NBER WORKING PAPER SERIES

HOURS WORKED AND LIFETIME EARNINGS INEQUALITY

Alexander Bick
Adam Blandin
Richard Rogerson

Working Paper 32997
http://www.nber.org/papers/w32997

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024, revised April 2025

We thank Jésus Fernández-Villaverde and Kjetil Storesletten in addition to seminar/conference participants at the Bureau of Labor Statistics, Deutsche Bundesbank, Federal Reserve Banks of Minneapolis and St. Louis, Goethe University, I-85 Macro Workshop, LAEF Workshop on Labor Markets and Macroeconomic Outcomes, LAEF and St. Louis Fed Human Capital Conference, LMU Munich, NBER Conference on Economic Activity Across Time and Space, NBER Summer Institute Micro Data and Macro Models, Society for Economic Dynamics Annual Meeting in Barcelona, Stanford Institute for Theoretical Economics Labor Market and Policies Workshop, University of Connecticut, Carleton University, University of Bonn, University of North Carolina at Charlotte, MIT, Princeton University, University of Pennsylvania and Vanderbilt University for helpful comments and suggestions. Kevin Bloodworth II, Elizabeth Harding, and Siyu Shi provided outstanding research assistance. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w32997

Hours Worked and Lifetime Earnings Inequality
Alexander Bick, Adam Blandin, and Richard Rogerson
NBER Working Paper No. 32997
September 2024, revised April 2025
JEL No. D15, E20, J22, J24

ABSTRACT

We document large differences in lifetime hours of work using data from the NLSY79 and argue that these differences are an important source of inequality in lifetime earnings. To establish this we develop and calibrate a rich heterogeneous agent model of labor supply and human capital accumulation that allows for heterogeneity in preferences for work, initial human capital and learning ability, as well as idiosyncratic shocks to human capital throughout the life-cycle. Our calibrated model implies that almost 20 percent of the variance in lifetime earnings is accounted for by differences in lifetime hours of work, with over 90 percent of this effect due to heterogeneity in preferences. Higher lifetime hours contribute to lifetime earnings via two channels: a direct channel (more hours spent in production at given productivity) and a human capital channel (more hours spent investing in human capital, which increases future productivity). Roughly one-half of the effect of lifetime hours on lifetime earnings is due to the human capital channel. Higher lifetime hours are also an important source of upward earnings mobility over the life-cycle for many workers.

Alexander Bick
One Federal Reserve Bank Plaza
St. Louis, MO 63102
USA
alexander.bick@mailbox.org

Adam Blandin
Department of Economics
Vanderbilt University
VU Station B #351819
2301 Vanderbilt Place
Nashville, TN 37235-1819
adam.blandin@vanderbilt.edu

Richard Rogerson
Department of Economics
JRR Building Room 292
Princeton University
Princeton, NJ 08544
and NBER
rdr@princeton.edu

1 Introduction

What are the quantitatively important forces that shape earnings inequality? One basic determinant of an individual's earnings is the time they spend working, and it is well documented that cross-sectional variation in hours worked is an important contributor to cross-sectional earnings inequality.[1] But economists have long recognized that, for many questions, inequality in lifetime earnings is more relevant than inequality at a given point in time; see, e.g., Farr (1853). While recent empirical work has leveraged administrative data to document inequality in lifetime earnings (e.g., Guvenen, Kaplan, Song, and Weidner (2022)), little is known about the extent of differences in lifetime hours of work. The goal of this paper is to document inequality in lifetime hours of work and assess its contribution to lifetime earnings inequality.

We study lifetime earnings from the perspective of human capital theory. Textbook human capital models robustly predict that current hours affect future wages. Assessing the impact of lifetime hours on lifetime earnings therefore requires not only data on lifetime hours, but also a model that accounts for the effect of life-cycle hours on life-cycle wages. Our analysis extends previous life-cycle human capital models to capture differences in lifetime hours.

