HCEO WORKING PAPER SERIES

Human Capital Inequality: Empirical Evidence

Brant Abbott
Giovanni Gallipoli

Working Paper

2018-085

11/2018

HUMAN CAPITAL AND
ECONOMIC OPPORTUNITY
GLOBAL WORKING GROUP

The University of Chicago
1126 E. 59th Street Box 107
Chicago IL 60637

www.hceconomics.org

Human Capital Inequality: Empirical Evidence*

Brant Abbott

Queen's University

Giovanni Gallipoli

Vancouver School of Economics, UBC

September 14, 2018

*We acknowledge support through an Insight grant from the SSHRC of Canada. We are grateful to Sarah O'Brien, Robin Li, Pietro Montanarella and Lily Suh for excellent research assistance.

1 Introduction

Persistent differences in economic outcomes, such as income, wealth, employment and consumption, have received continuing attention in the academic debate. Part of this attention is motivated by the well documented increase in cross-sectional economic inequality that began in the 1980s. Alternative measures of inequality, such as variances, inter-percentile ranges and concentration at the top, all suggest that this increase has occurred across a range of measurable outcomes (earnings, wealth, health).

Wealth inequality, in particular, has received considerable attention, with mounting evidence of steady and economically meaningful changes in the concentration of wealth ownership. By definition, wealth inequality captures disparity in the ownership of productive capital and other non-labor factors of production. In contrast, in this article we focus on the distribution of human capital and its implications for the accrual of economic resources to individuals and households. Human capital inequality can be thought of as measuring disparity in the ownership of labor factors of production, which are usually compensated in the form of wage income. While not directly observable, the notion of a stock of human capital and its main properties have been formally spelled out in the original work of Gary Becker and coauthors (e.g. Becker, 1962, 1964).

Earnings inequality is tightly related to human capital inequality. However, as we will make clear in what follows, earnings dispersion can only provide a partial and incomplete perspective on the underlying distribution of productive skills and on the income generated by way of them. This follows from the fact that, at any point in time, earnings dispersion captures an isolated snapshot of the prevailing disparity in payments to labour. In fact, we would rather learn about the distribution of the underlying stocks of human capital that generate income, and about their market value.

This tension between measures of current earnings and measures of long-term ability to earn has become a distinguishing feature of the literature on inequality and human capital. Despite its shortcomings, a fairly common way to gauge the distributional implications of human capital inequality is to examine the distribution of labor income. While it is not always obvious what accounts for returns to human capital, an established approach in the empirical literature is to decompose measured earnings into permanent and transitory components as in, for example, Gottschalk, Moffitt, Katz, and Dickens (1994), Heathcote, Perri, and Violante (2010) and recently Lochner and Shin (2014). The decomposition approach finds its theoretical motivation in the observation by Friedman (1957) that the underlying ability to earn associated with human capital must be related to the expected flow of income over an individual's life.

Because permanent components of current earnings should, by definition, be detectable in every realization of an individual's earnings process, one can regard an increase in the estimated variance of permanent earnings shocks as a proxy for the underlying changes in the value of human capital.

A second approach to the measurement of human capital inequality focuses on the lifetime present value of earnings. An early example of this is Lillard (1977), while more recent contributions are by Bowlus and Robin (2004) and Guvenen, Kaplan, Song, and Weidner (2017). Lifetime earnings are, by definition, an ex-post measure that is only observable at the end of an individual's working lifetime. This measure captures how an individual's position in the labour market changes over their lifetime through the evolution of their wage and employment status, rather than their position at a single point in time. One limitation of this approach is that it assigns a value based on one of the many possible realizations of human capital returns. That is, since it is impossible to re-run the working life of the same individual multiple times, this measure can only capture the value associated with the particular earnings realization that we observe for each individual. Arguably, this provides a partial view of human capital because it ignores the option value associated with alternative, but unobserved, potential earning paths that may be valuable ex-ante. Hence, ex-post lifetime earnings reflect both the genuine value of human capital and the impact of the particular realization of unpredictable shocks (luck): crucially these two components are not separately observable.

A different but related measure focuses on the ex-ante value of expected lifetime earnings, which differs from ex-post (realized) lifetime earnings insofar as they account for the value of yet-to-be realized pay-offs along different potential earning paths. Ex-ante expectations reflect how much an individual reasonably anticipates earning over the rest of their life based on their current stock of human capital, averaging over possible realizations of luck and other income shifters that may arise. The discounted value of different potential paths of future earnings can be computed using a risk-less discount factor as, for example, Jorgenson and Fraumeni (1989) and Cunha, Heckman, and Navarro (2005) do. This requires a way to estimate the distribution of future earnings outcomes. However, when valuing such outcomes it may be more appropriate to discount their value in a state-dependant fashion, as suggested in Huggett and Kaplan (2016) and Abbott and Gallipoli (2018). These more theory-focused approaches take estimates of the expectation of the present value of lifetime earnings as their yardstick measure of human capital, providing an indirect way to account for some of the key empirical issues highlighted above.

The empirical literature on economic inequality is so expansive that one cannot hope to summarize it all within a few pages. Instead, we focus on a narrow aspect of inequality –

namely human capital disparities – and attempt to provide an in-depth review of this particular dimension of inequality within a short article.[1] Our approach is to dissect and categorize existing evidence into different metrics of human capital, and then summarize and convey the critical information from each of those threads of the literature. While we cannot cover every article that might fit under each alternative measurement approach, we do try to provide an overview of relevant and recent advances within those threads. Combining the evidence we report here with evidence on other dimensions of inequality, such as wealth and consumption, should allow the reader to obtain a more complete understanding of overall economic inequality.

Finally, in addition to reporting evidence on human capital inequality in the United States, we also briefly outline some international comparisons and provide a cursory discussion of ideas from the ever expanding literature on the determinants of human capital heterogeneity.

2 Concepts of Human Capital and their Measurement

While human capital is not directly observable, one can measure the distribution of returns to human capital utilization in the form of earnings. This allows researchers to draw inferences about the underlying distribution and value of human wealth. Moreover, given the multifaceted nature of human capital, we are also able to gauge variation in some of its most relevant features, such as educational achievement, health status and longevity. In what follows we discuss relevant measurement issues and overview key results on the distribution of human capital.

2.1 Earnings Inequality

Earnings are observed at the individual level, but are sometimes analyzed at the household level. Earnings capture the part of income that is payment for labour. Thus earnings can be interpreted as the annual dividend paid by one's human capital.[2] As we argue below, this dividend is not the best measure of human capital, but it is by far the most popular.

[1] For example, a closely related literature on the relationship between consumption inequality and income inequality exists. Examples include Blundell, Pistaferri, and Preston (2008), Aguiar and Bils (2015) and Blundell, Pistaferri, and Saporta-Eksten (2016) among many others, as also summarized in the interesting book of Jappelli and Pistaferri (2017).

[2] An obvious measurement issue is that the remuneration of human capital can be recorded as a payment for business activities and possibly be treated like capital income. Active business activity often takes the form of sole proprietorship of corporations organized as 'limited liability' legal entities that pay corporate income tax on annual taxable income. In such circumstances the owners can be active earners rather than passive rentiers, as pointed out by Smith, Yagan, Zidar, and Zwick (2018) and Dyrda and Pugsley (2018). This complicates the measurement and taxation of human capital returns.

