction in the wealth to income ratio.[16] Finally we observe that the underlying assumption that the increases in unequal growth are not anticipated is quantitatively important for wealth dynamics during the transition. If high earnings households anticipated that they will experience faster earnings growth for many years to come, they would increase their consumption in response (by saving less), and doing so would further reduce the wealth to income ratio during the transition.

[16] Note that unequal growth also implies that households at the bottom of the income distribution face more risk and have higher saving rates. However since these households control a small fraction of total wealth their impact on wealth is dominated by the impact of high earnings households

6 Conclusions

We have shown that a statistical process for household earnings that involves "unequal growth" that is, high earnings households growing (over time) faster and low earnings growing (over time) slower, can account well for the evolution of the earnings distribution in the United States, as captured by the PSID, over the past 50 years. We have also shown that "unequal growth" has had a positive (around 0.3% per year over the period 1978-2007) effect on aggregate growth and, when markets are incomplete, a potentially large (as high as 15% of lifetime consumption) negative welfare effect. The natural next question is, what is the driver of this increase in unequal growth? There is interesting ongoing work that digs deeper into the sources of unequal growth (see Fogli et al. (2021) and Moll et al. (2022) for two recent examples of such work). We believe that integrating our framework with those papers will yield a better understanding of the aggregate consequences of the changes in individual earnings. We also find that, with the increase in unequal growth, the social value of better (private or public) insurance mechanisms increase significantly, and thus another relevant research direction is how to improve such mechanisms.

References

Abbott, B. and G. Gallipoli (2022). Permanent-income inequality. *Quantitative Economics 13*(3), 1023–1060.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics 109*(3), 659–684.

Atkinson, A. B., T. Piketty, and E. Saez (2011). Top incomes in the long run of history. *Journal of Economic Literature 49*(1), 3–71.

Benabou, R. (1996). Inequality and growth. In B. Bernanke and J. Rotemberg (Eds.), *NBER macroeconomics annual 1996,*, pp. 11–92. The MIT Press.

Benabou, R. and J. Tirole (2016). Bonus culture: Competitive pay, screening, and multitasking. *Journal of Political Economy 124*(2), 305–370.

Benhabib, J. and A. Bisin (2018). Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature 56*(4), 1261–1291.

Bowlus, A. J. and J.-M. Robin (2004). Twenty years of rising inequality in U.S. lifetime labour income values. *Review of Economic Studies 71*(3), 709–742.

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.

Elsby, M. W. L., B. Hobijn, and A. Sahin (2013). The decline of the U.S. labor share. *Brookings Papers on Economic Activity 44*(2), 1–63.

Fogli, A., V. Guerrieri, and M. Ponder (2021). The end of the American dream? Inequality and segregation in US cities. Working paper, Federal Reserve Bank of Minneapolis.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

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

Garicano, L. and E. Rossi-Hansberg (2006). Organization and inequality in a knowledge economy. *Quarterly Journal of Economics 121*(4), 1383–1435.

Gordon, R. J. (2012). Is U.S. economic growth over? Faltering innovation confronts the six headwinds. Working Paper 18315, National Bureau of Economic Research.

Grigsby, J. (2021). Skill heterogeneity and aggregate labor market dynamics. Working paper, Princeton University.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Haskel, J., R. Z. Lawrence, E. E. Leamer, and M. J. Slaughter (2012). Globalization and U.S. wages: Modifying classic theory to explain recent facts. *Journal of Economic Perspectives 26*(2), 119–40.

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, G. L. Violante, and L. Zhang (2022). More unequal we stand? inequality in the United States from the great recession to the COVID pandemic. Working paper, Federal Reserve Bank of Minneapolis.

Heathcote, J., K. Storesletten, and G. L. Violante (2010). The macroeconomic implications of rising wage inequality in the United States. *Journal of Political Economy 118*(4), 681–722.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control 17*(5-6), 953– 969.

Jones, C. I. and J. Kim (2018). A Schumpeterian model of top income inequality. *Journal of Political Economy 126*(5), 1785–1826.