Our empirical analysis uses the National Longitudinal Study of Youth 1979 (hereafter, NLSY79) to build a balanced panel of individuals from ages 25 to 55 with annual data on earnings and hours worked. We document four key properties. First, dispersion in lifetime hours is large: the interquartile range for lifetime hours is about 25 percent of median lifetime hours for men and 50 percent for women. Second, annual hours worked at the individual level are persistent, even at very long horizons, and are not well-approximated by an AR(1) process. Third, the relative importance of years worked, weeks worked per year, and hours worked per week varies systematically across the lifetime hours distribution: in the lower part of the lifetime hours distribution, the dominant margin is years worked, while in the upper part of the distribution the dominant margin is hours per week. Fourth, there is a strong positive relationship between lifetime hours and life-cycle earnings growth. These patterns are very similar for men and women, except that women's hours have a lower mean and greater variance.

We show analytically in a simplified Ben-Porath model that higher future hours of work increases the incentive for human capital accumulation today.[2] As a result, differences in expected future hours lead to differences in human capital accumulation, which in turn generate additional variation in lifetime earnings beyond the direct effect of (future) hours worked. We then embed

[1]For example, in the 2024 March CPS, 27.9 percent of the variance in log annual earnings for men aged 25–55 comes from the variance of log annual hours and 2.0 percent from the covariance between log annual hours and log wages.

[2]A similar result is shown in Neal and Rosen (2000) and would also arise in a model where human capital accumulates through learning by doing. While our analysis could be conducted within such a framework, we adopt a Ben-Porath specification to facilitate comparison with the existing literature.

this mechanism into a rich heterogeneous agent life-cycle model. Like the existing literature, we allow for heterogeneity in initial human capital and learning ability, as well as permanent shocks to human capital. A key feature of our model is that we introduce heterogeneity in tastes for work in order to replicate the salient facts about hours heterogeneity. These taste differences may capture a variety of factors, including differences in health, the value of leisure, the productivity of non-market time, family demands on time, and social norms. We calibrate the parameters of this model to match a large set of moments that characterize the distribution of earnings and hours over the life-cycle for a sample of highly attached men that represents 80 percent of our overall male sample. Importantly, we include both transitory and permanent differences in the taste for work to capture both the transitory and permanent components of cross-sectional variation in hours.

We use the calibrated model to quantitatively assess the factors that generate inequality in lifetime earnings and hours. Five key results emerge. First, eliminating hours heterogeneity reduces the variance of log lifetime earnings by nearly one-fifth (18.3 percent). Second, roughly one-half of this effect is attributed to human capital accumulation. Third, preference heterogeneity accounts for virtually all of the dispersion in lifetime hours. That is, heterogeneity in initial human capital, learning ability, and idiosyncratic shocks to human capital effectively generate no dispersion in lifetime hours. Fourth, the contribution of hours to lifetime earnings inequality operates largely through the permanent component of preference heterogeneity. If we eliminate the permanent component of preference heterogeneity, our model still generates dispersion in lifetime hours, but this variation in lifetime hours contributes essentially nothing to the variation in lifetime earnings. Fifth, working long hours is an important channel of upward earnings mobility over the life-cycle for many workers with low initial human capital.

Our paper contributes to several related strands of literature. Our empirical work relates to Morchio (2020) and Glover, Mustre-del Río, and Pollard (2022). Morchio (2020) uses the NLSY79 to study the distribution of lifetime unemployment for men and shows that a small share of individuals account for a large share of unemployment spells. Consistent with this, we find that variation in years and weeks worked is an important source of variation in lifetime hours for a relatively small part of the male population. Morchio (2020) does not study the contribution of lifetime unemployment to lifetime earnings inequality and his model does not feature human capital. Glover et al. (2022) document that differences in lifetime years of work between black and white men are an important contributor to the lifetime earnings differences between those groups, even after conditioning on differences in education.[3]

Our paper complements the analysis in Hosseini, Kopecky, and Zhao (2025). Like us, they study how differences in lifetime labor supply contribute to inequality in lifetime earnings. While they highlight the role of health shocks via their impact on participation rates of older individuals, we highlight the role of permanent heterogeneity on labor supply along the intensive margin among

[3]Rauh and Valladares-Esteban (2023) emphasize the importance of differences in initial human capital and learning ability to account for the black-white wage gap for men in a life-cycle model with learning by doing.

highly attached males aged 25-55. Our analysis stresses the importance of endogenous human capital accumulation, whereas they assume that human capital accumulation is exogenous.