Some of the highest quality earnings data are sourced from government tax or social security data; however, such data cannot be freely accessed for analysis and are usually maintained in access-controlled research data centres. Earnings are also observable in survey data, such as the Census, though the reliability of such sources may be less than ideal.[3]

One way to measure earnings inequality is to compute ratios of earnings of individuals at different percentiles of the income distribution. Researchers usually report the 90-10 ratio (90th percentile relative to 10th percentile), as well as the 90-50 and 50-10 ratio to further decompose what half of the distribution is driving changes in inequality. Such statistics are easily computed as one simply determines earnings at the relevant percentiles and takes their ratios.

Another common measure of earnings inequality is the top share of earnings. This is the share of total earnings in the economy accruing to those with the highest earnings, such as the highest 1% or 10% of earnings. Such statistics are also easy to compute as one determines the relevant percentile threshold (e.g. the 99th percentile) and then computes the total earnings of those above this threshold. Then, dividing this number by total earnings in the entire economy one arrives at the top-share statistic.

One of the most frequently cited measures of earnings inequality is the Gini coefficient. The interpretation of the Gini coefficient can be related to the Lorenz curve, which describes the proportion of economy-wide earnings attributed to those at, or below, a given percentile of the earnings distribution. When earnings are equally distributed (every person has identical earnings) the Gini coefficient is zero. In contrast, when earnings are perfectly unequal (one person gets everything, all others get nothing) the Gini coefficient is one. Earnings are a non-negative quantity, which ensures that the Gini coefficient is between zero and one. However other variables, such as net worth, can assume negative values and their distributions can therefore exhibit a Gini coefficient exceeding one.

A tremendous amount of data and measurements exist on income and earnings inequality in the United States. When analyzing inequality in human capital in the US one should be careful to source out facts based on Social Security Administration (SSA) data, for which reported amounts are labor income and the unit of analysis is the individual. Another popular source for computing inequality facts is the IRS (U.S. Internal Revenue Service) database. However IRS

[3]Non-trivial issues of measurement become apparent when looking at labor income payments, as pointed by Guveven and Kaplan (2017). These authors study top income inequality, measured as the share of incomes accruing the individuals in the top percentiles of the income distribution: crucially they use a combination of total income data from the Internal Revenue Service (IRS) and labour income information from the Social Security Administration (SSA), starting from 1981. These two data sets allow them to isolate differences in top income inequality based on the definition of income (labour income vs total income) and to differentiate with respect to the unit of analysis (tax unit vs individual).

records report total income and the unit of analysis is the tax-unit.

An excellent description of facts and figures, based on SSA data, is in Kopczuk, Saez, and Song (2010). These authors consider 80-50 and 50-20 ratios rather than the usual 90-50 and 50-10 ratios. They find that while the 80-50 ratio dipped after the Second World War (WWII) to as low as 0.41, a continuous steady rise in this ratio followed: the ratio climbed above 0.6 by the mid 2000s. This pattern reflects a steady increase in above-median inequality in the post-war era. The 50-20 ratio fell overall in this period, starting at about 0.9 at the end of WWII, and hovering around 0.75 by the mid-2000s, indicating a decrease in below-median inequality; however, this decline was not steady, with periods of increase occurring in the early 1980s and early 2000s. The overall inequality measure (the 80-20 ratio) showed a drop after WWII, and substantial increases thereafter, as the above-median increases in inequality outpaced the lower inequality in below-median earnings. These authors also report the time series of the Gini coefficient for earnings, which dropped to as low as about 0.34 in the early 1950s and then steadily rose to just under 0.45 by the mid-2000s. Figure 1 reproduces a figure from the original study of Kopczuk, Saez, and Song (2010), which displays Gini coefficients for earnings by gender and over time. A similar dynamic is observed in top earnings shares: the top 1% share of earnings dropped from 9.55% in 1939 to 5.92% in 1960, but then more than doubled to 12.28% by the mid 2000s. These general patterns and magnitudes are consistent with many other studies and data sources, for example Heathcote, Perri, Violante, et al. (2010) and Eckstein, Nagypal, et al. (2004), who use Current Population Survey data for similar time periods, or Rios-Rull and Kuhn (2013), who use data from the Survey of Consumer Finances between 1989 and 2013.[4].

Lastly, recent evidence has suggested that changing patterns of inequality are not a general phenomenon, but rather are linked to growth of inequality within and across groups. For example, Lemieux (2006c) shows that much of the previously discussed increases in inequality are attributable to the increasing return to postsecondary education.[5] In a related line of research Song, Price, Guvenen, Bloom, and Von Wachter (2015) find that most of the observed increase in earnings inequality can be attributed to growing differences between firms in the wages they pay to workers. Similarly, Barth, Bryson, Davis, and Freeman (2016) find that rising inequality can be partly attributed to increasing inequality across industries.[6]

[4]See also Quadrini and Ríos-Rull (2015)

[5]Furthermore, Lemieux (2006a) shows that inequality that is not explained away by observables, such as education, increased very little.

[6]Song, Price, Guvenen, Bloom, and Von Wachter (2015) suggest that the firm, rather than the industry, is the key dimension.

2.2 Lifetime Earnings Inequality

In theory, lifetime earnings correspond to the present discounted value of all labour-related earnings over the course of one's entire life. This concept is closer to our definition of human capital than annual earnings is, as it captures the entire lifetime of dividends paid by one's human capital. Importantly, these are the actual realized earnings of the individual, rather than forecasted earnings, and so, to obtain a measure of this quantity, the individual's entire lifetime of earnings must be observed. In practice, lifetime earnings can be approximated by the present discounted sum of annual earnings over some arbitrary period of time. As examples, Lillard (1977) studies a narrow data set on post-schooling earnings, Guvenen, Kaplan, Song, and Weidner (2017) study extensive administrative earning records for workers between ages 25 to 55. When the object of interest is differences in lifetime income across cohorts, like in Guvenen, Kaplan, Song, and Weidner (2017), researchers can avoid specifying a discount rate by assuming that all cohorts discount at the same rate. In some instances, such as Bowlus and Robin (2004) and Low, Meghir, and Pistaferri (2010), theoretical models of life-cycle earnings are used to interpret data panels and generate estimates of realized lifetime earnings distributions. Once estimates of lifetime earnings are obtained for a cross-section of individuals, researchers can compute the same measures of inequality described above in the context of current (annual) earnings.

Guvenen, Kaplan, Song, and Weidner (2017) use 57 years of panel data (1957-2004) from the Social Security Administration to study the lifetime earnings of 27 birth cohorts. Each cohort is defined by the year that they turn 25, and the sample period spans 31 years. Annualized lifetime earnings are defined as the sum of real annual labour income between ages 25 and 55, divided by 31. Lifetime earnings of men and women are examined separately to reach two broad conclusions: first, men have experienced a stagnation in lifetime income while women have seen a rise in lifetime income; second, although lifetime income inequality within each gender group has risen, the closing of the gender gap has kept overall lifetime inequality nearly flat. Indeed, the 75-25 inter-percentile ratio in lifetime earnings was fairly steady (at about 3) for cohorts turning 25 between 1957 and 1983. However, looking separately at men and women the respective 75-25 ratios increased from about 2.3 to 2.7. This striking pattern is displayed in panel (b) of Figure 2, which reproduces Figure 11 in Guvenen, Kaplan, Song, and Weidner (2017). Although overall inequality in lifetime earnings changed very little, at the top of the distribution we did observe an increase in inequality: the 90-50 percentile ratio moved up from 2.3 to 2.7, while inequality at the bottom of the distribution fell, as indicated by the decreasing 50-10 percentile ratio, which dropped from 3.1 to 2.9.