Jovanovic, B. (2014). Misallocation and growth. *American Economic Review 104*(4), 1149–71.

Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *Quarterly Journal of Economics 107*(1), 35–78.

Kopczuk, W., E. Saez, and J. Song (2010, 02). Earnings inequality and mobility in the United States: Evidence from Social Security data since 1937. *Quarterly Journal of Economics 125*(1), 91–128.

Krueger, D. and F. Perri (2004). On the welfare consequences of the increase in inequality in the United States. In E. H. M. Eichembaum and J. Parker (Eds.), *NBER Macroeconomics Annual 2003*, pp. 83–138. The MIT Press.

Lochner, L. and Y. Shin (2014). Understanding earnings dynamics: Identifying and estimating the changing roles of unobserved ability, permanent and transitory shocks. Working Paper 20068, National Bureau of Economic Research.

Luttmer, E. G. J. (2011). On the mechanics of firm growth. *Review of Economic Studies 78*(3), 1042–1068.

Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica 72*(1), 1–32.

Moll, B., L. Rachel, and P. Restrepo (2022). Uneven growth: Automation's impact on income and wealth inequality. *Econometrica 90*(6), 2645–2683.

Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica 64*(6), 1263–1297.

Piketty, T. and E. Saez (2003). Income inequality in the United States, 1913–1998. *Quarterly Journal of Economics 118*, 1–41.

Piketty, T., E. Saez, and G. Zucman (2018). Distributional national accounts: Methods and estimates for the United States. *Quarterly Journal of Economics 133*(2), 553–609.

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. Working paper, Harvard University.

Summers, L. H. (2015). Demand side secular stagnation. *American Economic Review 105*(5), 60–65.

A The PSID sample and variables

This appendix describes in detail the sample selection and the construction of the variables used for the empirical analysis in Section 4.

Sample selection

Our initial sample includes all households in the Panel Study of Income Dynamics from the first year of the sample (1967) the last year available (2018). From the initial sample, we first exclude the following household/year observations:

1. Household/year observations with zero weight in PSID (this includes the entire Latino sample)

2. Household/year observations that report total money income below 20% of the poverty thresholds for a single person household for that year (poverty thresholds are from the U.S. Census)

3. Household/year observations with substantial imputations (as reported by the relevant PSID flag) in labor income.

We perform the first exclusion because households with zero weight cannot be used to compute weighted statistics, and they also typically have very incomplete income information. The second set of observations is excluded done because we deem implausibly low levels of income as driven by measurement error. The final exclusion is done because over time, imputation procedures in the PSID change, and having households whose income is not imputed in one year and imputed in a subsequent year would create bias in the measure of growth rate of household income, which is our main object of interest. The remaining households constitute the sample we use to construct the statistics reported in Figure 1. From this sample, we further exclude households whose heads are younger than 25 and older than 60. These households constitute the sample we use to construct the inequality statistics reported in Figure 2. This exclusion is motivated by the fact that earnings dynamics for younger and older households are driven by entry and exit in the labor market, which are not captured in our simple model of earning dynamics. The final sample selection is dictated by the necessity of computing decile-specific growth rates over the 11 year window; this implies that in each period, we select households that are also in the sample in period $t-4$, $t-2$, $t+2$, $t+4$ and $t+6$. This final sample is the one used to compute statistics reported in Figures 3 and 4. In table A.1 we report sample sizes and the impact of different sample selection choices.

Table A.1: PSID samples

	Dropped Obs	Sample size	Used for:
Full PSID Sample	0	302,097	
HHs with positive weight	14,454	287,643	
HHs with plausible total money income	4,665	282,978	
HHs without substantial imputation	10,022	272,956	Figure 1
HHs with head age 25-60	76,242	196,714	Figure 2
HHs which are in sample from $t-4$ to $t+6$	77,110	119,604	Figures 3 and 4

Variable construction

The key variable we use in our empirical analysis in Section 4 is household earnings. We construct total household earnings as the sum of four PSID variables:

1. Head labor income, which includes wages and salaries, any separate reports of bonuses, overtime, tips, commissions, professional practice or trade, additional job income, miscellaneous labor income plus 50% of business income