Kaplan (2012) and Heathcote, Storesletten, and Violante (2014) study hours heterogeneity in life-cycle models. Kaplan (2012) specifically focuses on the reduction in cross-sectional variance of hours in the early part of the life-cycle, an issue that we do not explicitly address, as this is not such a prominent feature in our sample of men with high labor market attachment. Heathcote et al. (2014) study cross-sectional inequality in both wages and hours over the life-cycle. Both papers focus on differences in hours in the cross-section and do not consider the distinction between variation in hours at a point in time and variation in lifetime hours. Neither paper allows for human capital accumulation, nor do they study inequality in lifetime earnings.

Several papers use Ben-Porath models to quantitatively study aspects of inequality in a life-cycle setting. Among these, our quantitative analysis is most similar to Huggett, Ventura, and Yaron (2006, 2011). Like us, they study lifetime income inequality through the lens of a heterogeneous agent life-cycle model that features a Ben-Porath human capital accumulation technology. But they assume no variation in hours worked either across individuals or over time. Relative to them, our key contribution is to assess the role of hours inequality for lifetime earnings inequality.

Our analysis is also similar to the contemporaneous and independent study by Fillmore and Gallen (2023). They use the NLSY79 to document heterogeneity in mean total hours for men between the ages 30 and 44 who work in at least 14 out of these 15 years. In contrast, our empirical analysis covers a longer age span (25-55), considers both men and women, and excludes only individuals who essentially never work during these 31 years. Their quantitative work also uses a Ben-Porath model to understand how heterogeneity in hours affects heterogeneity in earnings via its effect on human capital, but abstracts from shocks to both human capital and preferences. This likely explains why they find a much stronger role for hours worked than we do. When they eliminate preference heterogeneity, the variance of earnings inequality at age 44 drops by 75 percent, versus 19 percent in our model.

Guvenen, Kuruscu, and Ozkan (2014) and Erosa, Fuster, and Restuccia (2016) both use a Ben-Porath model with endogenous labor supply to study aspects of cross-sectional wage inequality.[4] Guvenen et al. (2014) use a heterogeneous agent model to study the effect of progressive taxation on human capital investment and the evolution of cross-sectional wage inequality over the life-cycle. Although they include an endogenous decision about hours of work, they do not include any preference heterogeneity. We show that their model generates essentially none of the variation

[4] A large literature studies labor supply in models of human capital accumulation. While not unrelated to inequality, many of these papers have focused on labor supply elasticities and the effects of policy. Heckman (1976a,b) and Shaw (1989) are early contributions. Subsequent contributions of note include Keane and Wolpin (1997), Imai and Keane (2004), Wallenius (2011, 2013), Keane and Wasi (2016), Blundell, Dias, Meghir, and Shaw (2016)), and Blundell, Costa-Dias, Goll, and Meghir (2021). Stantcheva (2017) and Badel, Huggett, and Luo (2020) study optimal taxation of labor income in models with human capital accumulation.

in lifetime hours found in the data. Relative to them, our key contribution is to assess the role of heterogeneity in lifetime hours on human capital accumulation. Relative to Erosa et al. (2016), our quantitative analysis adds several important features: endogenous variation of hours on the intensive margin, permanent and transitory preference heterogeneity, heterogeneity in learning ability and initial human capital, and permanent shocks to human capital. Additionally, whereas they focus on across-group differences (how gender gaps in hours lead to gender gaps in wages), we focus on within-group differences for men. Importantly, our empirical work shows that the interquartile range for lifetime hours among men is even larger than the gender gap in median lifetime hours.

A key mechanism in our model is the effect of lifetime hours on human capital accumulation. This mechanism is also central in Manuelli, Seshadri, and Shin (2012) and Fan, Seshadri, and Taber (2024). These papers use Ben-Porath models to study how tax and transfer programs affect human capital accumulation by changing retirement decisions. Whereas we abstract from endogenous retirement decisions, these papers abstract from the intensive margin of labor supply.[5] Neither of these papers studies lifetime earnings inequality.

A distinct strand of literature studies lifetime earnings inequality in life-cycle models featuring frictional labor market models with heterogeneous firms and workers. Examples include Bagger, Fontaine, Postel-Vinay, and Robin (2014) and Ozkan, Song, and Karahan (2023).[6] We view our work as complementary to these papers. Whereas they abstract from heterogeneity in hours worked and model human capital as an exogenous process in order to focus on the role of job ladder dynamics in a frictional labor market, we capture job ladder dynamics as exogenous shocks to human capital in order to focus on the role of endogenous choices of hours and investment in human capital.