Bowlus and Robin (2004) also study the evolution of lifetime earnings inequality over time. They note that lifetime income is dependent not only on an individual's current position in the labour market, but also on the evolution of their wage and employment status. Therefore they develop a model incorporating wage mobility and employment transitions. For each individual in the sample they simulate a possible lifetime earnings trajectory, assuming that individuals are subject to the same distribution of shocks as older workers today, and thereby estimate the distribution of lifetime earnings realizations. Their sample consists of white males aged 16-65 from the March CPS between 1978 and 1999. They use this sample to estimate the experience and education specific wage and employment transition probabilities. The key finding from this exercise is that lifetime earnings inequality is roughly 40% lower than annual earnings inequality. This implies that measures of cross-sectional inequality are significantly reduced when one factors in future employment and wages. This is partly due to the fact that the young, while exhibiting lower current wages, benefit disproportionately more from upward wage mobility relative to the old. It is also shown that inequality has been rising over time within education and experience groups, especially for university graduates. Both within and across group variation is important in explaining rising earnings inequality.

As mentioned above, using a narrow data sample from the NBER-TH survey, Lillard (1977) also examines lifetime dispersion of earnings and draws comparisons to cross-sectional earnings inequality. The sample for this classic study includes only males who volunteered for the U.S. Air Force pilot, navigator and bombardier programs in the last half of 1943. Lillard explores the relative importance of schooling, measured ability and family background for both annual and lifetime earnings. His broad conclusions are that cross-sectional earnings inequality is about 50-80% higher than lifetime earnings inequality. Annual earnings inequality is high, both overall and by age group. These inequality measures are not sensitive to discounting, nor to the length of working life. Another finding of this pioneering work is that the contributions of schooling, ability, and background to variation in lifetime earnings are similar to their respective impacts on the variation of annual earnings within age groups; age itself is the most important factor in explaining annual earnings. These results, as well as those from the previously discussed articles, are echoed in a later study of Norwegian data by Aaberge, Mogstad, et al. (2014). This study also shows that lifetime earnings are much more equally distributed than annual earnings. Life-cycle bias, which is assessed by comparing within-age earnings inequality to aggregate earnings inequality, is a key contributor to the discrepancy between annual and lifetime earnings.

2.3 Expected Lifetime Income Inequality

At least since the seminal work of Modigliani and Brumberg (1954) and Friedman (1957), it has been understood that expected lifetime earnings (as opposed to realized, and observed, lifetime earnings) are the main drivers of economic decisions and utility. This discrepancy between ex-ante (unobserved) magnitudes and ex-post realized outcomes is central to make sense of the choices and behaviours of economic agents. As Cunha, Heckman, and Navarro (2005) emphasize, differences in expected lifetime earnings reflect ex-ante heterogeneity in human capital, whereas un-forecasted differences in lifetime earnings reflect luck and uncertainty. Assuming concave utility functions and a lack of full insurance, greater inequality due to uncertainty reduces the welfare of all agents. However, greater inequality in expected (hence predictable) resources, holding the mean constant, would imply that the rich become better-off and poor become worse-off in welfare terms. Therefore, understanding how lifetime income inequality breaks down into its different components is crucial for understanding welfare.

To assess inequalities in human capital we primarily need to estimate expectations of lifetime earnings. That is, we need to find ways to gather and aggregate all information available to individuals when making forward-looking judgements about their future stream of resources. These judgements about future earnings do not include variation due to the residual component of lifetime earnings, such as luck and other unforecastable factors, as these are, by definition, unrelated to a worker's skills or potential to earn income. Two individuals can be identical in their education and earning abilities, yet they may realize very different earnings in the future due to luck; however, despite these differences due to chance, we should consider these individuals to have possessed the same ex-ante human capital when younger, and indeed to have been equal in their opportunities. Two lines of research that use different approaches to quantify inequality in the distribution of ex-ante human capital have been advanced by Huggett and Kaplan (2016) and Abbott and Gallipoli (2018).[7] The central questions in these lines of inquiry relate to how agents discount future payoffs when forming expectations, given different possible realizations of future earnings; and how much prior information (possibly unobservable to researchers) agents possess about their future earnings. With respect to the former question, recent advances in asset pricing theory suggest that agents employ state-dependent stochastic discount factors, which can be estimated non-parametrically using consumption panel data. With respect to the latter question, work by Cunha, Heckman, and Navarro (2005) suggests the use of early life choices, such as the decision to complete education, as proxies of ex-ante information available to agents. Credible answers to the question whether (and how much) re-

[7] Blundell, Pistaferri, and Saporta-Eksten (2016) also estimate human wealth, i.e. expected lifetime earnings, but crucially they do no consider state-dependent stochastic discounting.

alized lifetime earnings inequality overstates ex-ante human capital inequality will be a critical, and likely contentious, issue in the ongoing debate on human capital inequality.

2.4 Permanent vs. Transitory Earnings Decompositions

A popular approach to studying human capital inequality involves decomposing earnings into underlying permanent and transitory components. While the transitory component reflects temporary, unpredictable variation or measurement error in wages, the permanent component can be thought of as a measure of human capital, akin to the flow value of the expected future earnings discussed in the previous section. This decomposition turns out to be extremely useful for two reasons: first, one can study trends in human capital inequality by estimating how the variance of the permanent component of earnings has changed over time; second, the actual implementation of this method can be carried out using fairly simple variance-covariance matrix estimators. This approach does, however, require a stylized set of assumptions and cannot deliver point estimates of each individual's permanent component value, nor does it suggest how to discount across possible future realizations of earnings to generate human wealth estimates.[8] That said, much has been learned about the nature of trends in human capital inequality through such decompositions.

To fix ideas, denote an individual's log earnings by y_{it}, and assume they are the sum of a transitory component μ_{it} and permanent component ν_{it}: $y_{it} = \mu_{it} + \nu_{it}$. The transitory component is such that realizations μ_{it} are independent of past and future realizations of the same component, or more precisely $E(\mu_{it}\mu_{it+\tau}) = 0$ for $\tau \neq 0$. In contrast, the distribution of realizations of the permanent component explicitly depends on past realizations. It is common for the permanent component to be modelled as either an $AR(1)$ or unit-root process; however, econometric advances also allow for non-linear dynamics, which may be a more appropriate assumption.[9] Typically, these models can be identified with relatively short data panels, e.g. three year for the AR(1) specification, but longer panels are needed for more reliable estimates. If longer longitudinal data sets are available, windows of data can be used to produce time-varying estimates of the variances of the key parameters of these models.[10]

Textbook examples of this approach are Gottschalk, Moffitt, Katz, and Dickens (1994) and

[8] To learn about inequality across groups, for example, one must estimate a separate wage process for each partition. One such example is the work of Abbott, Gallipoli, Meghir, and Violante (forthcoming) who estimate and simulate separate wage processes for six groups: men and women, each across three education levels.

[9] See Arellano, Blundell, and Bonhomme (2017) or Guvenen, Karahan, Ozkan, and Song (2015) for recent research on nonlinear implementations.

[10] What we can learn about the nature of earnings dynamics from longer auto-covariances is discussed in Meghir and Pistaferri (2006).

Moffitt and Gottschalk (2012). The authors decompose the earnings residuals of prime age males in the Panel Study of Income Dynamics (PSID) from 1970 to 2004 into permanent and transitory components.[11] As a baseline, they find that inequality measured by the variance of log-earnings residuals nearly doubled during this period, and between 51 and 69 percent of the increase is attributable to the permanent component of earnings. These findings are displayed in Figure 3 of Moffitt and Gottschalk (2012), which is reproduced in Figure 3 here. Because the permanent component better reflects human capital, as opposed to luck and other random variation, we conclude that a large increase in human capital inequality occurred over the sample period, although only about half to two-third of the rise in earnings inequality is genuinely attributable to changes in the distribution of human capital. Their estimates for 2004 indicate that about 60% of earnings inequality is attributable to a permanent component.