2. Spouse labor income, which is computed in the same fashion as labor income of the head

3. Farm household income

4. Other family members' taxable income

B Additional Figures

Figure B.1: Decomposition of aggregate growth: Terms in equation 1

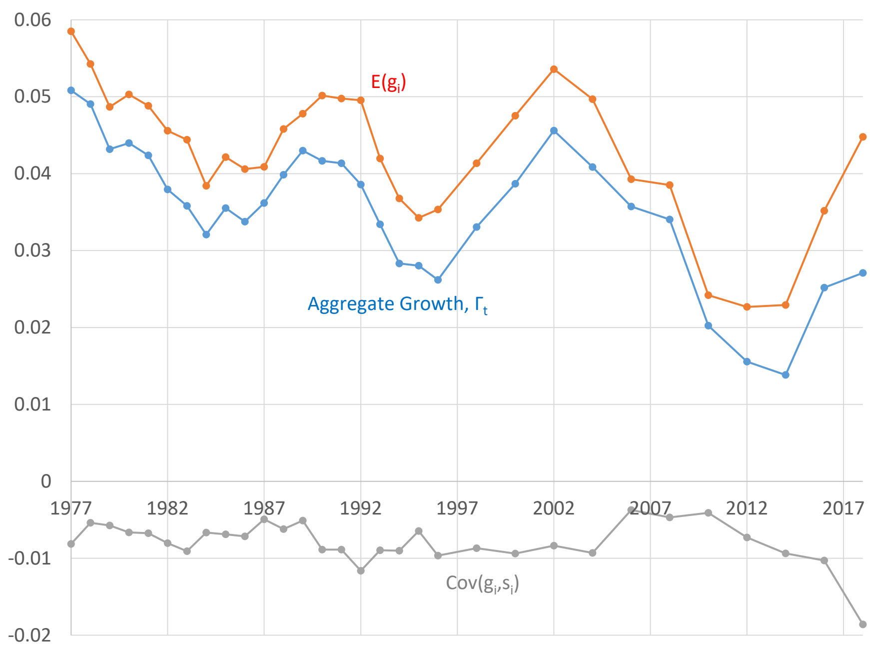

Figure B.2: Decomposition of aggregate growth: Terms in equation 2

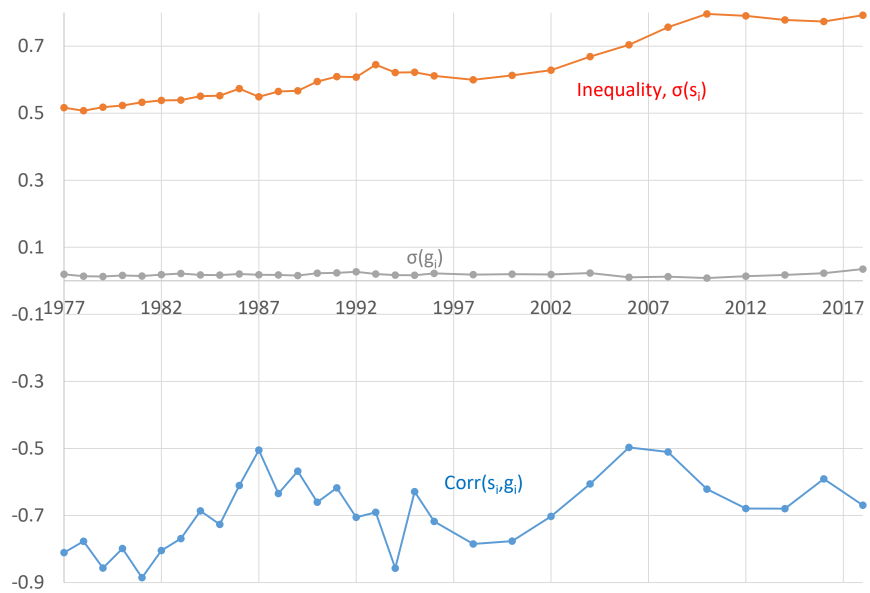

Figure B.3: Earnings growth in PSID: Two samples