An outline of the paper follows. Section 2 describes how we use the NLSY79 to create a balanced panel of observations on annual hours and annual earnings. Section 3 documents properties of lifetime hours and earnings using the NLSY79. Section 4 establishes analytically that higher future hours lead to higher human capital accumulation in a simple Ben-Porath model. Section 5 describes our quantitative model and Section 6 calibrates it. Section 7 presents our main results about the role of lifetime hours for lifetime earnings inequality. Section 8 concludes.

[5] Fan et al. (2024) briefly consider an extension which distinguishes between part-time and full-time work. This margin is not important for highly attached men, which is the focus of our quantitative analysis.

[6] Other recent examples of papers that use search models to study heterogeneous earnings profiles include Papageorgiou (2014); Lise and Postel-Vinay (2020); Jarosch, Oberfield, and Rossi-Hansberg (2021); Jarosch (2023); Herkenhoff, Lise, Menzio, and Phillips (2024).

2 Data

This section describes how we use the NLSY79 to create a balanced panel of individual observations on annual earnings and annual hours worked from age 25 to 55.

2.1 The NLSY79

The NLSY79 is a longitudinal study of 12,686 individuals born between 1957 and 1964.[7] Respondents were recruited and initially interviewed in 1979, when they were between 14 and 22 years old. Interviews were conducted annually through 1994, after which they were conducted biennially, occurring in even-numbered years. We use data up to and including the 2020 interview year. This interview provides earnings data for 2019, at which point the youngest individuals in the sample are 55 and the oldest are 63. All statistics reported in the paper use the initial NLSY79 sample weights.

Each NLSY79 interview records the start and end dates of all jobs held since the individual's most recent interview, as well as the individual's usual weekly hours worked at each job. Therefore, even though interviews are biennial after 1994, we can construct a weekly history of employment status and hours worked spanning 1978-2019. We aggregate this information to produce annual measures of weeks worked, usual weekly hours, and total annual hours.

Respondents report annual earnings for the calendar year preceding the interview year. Thus, whereas employment histories are collected for all years, earnings are not collected for even-numbered years starting in 1994. Information is collected separately for two categories of earnings: (i) income from wages, salary, commissions, or tips from all jobs before deductions for taxes or anything else last year and (ii) income received from a farm/business owned last year. Our measure of earnings is the sum of these two components. Following Guvenen et al. (2022) we deflate earnings with the Personal Consumption Expenditures index normalized to one in 2013. We measure hourly wages as annual earnings divided by annual hours.[8]

[7] In addition to this core sample there were two supplemental samples (a military sample and an economically disadvantaged non-Black, non-Hispanic youth sample) that were subsequently discontinued. 201 respondents randomly selected from the military sample remained in the survey. We keep these 201 individuals in our baseline sample but do not use any other respondents from the discontinued samples.

[8] Following the procedure in Bick, Blandin, and Rogerson (2025), we make two adjustments to control for hourly wage outliers. First, wages below half of the Federal minimum wage are set equal to that value, with earnings adjusted accordingly. Second, motivated by the evidence in Bick et al. (2025), we treat wage observations in the top 0.1 percent of the wage distribution as due to misreported hours, and set weeks worked and weekly hours to missing for these observations.

2.2 Creating a Balanced Panel

We use this annual data at the individual level to create a balanced panel that covers individuals from age 25 to 55. We stop at age 55 because that is the age of the youngest members of the NLSY79 in 2019, and so dictates the longest balanced panel that we can create. We start at age 25 to focus on outcomes after formal education is complete for most individuals. Guvenen et al. (2022) and Ozkan et al. (2023) also focused on the age range of 25 to 55 in their analysis of lifetime earnings.

To create this balanced panel we need to address the issue of missing values. In addition to the years in which earnings data are not collected, there is also missing data for individual responses to hours and/or earnings questions in some interviews. We impute missing values using the interpolation procedure described in Bick et al. (2025). Given that we rely on interpolation to fill in missing values and our goal is to have a balanced panel through age 55, we remove from our sample any individuals who do not have a complete interview at age 55 or older. In particular, this will remove all individuals who leave the sample prior to age 55. We also exclude individuals who lack sufficiently nearby observations for use in the imputation procedure. Specifically, we impose for any year in which employment status, hours worked, or earnings are missing for the entire year, there must be at least one observation within the previous or next five years.[9]

Implementing our procedure produces a balanced panel data set with annual data on weeks worked, usual weekly hours, total hours, earnings, and average hourly wages for individuals between ages 25 and 55. We drop 74 individuals from our sample that work less than one hour per year on average. This leaves us with a sample of 6261 individuals.