Lochner and Shin (2014) also study the evolution of inequality in male earnings using PSID data from 1970 to 2008. They specifically focus on the returns to human capital and show that the pricing of unobserved skills has changed dramatically over time - returns to unobserved skills have increased in the 1970s and 80s before reversing back by the late 90s. Moreover, they observe that the variance in log earnings residuals remained stable until the 90s before rising. These discrepancies in trends point to changes in the variance of permanent shocks and transitory shocks. The authors also look at changes in returns at various points in the distribution of earnings (poorer vs richer individuals), and find that the increase in returns to unobserved skills is not detectable at the top of the earnings distribution.

Kopczuk, Saez, and Song (2010) take earnings averaged over a five-year period to be a measure of the permanent component of earnings in the SSA data. Following this alternative approach they find that the permanent component is a much larger fraction of the variance of log-earnings, over 80% by the early 2000s. This approach also indicates that the entire rise in earnings inequality is due to a rise in the variance of the permanent component. It is not clear whether the differences in findings between this research and Moffitt and Gottschalk (2012) is due to the definition of permanent components, working with raw data rather than residuals, or the use of a different data sets.

While the previously mentioned articles have studied earnings (which depend both on wages and on labor supply), a closely related literature focused on the evolution of wages alone. The permanent component of one's wage rate is also a measure of human capital, under a simple proportionality assumption. Using a permanent-transitory wage decomposition, estimated from PSID data, Heathcote, Storesletten, and Violante (2010) study the evolution of

[11] Residuals are defined as the component of log-earnings not predicted by education, race, an age polynomial, and interactions among these variables.

the parameters of an AR(1) process describing the evolution of permanent wage factors. These authors find that both the permanent and transitory components of wages increased similarly from the mid 1960s to early 2000s, and that the permanent component is significantly larger than the transitory component. At the end of their sample period the transitory component accounts for only about 30% of the total variance of log-wage residuals.

Heathcote, Perri, and Violante (2010) also work on residual wage inequality in the United States using the PSID data. Echoing the previously discussed literature, this research finds an increase in residual wage inequality since the mid 1960s, and finds that the rise is roughly 50% attributable to an increase in the permanent component of wages and 50% attributable to the transitory component. One key aspect discussed in this paper is the impact of alternative ways to identify the components of the wage process. Identification can be achieved using either a set of moments on the level of log-wages or a set of moments on first-differences of log-wages. Although the main findings just mentioned are not affected, the authors do note other important differences in the results obtained under the two identification strategies. This disagreement, they say, "indicates that the permanent-transitory model is misspecified." This relates closely to our earlier comment suggesting that the literature has now directed itself towards studying non-linear dynamics, discussed next.

Guvenen, Karahan, Ozkan, and Song (2015) study earnings dynamics using SSA data, and find that earnings shocks exhibit "strong negative skewness and extremely high kurtosis," hence contradicting the log-normality assumption underlying the previously discussed literature.[12] While these authors do not decompose earnings into permanent and transitory components in the same manner described above, they do study these components, and their relationships with inequality, using nonparametric methods.[13] Their findings are very interesting. For high-earning individuals further increases in earnings tend to be transitory and decreases in earnings tend to be permanent. For low-earning individuals the opposite is true in that increases in earnings tend to be persistent and decreases tend to be transitory. It is not yet clear how this affects our understanding of the evolution of residual wage inequality; however, it does provide a further indication of why current (contemporaneous) earnings inequality is a poor measure of inequality in the stock of human capital.

[12] As we stated above, we believe much has been learned from that previous literature, and many of those lessons still stand.

[13] Theoretical work by Huggett, Ventura, and Yaron (2006) shows that some of these patterns can be replicated in a model of human capital dynamics.

3 International Perspective

Most of the studies we have mentioned refer to the United States. This is partly a reflection of the enormous amount of work that has been done using US data. Nonetheless, it is interesting to assess how the patterns of cross-sectional earnings inequality compare across different countries, and to pinpoint common trends and differences.

A good source of information about cross-sectional earnings dispersion in different countries is the special issue on inequality published by the Review of Economic Dynamics in 2010. This issue contains a selection of papers examining inequality patterns in nine countries (Canada, Germany, Italy, Mexico, Russia, Spain, Sweden, US and UK). A study of this detail and scale is rare, thus the findings are of some significance, as they allow one to compare aspects of inequality across the subset of countries considered. Most of the studies in this special issue are based on data for the period between the 1970s and the 2000s: this makes it possible to draw comparisons and assess similarities and discrepancies. Another complementary, and equally rich, source of information can be found in the work of Atkinson and Piketty (2007).[14]

We begin by highlighting the most recurrent patterns. Most studies present similarities in the evolution of earnings and income inequality over the period considered: (1) in most countries we observe rising earnings inequality starting in the late 1980s, with the notable exceptions of Spain and Russia; (2) income inequality is significantly higher than consumption inequality, consistent with the notion that cross-sectional insurance and redistribution play a key role in all countries considered in the study; (3) earnings and income inequality grew remarkably during the 1990s in most countries, suggesting that common technological or institutional factors[15] may have been key for inequality growth in that period; (4) residual inequality (wage and earnings dispersion after controlling for observable characteristics) played a key role in the rise of overall inequality, with transitory shocks accounting for a large share of this surge;[16] (5) a commonality found across all countries considered is that the variance of permanent wage shocks is smaller than the variance of transitory wage shocks; (6) during recessions, inequality in earnings increased quite sharply everywhere, especially at the bottom of the distribution of earnings. This suggests a link between unemployment (or limited employment) and earnings of the poorest workers, who seem to carry most of the burden of regressions in terms of labor income.

Important differences also become apparent when comparing the inequality experience of different countries. For example, while labour supply plays a central role in shaping household

[14] See also Roine and Waldenström (2015) and Morelli, Smeeding, and Thompson (2015).

[15] See DiNardo, Fortin, and Lemieux (1996).

[16] For a discussion of possible explanations, see Lemieux (2006b).

level inequality dynamics in countries like the US and Canada,[17] this is less true in European countries like Italy, Spain and Germany. This discrepancy highlights the importance of heterogeneity in human capital utilization in different economies: whether labour supply changes at the household level impact the dynamics of aggregate inequality partly depends on the rate of female labour participation and on the distribution of household types (marital status, number of children, education of spouses). Some European countries exhibit a low level of female labour participation and, given the relatively inelastic male labor supply, this may weaken the pass-through from changes in hours worked to overall earnings dispersion. In addition, wage inequality in countries like Italy, Spain and Germany is measured to be much lower than in the US, which can make any heterogeneity in labor supply less salient for earnings dispersion.

The country that most closely mimics the patterns observed in the US appears to be Canada.[18] As documented in the work of Brzozowski, Gervais, Klein, and Suzuki (2010), the time paths for wage dispersion, hours dispersion and wage-hour correlation since 1975 are fairly close to those observed in the US. However, despite the growth in income inequality, disposable income and consumption inequality did not grow as much in Canada, suggesting that redistribution trough taxation and benefits may have been effective in mitigating the effects of growing economic discrepancies. The relative mitigation effects seem, however, to have become weaker over time: during the 1980s there was a strong rise in before-tax income inequality in Canada, which was mostly absorbed by the tax and transfer system. However, after the 1990s, when the before-tax income inequality rose again, the tax and transfer system has been less effective in offsetting the rise in income inequality. This resulted in a more pronounced increase after-tax income inequality.