For ease of interpretation, in what follows we will report all lifetime measures as annualized values. That is, we measure lifetime hours and lifetime earnings by summing values over the 31 years of our panel and dividing by 31. This is the same measure of lifetime earnings that has been used by Guvenen et al. (2022) and Ozkan et al. (2023).[10]

2.3 Validating the Balanced Panel

One potential concern with the balanced panel we create is that the effects of attrition and our selection criterion based on missing values may not be random. In Bick et al. (2025), we argue that selection effects are not significant and that our imputation method works well. In particular,

[9] In Bick et al. (2025), we document the effects of varying this threshold and conclude that five years strikes a reasonable balance between maximizing sample size and minimizing measurement error in our imputation procedure.

[10] This method implicitly assumes a zero interest rate when computing lifetime earnings. We have also computed lifetime earnings using an interest rate of four percent. This had essentially no impact on any of our results.

Figure 1: Comparison of the NLSY79 and SSA Data by Ozkan et al. (2023)

(a) Lifetime Earnings (b) Life-cycle Earnings Growth

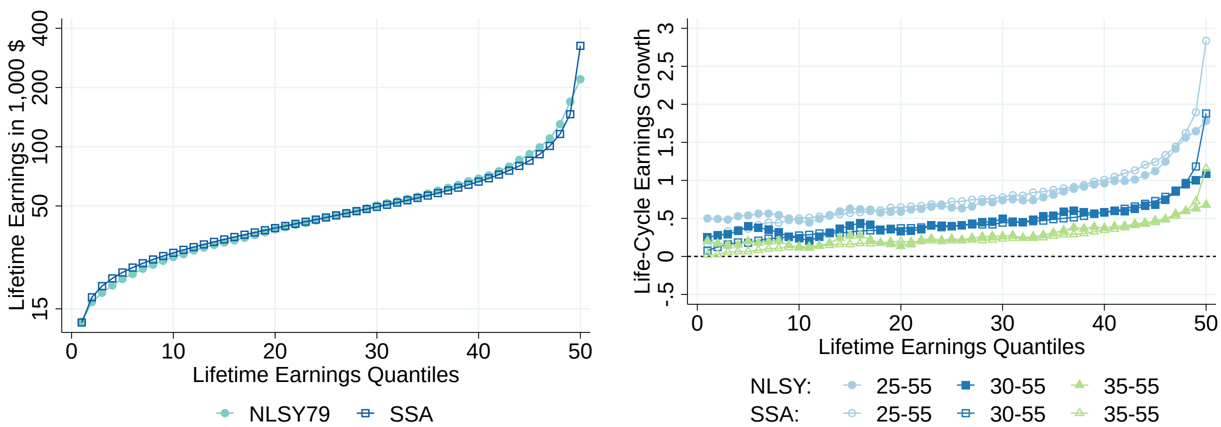

Notes: SSA data are from Figure 1 in Ozkan et al. (2023). In Figure 1a, we normalize average annualized lifetime earnings in the SSA data so that median earnings match those of the NLSY79 sample. In Figure 1b, life-cycle earnings growth is defined as the log difference between average earnings at age 55 and average earnings at ages 25, 30, and 35, respectively. For the NLSY79, we plot a centered moving average with two lags and leads of these growth rates over the lifetime earnings quantile distribution.

we show that life-cycle profiles of employment, hours worked, and earnings (both means and standard deviations, where applicable) in our balanced panel are comparable to those in the Current Population Survey (CPS) for the same cohorts. We also show that the balanced panel closely aligns with both men's and women's lifetime earnings distributions in Social Security Administration (SSA) data, as documented by Guvenen et al. (2022). This is particularly reassuring given the larger sample size and administrative nature of their data, which does not rely on self-reported earnings.