One interesting aspect that differentiates the US and Canada is the variation in the health dimension of human capital. Arguably, one of the key aspects determining the value of an individual's human capital is the relative health enjoyed by that person. The value of human wealth crucially depends on the ability to generate income through labor supply, which becomes harder in poor health. In the United States a sharp decline in labour market, marriage, and health outcomes has been documented for relatively poorer white non-Hispanics. Case and Deaton (2017) find that, for birth cohorts after 1940, this demographic group, especially those with less than a four-year college degree, experienced a decline in real wages that is more pronounced with each successive cohort. This decline in real wages is accompanied by rising mortality from drugs and alcohol poisoning, suicide, higher risk of heavy drinking, chronic pain, labour force detachment, and declining marriage outcomes. These trends are not common

[17] See Abowd and Card (1989).

[18] For an informative comparison see Green and Sand (2015).

to those with a bachelor's degree; educated men have seen limited changes in health, mental health and marriage outcomes and have flat profiles for labour force participation, suicide and drug mortality. The decline in real wages is also not identical across education groups: controlling for age, real wages for those with degrees are on average 10% higher for the cohort born in 1980 relative to the cohort of 1940, while wages for those without a degree are 10% lower.

In related work, Milligan and Schirle (2018) present evidence that health inequality changes were not as extreme in Canada. Using a comprehensive administrative dataset of Canadian men and women spanning a half century, Milligan and Schirle examine the relationship between income and health, estimating a gap in life expectancy between the lowest and highest earners of about 11 percent (an eight years difference in lifespan for men). Comparing this gap to the one in the US, it is only about 3/4 as large as that estimated recently by Chetty, Stepner, Abraham, Lin, Scuderi, Turner, Bergeron, and Cutler (2016). Crucially, these authors do not find the same reversal in survival rates for mid-life males that has been documented in US studies like Case and Deaton (2017). In contrast to the US experience, it appears that the evolution of the earnings-longevity gradient can be described as a fairly uniform shift in Canada, with equal improvements among both high and low earners. The fact that in the United States there is a growing mortality gap between the top and bottom of the income distribution, while in Canada the mortality gap remains fairly constant suggests the possibility that institutional factors may help mitigate the growth of human capital and earnings inequality. While no clear consensus exists on what exactly can account for such differences, some potential explanations could relate to differences in access to health-care and education, as well as the incidence of long-run stress and hardship associated with job uncertainty.[19]

The importance of income stabilizers is confirmed when examining the experience of Sweden, where earnings inequality increased in the early 1990s. As discussed in Domeij and Floden (2010) this growth is largely attributable to movements in and out of employment. These authors also find that earnings inequality in Sweden was mostly due to increasing residual earnings dispersion, triggered by increasing volatility of persistent shocks. Like in the case of Canada, inequality in disposable income (after taxes and transfers) increased much less than for raw incomes, indicating the presence of an effective welfare system. There is no evidence of an increasing trend in household level consumption inequality and, since persistent shocks are difficult to insure against, it is reasonable to conclude that the Swedish tax-and-transfer was

[19]One aspect worth highlighting is that the concentration of wealth in Canada is less severe than that in the US: the top 5 percent of richest households holds 35 percent of wealth in Canada, which is similar to the share held by the top 1 percent of richest households in the US. A related literature examines the implications of changes in tax and transfers system for inequality, see Kaymak and Poschke (2016) and Hubmer, Krusell, and Smith Jr (2016).

able to absorb and mitigate the effects rising income inequality.[20]

Britain also exhibits patterns that are qualitatively similar to those observed in the US: income inequality in the UK rose dramatically during the 1980s and continued its growth over the 1990s, while consumption inequality did rise but at a slower pace. Interestingly, Blundell and Etheridge (2010) show evidence that the surge of inequality in Britain during the 1980s is attributable to the strong growth of the volatility of permanent shocks perturbing labour income. In turn, this suggests a role for structural transformation that affected both the cross-sectional distribution of, and returns to, human capital: as they point out, changing education differentials had a big role in this episode, especially for the growth of inequality over the 1980s and early 1990s. The subsequent growth can instead be attributed to changes in transitory volatility, like in many other countries in this comparative study. This evidence is consistent with earlier results in Gosling, Machin, and Meghir (2000).[21]

One important caveat is that caution must be exercised when examining the drivers of inequality in different contexts, as local conditions and arrangements can affect both the distribution of human capital and its returns. One example of this is provided by the Mexican experience, as documented in Binelli and Attanasio (2010). Mexico's market structure is quite different from that of its northern neighbours US and Canada. A distinctive feature of Mexico's labor market is the relative size of its informal sector as well as one of the highest levels of measured income inequality in the world, exceeding that in the US. Inequality in Mexico rose most prominently during the 90s, when the 'Peso crisis' occurred and a possible explanation is that the share of workers in the informal sector responded to changes in aggregate conditions: given the unregulated labor market, unskilled workers are more vulnerable during economic downturns and when unemployment hits during an economic crisis, many workers see their wage negatively impacted as they accept informal employment. The cyclical sensitivity of the returns to labor has therefore a large effect on the poorer sections of the working population and affects inequality by stretching out the bottom end of the earnings distribution.

Another example of the central importance of local dynamics is Germany, a country with a unique historical background. In this case a major event - the reunification in 1990 - has shaped the recent evolution of inequality. Earnings dispersion was stable in the years before 1990, but after West and East Germany were rejoined, both wage and earnings inequality increased almost mechanically. Interestingly, the work by Fuchs-Schündeln, Krueger, and Sommer (2010)

[20]Questions remain about the process of joint determination of human capital outcomes and redistributive policies. Abbott and Gallipoli (2017) explore the possibility that the shape of the distribution of human capital may itself affect optimal public policies and redistribution. They provide microdata evidence that this may in fact be the case for a set of developed countries.

[21]See also Blundell, Gosling, Ichimura, and Meghir (2007).

documents that the German public welfare system was rather effective in mitigating the impact of growing inequality, as disposable income and consumption inequality increased only modestly.

Two countries that bucked the overall trend of increasing inequality between the 1980s and early 2000s were Spain and Russia. In the case of Spain, Pijoan-Mas and Sánchez-Marcos (2010) show that inequality in individual net labour earnings and household net disposable income decreased significantly over the study period (1985-2000). A partial explanation can be found in the fact that the unemployment rate was extremely high at the beginning of the sample period, despite a continued economic expansion. Two key changes shaped the evolution of inequality in Spain after 1985: first, the tertiary education premium fell (in contrast to many other countries); second, the unemployment level fell from 24 percent in 1985 to 13 percent in 2000. These factors ultimately had the effect to reduce inequality in labour earnings in Spain, by compressing both the top and bottom ends of the distribution. The study also finds that the decrease in income inequality was driven mostly by changes in the distribution of permanent components of earnings. This is in sharp contrast to the case of Russia, the other country where, in recent periods, recorded income inequality exhibited a downward trajectory. Analysis of Russian micro-data by Gorodnichenko, Sabirianova-Peter, and Stolyarov (2010) suggests that a moderation in the volatility of transitory shocks during a period of strong economic recovery was responsible for this pattern. Interestingly, expenditure and income inequality in Russia are not far apart, which potentially indicates a lack of cross-sectional insurance and an ineffective transfers system that is unable to mitigate fluctuations in disposable income and consumption.