Figure 1 provides an additional comparison between our sample and SSA data. Ozkan et al. (2023) produce estimates for the distribution of lifetime earnings and life-cycle earnings growth for US birth cohorts born between 1953-1960. We apply their sample selection criterion to our sample of birth cohorts born from 1957-1964. This leaves us with 47.1 percent of the initial balanced panel sample, very similar to the 45.5 percent reported by Ozkan et al. (2023) for the SSA data.

Figure 1a displays lifetime earnings distributions in each sample. Median lifetime earnings in the SSA data are higher than in our sample (by 14 percent). To focus on earnings inequality, the figure normalizes the SSA data so that median earnings match that of the NLSY sample. The resulting distributions essentially lie on top of one another, with the exception of the top 2 percent of earners, which are higher in the SSA data.

Figure 1b plots earnings growth from ages 25-55, 30-55, and 35-55 by lifetime earnings quantiles. Given the smaller sample size of the NLSY79, we present results using a centered moving average with two lags and leads over the lifetime earnings quantile distribution. With the exception of somewhat higher growth rates for individuals with very low lifetime earnings and somewhat lower growth rates for individuals with very high lifetime earnings, the patterns between the two datasets align closely.

Taken together, the CPS comparisons of life-cycle profiles of employment, hours worked, and earnings, along with SSA comparisons of lifetime earnings and earnings growth, suggest that our NLSY79 balanced panel is nationally representative along the key dimensions of interest.

3 Facts About Lifetime Hours and Lifetime Earnings

In this subsection we document four key facts. First, dispersion in lifetime hours is large. Second, hours at the individual level are persistent, even at long horizons, and are not well approximated by an AR(1) process. Third, the importance of variation along three margins—years with positive hours, weeks worked per year and hours worked per week worked—varies systematically along the lifetime hours distribution. Fourth, there is a positive relationship between lifetime hours and life-cycle earnings growth, and the reduced-form elasticity of lifetime earnings with respect to lifetime hours is significantly greater than one. We document that these facts hold for both men and women, with the main difference being that the distribution of lifetime hours for women has a lower mean and higher variance.

3.1 Dispersion in Lifetime Hours

The left panel of Figure 2 displays the distribution of lifetime hours for both men and women when we sort individuals into hours bins that are 250 hours wide. The right panel shows the cumulative distribution function.

The distribution of lifetime hours for women is shifted to the left relative to men; median lifetime hours for women (1761 hours per year) are 20 percent lower than median lifetime hours for men (2189). The distribution of lifetime hours for women also displays more dispersion than the distribution for men; the standard deviation of log lifetime hours is 0.39 for men and 0.55 for women. But importantly, the amount of dispersion is large for both men and women: the interquartile ranges for men and women are 525 and 839, respectively. For men this value is almost 25 percent of median hours, and for women it is nearly half of median hours.

We note that the dispersion in annual hours in the cross-section is even larger than the dispersion in lifetime hours. To show this we compare dispersion in lifetime hours with dispersion in the pooled sample of observations on annual hours from our balanced panel. (Appendix Figure B.1

Figure 2: Lifetime Hours Worked Distribution

(a) PDF of Lifetime Hours (b) CDF of Lifetime Hours

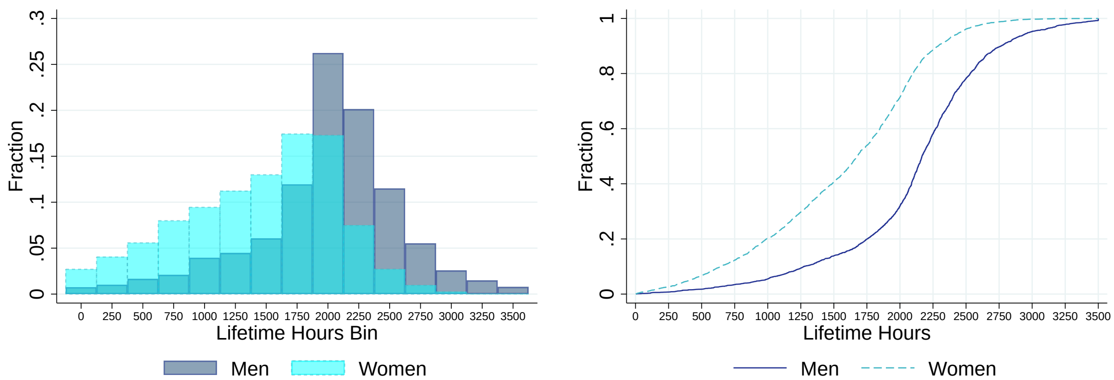

Notes: In panel (a), each lifetime hours bin starts at the value displayed, e.g., the 250 hours bin includes anyone with lifetime hours from 250 to 499. The 0 hours bin includes anyone with lifetime hours from 1 to 249 (as discussed previously we drop anyone with lifetime hours below 1), and the 3500 hours bin anyone with lifetime hours of at least 3500.

compares the full distributions directly). Because the annual data include some zero values, we compute the coefficient of variation for hours rather than the standard deviation of log hours. We find that the coefficient of variation for pooled annual hours is more than 50 percent larger than the coefficient of variation for lifetime hours for both men and women. A simple message is that the magnitude of dispersion in annual hours is a poor proxy for the dispersion in lifetime hours.

3.2 Persistence of Hours Worked Over Time

The fact that lifetime hours are less dispersed than annual hours implies that some differences in annual hours are transitory and tend to average out over time. To quantify this, we exploit the long panel dimension of our sample to compute the autocorrelation of hours at lags ranging from one to twenty years. To construct autocorrelations at lag length $t > 0$, we collect all length-t pairs of hours $\{(h_{i,a}, h_{i,a+t})\}_{a=25}^{55-t}$ from our sample of individual hours profiles $\{h_{i,a}\}_{a=25}^{55}$ and compute the pairwise correlation of this collection. Figure 3 plots the resulting correlation coefficients for lag lengths $t = 1, ..., 20$ for both men and women.[11]

Figure 3 shows that although there is some tendency for mean reversion, hours display considerable persistence over long horizons. The one year autocorrelation of hours is 0.81 for men and 0.84 for women. These values fall by about 40 percent, to 0.50 and 0.52, when moving to a lag length of five years, after which it continues to decrease, but at a slower rate. At twenty years, the

[11] One may worry that the survey design and our imputation procedure creates too much persistence in annual hours worked. We address this concern in Appendix Figure B.2, which shows that the autocorrelation profile excluding imputed values for annual hours is virtually identical to the profile shown in Figure 3.

Figure 3: Autocorrelation of Annual Hours worked

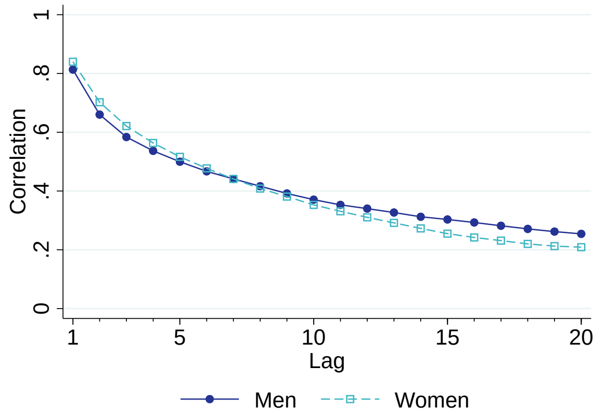

Notes: Person-year observations with zero hours are included.

autocorrelation is still substantially above zero, at 0.25 and 0.21, for men and women, respectively. The autocorrelations at one- and twenty-year horizons are not consistent with a simple AR(1) process. AR(1) processes with autocorrelations of 0.81 and 0.84 at a lag of one year would have an autocorrelation of effectively zero at a twenty-year lag (to be precise, 0.01 and 0.03, respectively). This observation will influence modeling choices later in the paper.

3.3 Components of Lifetime Hours: Years, Weeks per Year, and Hours per Week

Differences in lifetime hours may arise because of differences along three distinct margins: years of work (number of years with positive hours), average weeks worked in years with positive hours, and average usual hours worked per week worked. Panels (a) and (b) of Figure 4 show how the role of these three margins varies across the lifetime hours distribution for men and women, respectively. In these figures each variable is normalized by its respective gender-specific value in the 1750 lifetime hours bin (1750 lifetime hours corresponds to the 20th percentile of men's lifetime hours and the 54th percentile for women). Panel (c) reports this decomposition in levels for three selected lifetime hours bins.