While all these studies are useful to characterize the evolution of cross-sectional earnings inequality in a diverse set of countries, there is still relatively little work on the distribution of lifetime earnings and the present value of human capital in countries outside the US. As mentioned before, the lifetime notion of human wealth is appealing as it relates to extended flows of income associated with underlying productive skills. The lack of international evidence on lifetime human wealth inequality is possibly due to the fact that such studies require rich data and more complex analytical approaches. One notable non-US study on the relationship between current and lifetime income uses high quality administrative data from Norway (Aaberge, Mogstad, et al., 2014). Specifically, this study examines longitudinal data between 1942 and 2006, considering cohorts born in the interval 1942-1944. In this way the authors are able to capture the full working life of the sample members between ages 23/25 and 62/64. The measure of lifetime income is computed using the approach of Haider and Solon (2006), and is consistent with studies in the American context focusing on the annuity value of the discounted sum of real income. Findings from this study are broadly consistent with those

obtained for US data: they highlight the life-cycle bias in an individual's earnings profile, and show that inequality measured using lifetime income is much lower than inequality measured using cross-sectional income.

4 Summary and Conclusions

Heterogeneity in human capital is a key source of differences in economic well-being. This article provides a synopsis of the empirical approaches that have characterized the analysis of human capital inequality over the past few decades. In the process we overview various methodological and measurement issues, and summarize the existing evidence on the changing patterns in the distribution of returns to labor and human wealth across various countries. Our focus is on specific aspects of human capital measurement and on the empirical evidence gathered by the large literature on these topics. It is impossible to summarize this extensive research entirely within such a limited space.

One key aspect on which we remain silent is the analysis of the fundamental causes underlying the persisting, and often growing, disparities in human capital and earnings. Some of the potential causes examined in the literature relate to innate abilities, heterogeneity of early investments in human capital due to family background, as well as short term credit constraints that limit the ability to attend formal schooling, especially at the tertiary education level. Although much academic and policy focus has been on returns to schooling, several empirical studies have highlighted the importance of early environments in fostering cognitive and non-cognitive skills, which are key in determining the development of human capital and, eventually, economic success. This growing body of work has stressed the importance of childhood influences, especially at very young ages, on skill development, arguing that while institutional learning is an important aspect of skill development, it is not the only channel through which skills are developed, and not necessarily the most important (see Heckman, 2000; Aizer and Currie, 2014). The key message of this literature is that a child's environment is a major predictor of future success, as disadvantages arise from lack of cognitive and non-cognitive stimulus early in life or even from beneficial environments at or before birth itself. These disadvantages are then compounded as the child moves through subsequent stages of development (see Heckman, 2004; Cunha and Heckman, 2007). While distinct, all these explanations find their motivation in the underlying lack of equal economic opportunities offered to young individuals.

The broad set of measures and strands of research we discussed in this article all suggest that human capital inequality has grown since the late 1970s. However, since human capi-

tal is not directly observable, new results and empirical evidence can, and certainly will, be brought forward to further refine and advance our understanding of human capital heterogeneity. As we stressed all along, earnings inequality is not the same as human capital inequality, and comparisons between annual and lifetime earnings inequality show how much transitory earnings inequality can overstate underlying human capital disparities. Lifetime earnings are themselves not a satisfactory proxy for human capital, as they encompass a large component of realized shocks that are unrelated to human capital. We believe the direction of future research in this field will lean towards data and methods that shed new light on individuals' ex-ante, heterogeneous valuations of their human wealth. To understand the cross-sectional disparities in permanent income and human wealth it will be crucial to measure how much of realized lifetime income inequality is a predictable reflection of human capital at any point in time, as opposed to sheer luck or random unpredictable events. Answering these question has immediate and profound implications for tax and redistribution policies, as well as for our positive assessment of what drives individuals to vastly different long-term outcomes.

References

AABERGE, R., M. MOGSTAD, ET AL. (2014): "Income mobility as an equalizer of permanent income," *Discussion Papers*, 769.

ABBOTT, B., AND G. GALLIPOLI (2017): "Human capital spill-overs and the geography of intergenerational mobility," *Review of Economic Dynamics*, 25, 208 – 233, Special Issue on Human Capital and Inequality.

——— (2018): "Permanent Income Inequality," *Working Paper*.

ABBOTT, B., G. GALLIPOLI, C. MEGHIR, AND G. L. VIOLANTE (forthcoming): "Education policy and intergenerational transfers in equilibrium," *Journal of Political Economy*.

ABOWD, J. M., AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57(2), 411–445.

AGUIAR, M. A., AND M. BILS (2015): "Has Consumption Inequality Mirrored Income Inequality?," *The American Economic Review*, 105(9), 2725–2756.

AIZER, A., AND J. CURRIE (2014): "The intergenerational Transmission of Inequality: Maternal Disadvantage and Health at Birth," *Science*, 344, 856–861.

ARELLANO, M., R. BLUNDELL, AND S. BONHOMME (2017): "Earnings and consumption dynamics: a nonlinear panel data framework," *Econometrica*, 85(3), 693–734.

ATKINSON, A., AND T. PIKETTY (2007): *Top incomes over the twentieth century: a contrast between continental european and english-speaking countries*. Oxford University Press.

BARTH, E., A. BRYSON, J. C. DAVIS, AND R. FREEMAN (2016): "Its where you work: Increases in the dispersion of earnings across establishments and individuals in the United States," *Journal of Labor Economics*, 34(S2), S67–S97.

BECKER, G. S. (1962): *Investment in Human Capital: A Theoretical Analysis*pp. 9–49. The Journal of Political Economy Vol. LXX, No. 5, Part 2 (University of Chicago Press).

BECKER, G. S. (1964): *Human Capital*. NBER.

BINELLI, C., AND O. ATTANASIO (2010): "Mexico in the 1990s: the main cross-sectional facts," *Review of Economic Dynamics*, 13(1), 238.

BLUNDELL, R., AND B. ETHERIDGE (2010): "Consumption, income and earnings inequality in Britain," *Review of Economic Dynamics*, 13(1), 76–102.

BLUNDELL, R., A. GOSLING, H. ICHIMURA, AND C. MEGHIR (2007): "Changes in the distribution of male and female wages accounting for employment composition using bounds," *Econometrica*, 75(2), 323–363.

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.

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, 709–742.

BRZOZOWSKI, M., M. GERVAIS, P. KLEIN, AND M. SUZUKI (2010): "Consumption, income, and wealth inequality in Canada," *Review of economic dynamics*, 13(1), 52.

CASE, A., AND A. DEATON (2017): "Mortality and morbidity in the 21st century," *Brookings papers on economic activity*, 2017, 397.

CHETTY, R., M. STEPNER, S. ABRAHAM, S. LIN, B. SCUDERI, N. TURNER, A. BERGERON, AND D. CUTLER (2016): "The association between income and life expectancy in the United States, 2001-2014," *Journal of the American Medical Association*, 315(16), 1750–1766.

CUNHA, F., AND J. HECKMAN (2007): "The technology of skill formation," *American Economic Review*, 97(2), 31–47.

CUNHA, F., J. HECKMAN, AND S. NAVARRO (2005): "Separating uncertainty from heterogeneity in life cycle earnings," *Oxford Economic Papers*, 57(2), 191–261.

DINARDO, J., N. M. FORTIN, AND T. LEMIEUX (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64(5), 1001–1044.

DOMEIJ, D., AND M. FLODEN (2010): "Inequality trends in Sweden 1978–2004," *Review of Economic Dynamics*, 13(1), 179–208.

DYRDA, S., AND B. PUGSLEY (2018): "Taxes, private equity and evolution of income inequality in the us," Discussion paper, University of Toronto.

ECKSTEIN, Z., E. NAGYPAL, ET AL. (2004): "The evolution of US earnings inequality: 1961-2002," *Federal Reserve Bank of Minneapolis Quarterly Review*, 28(2), 10–29.

FRIEDMAN, M. (1957): *A Theory of the Consumption Function*. Princeton University Press.

FUCHS-SCHÜNDELN, N., D. KRUEGER, AND M. SOMMER (2010): "Inequality trends for Germany in the last two decades: A tale of two countries," *Review of Economic Dynamics*, 13(1), 103–132.

GORODNICHENKO, Y., K. SABIRIANOVA-PETER, AND D. STOLYAROV (2010): "Inequality and volatility moderation in Russia: Evidence from micro-level panel data on consumption and income," *Review of Economic Dynamics*, 13(1), 209–237.

GOSLING, A., S. MACHIN, AND C. MEGHIR (2000): "The changing distribution of male wages in the UK," *The Review of Economic Studies*, 67, 635–666.

GOTTSCHALK, P., R. MOFFITT, L. F. KATZ, AND W. T. DICKENS (1994): "The growth of earnings instability in the US labor market," *Brookings Papers on Economic Activity*, 1994(2), 217–272.

GREEN, D. A., AND B. M. SAND (2015): "Has the Canadian labour market polarized?," *Canadian Journal of Economics/Revue canadienne d'économique*, 48(2), 612–646.

GUVENEN, F., G. KAPLAN, J. SONG, AND J. WEIDNER (2017): "Lifetime incomes in the United States over six decades," Discussion paper, National Bureau of Economic Research.

GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2015): "What do data on millions of US workers reveal about life-cycle earnings risk?," Discussion paper, National Bureau of Economic Research.

GUVEVEN, F., AND G. KAPLAN (2017): "Top Income Inequality in the 21st Century: Some Cautionary Notes," mimeo.

HAIDER, S., AND G. SOLON (2006): "Life-cycle variation in the association between current and lifetime earnings," *American Economic Review*, 96(4), 1308–1320.

HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967 to 2006," *Review of Economic Dynamics*, 13(1), 15–51.

HEATHCOTE, J., F. PERRI, G. L. VIOLANTE, ET AL. (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., K. STORESLETTEN, AND G. L. VIOLANTE (2010): "The Macroeconomic Implications of Rising Wage Inequality in the US," *Journal of Political Economy*, 118(4), 681–72.

HECKMAN, J. (2004): "Skill Formation and the Economics of Investing in Disadvantaged Children," *Science*, 312 5782, 1900–1902.

HECKMAN, J. J. (2000): "Policies to foster human capital," *Research in economics*, 54(1), 3–56.

HUBMER, J., P. KRUSELL, AND A. A. SMITH JR (2016): "The historical evolution of the wealth distribution: A quantitative-theoretic investigation," Discussion paper, National Bureau of Economic Research.

HUGGETT, M., AND G. KAPLAN (2016): "How large is the stock component of human capital?," *Review of Economic Dynamics*, 22, 21–51.

HUGGETT, M., G. VENTURA, AND A. YARON (2006): "Human capital and earnings distribution dynamics," *Journal of Monetary Economics*, 53(2), 265–290.

JAPPELLI, T., AND L. PISTAFERRI (2017): *The economics of consumption: theory and evidence*. Oxford University Press.

JORGENSON, D., AND B. M. FRAUMENI (1989): "The Accumulation of Human and Nonhuman Capital," in *The Measurement of Saving, Investment, and Wealth*, ed. by R. E. Lipsey, and H. S. Tice, chap. 5, pp. 227–286. University of Chicago Press.

KAYMAK, B., AND M. POSCHKE (2016): "The evolution of wealth inequality over half a century: The role of taxes, transfers and technology," *Journal of Monetary Economics*, 77, 1–25.

KOPCZUK, W., E. SAEZ, AND J. SONG (2010): "Earnings inequality and mobility in the United States: evidence from social security data since 1937," *The Quarterly Journal of Economics*, 125(1), 91–128.

LEMIEUX, T. (2006a): "Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill?," *American Economic Review*, 96(3), 461–498.

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

——— (2006c): "Postsecondary Education and Increasing Wage Inequality," *The American Economic Review*, 96, 195–199.

LILLARD, L. A. (1977): "Inequality: Earnings vs. Human Wealth," *The American Economic Review*, pp. 42–53.

LOCHNER, L., AND Y. SHIN (2014): "Understanding earnings dynamics: Identifying and estimating the changing roles of unobserved ability, permanent and transitory shocks," Discussion paper, National Bureau of Economic Research.

LOW, H., C. MEGHIR, AND L. PISTAFERRI (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100(5), 1432–67.

MEGHIR, C., AND L. PISTAFERRI (2006): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

MILLIGAN, K., AND T. SCHIRLE (2018): "The Evolution of Longevity: Evidence from Canada," mimeo.

MODIGLIANI, F., AND R. BRUMBERG (1954): "Utility analysis and the consumption function: An interpretation of cross-section data," *Franco Modigliani*, 1, 388–436.

MOFFITT, R. A., AND P. GOTTSCHALK (2012): "Trends in the transitory variance of male earnings methods and evidence," *Journal of Human Resources*, 47(1), 204–236.

MORELLI, S., T. SMEEDING, AND J. THOMPSON (2015): "Post-1970 Trends in Within-Country Inequality and Poverty: Rich and Middle-Income Countries," in *Handbook of Income Distribution*, ed. by A. B. Atkinson, and F. Bourguignon, vol. 2 of *Handbook of Income Distribution*, pp. 593 – 696. Elsevier.

PIJOAN-MAS, J., AND V. SÁNCHEZ-MARCOS (2010): "Spain is different: Falling trends of inequality," *Review of Economic Dynamics*, 13(1), 154–178.

QUADRINI, V., AND J.-V. RÍOS-RULL (2015): "Inequality in Macroeconomics," in *Handbook of Income Distribution*, ed. by A. B. Atkinson, and F. Bourguignon, vol. 2 of *Handbook of Income Distribution*, pp. 1229 – 1302. Elsevier.

RIOS-RULL, J., AND M. KUHN (2013): "Update on the US earnings, income, and wealth distributional facts: A view from Macroeconomics," *Quarterly Review,(April)*, pp. 1–75.

ROINE, J., AND D. WALDENSTRÖM (2015): "Long-Run Trends in the Distribution of Income and Wealth," in *Handbook of Income Distribution*, ed. by A. B. Atkinson, and F. Bourguignon, vol. 2 of *Handbook of Income Distribution*, pp. 469 – 592. Elsevier.

SMITH, M., D. YAGAN, O. ZIDAR, AND E. ZWICK (2018): "Capitalists in the Twenty-First Century," mimeo.

SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2015): "Firming up inequality," Discussion paper, National Bureau of Economic Research.

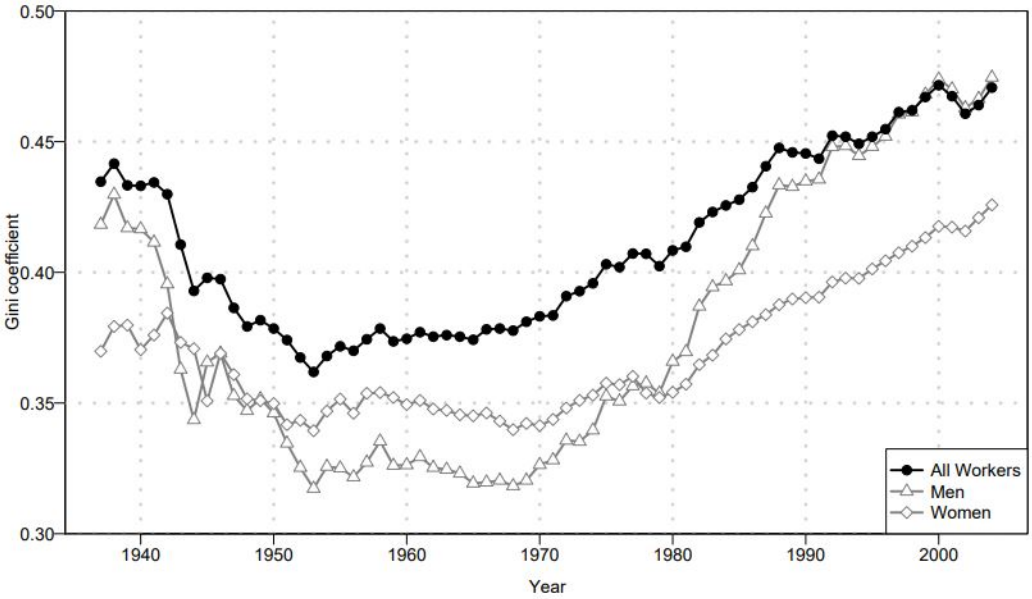

Figure 1: This is a reproduction of Figure 1 in Kopczuk, Saez, and Song (2010). It displays the evolution of the earnings Gini index in SSA data from 1937 to 2004, both in aggregate and by gender.

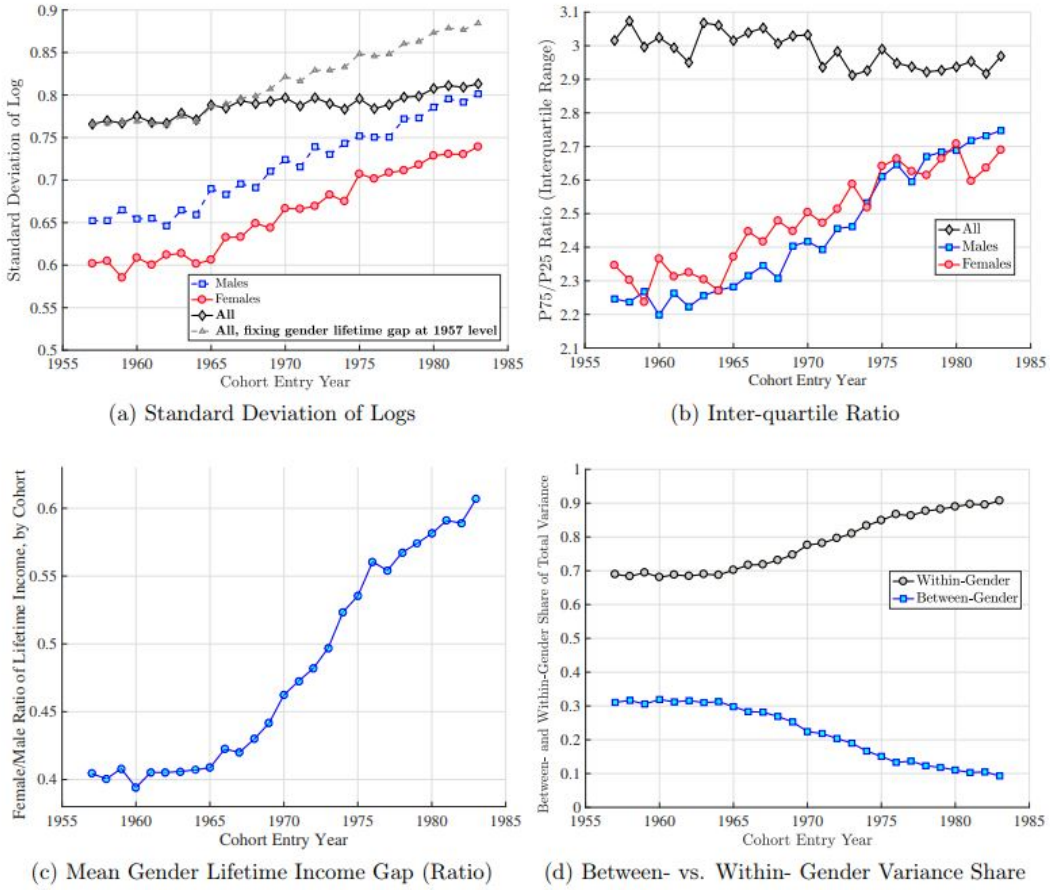

Figure 2: This is a reproduction of Figure 11 in Guvenen, Kaplan, Song, and Weidner (2017). Panel (a) displays general patterns of earnings inequality measured by the standard deviation of log-earnings across genders and over time. Panel (b) displays 75-25 earnings ratios by gender and overall, showing how the pattern changes significantly if one considers only one gender or all workers. Panel (c) displays the evolution of the gender earnings gap, and panel (d) displays how within-gender inequality has risen while between-gender inequality has fallen.

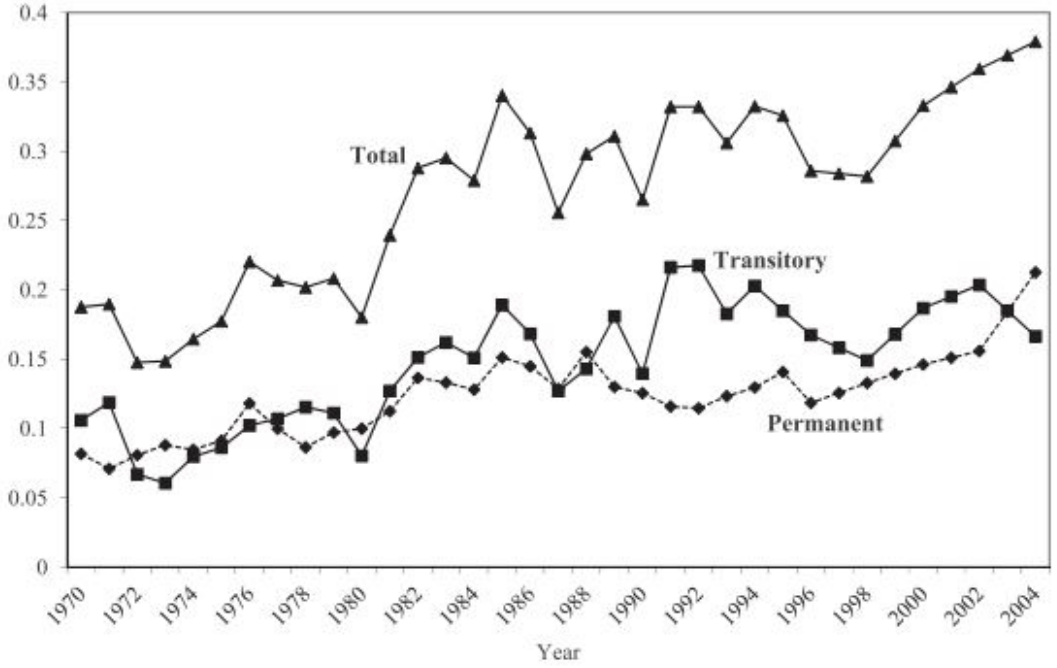

Figure 3: This is a reproduction of Figure 3 in Moffitt and Gottschalk (2012). It displays the evolution of the total variance of log-earnings' residuals from 1970 to 2004, as well as the evolution of both its permanent and transitory components.