The figure exhibits very similar patterns for women and men. Below 1750 hours per year, differences in lifetime hours are primarily due to differences in years worked. For example, men in the 500 lifetime hours bin work 61 percent fewer years relative to men in the 1750 hour bin, compared with 28 percent fewer weeks per year worked and 13 percent fewer hours per week worked. By contrast, above 1750 hours per year, differences in lifetime hours are almost entirely due to differences in hours per week worked. For example, men in the 3000 lifetime hour bin work

Figure 4: The Distribution of Lifetime Hours and Its Components

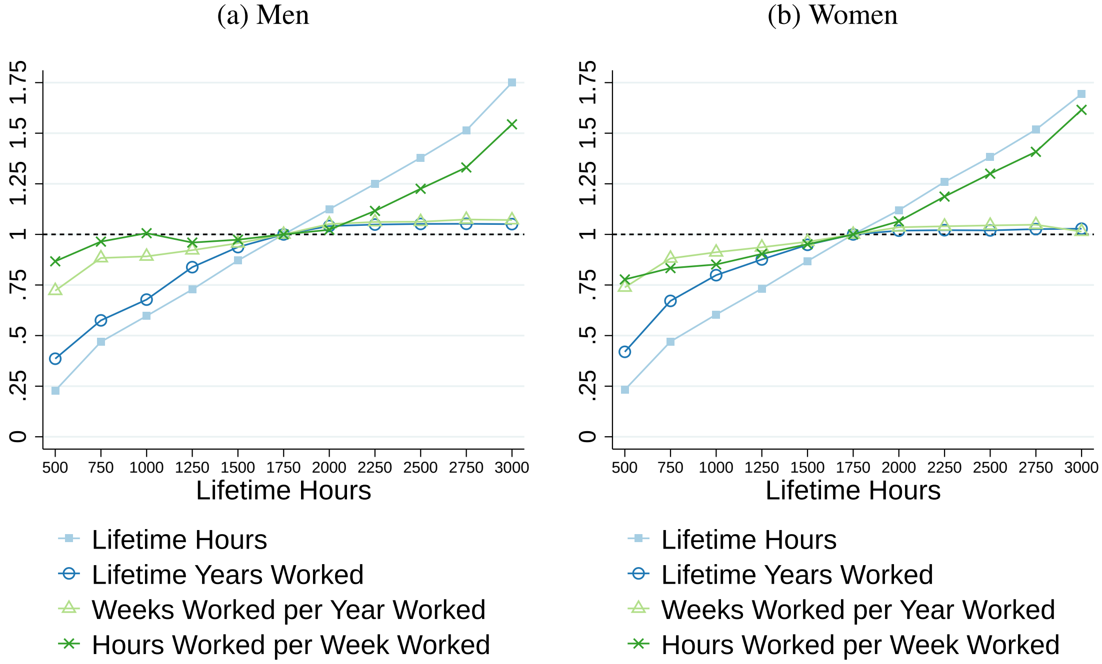

(c) Level Information for Selected Lifetime Hours Bins

	Men			Women		
	500	1750	3000	500	1750	3000
Annualized Lifetime Hours	432.1	1892.1	3310.9	436.9	1880.3	3184.8
Lifetime Years Worked	11.3	29.3	30.8	12.7	30.2	31.0
Weeks Worked per Year Worked	34.4	47.6	51.0	35.8	48.6	49.2
Hours Worked per Week Worked	36.6	42.2	65.2	30.9	39.7	64.2

Notes: Figure 4 displays each variable normalized by the respective gender-specific value in the 1750 hours bin. The 500 hours bin includes anyone with lifetime hours between 1 to 749 hours, and the 3000 hours bin includes anyone with lifetime hours of at least 3000 hours. The table in panel (c) reports the corresponding information in levels for the 1750 hours bin and for reference also the respective values for the bins at each end of the lifetime hours distribution.

55 percent more hours per week worked relative to men in the 1750 hour bin, compared with 7 percent more weeks per year worked and 5 percent more years.

3.4 Lifetime Hours and Lifetime Earnings

A key objective of our analysis is to examine the role of dispersion in lifetime hours for lifetime earnings inequality and, in particular, the extent to which differences in lifetime hours affect lifetime earnings via their effect on human capital accumulation. In this subsection we present information on two moments that relate lifetime hours and earnings: the (reduced-form) elasticity of lifetime earnings with respect to hours, and the (reduced-form) elasticity of earnings growth

Figure 5: Lifetime Earnings and Lifetime Hours

(a) Lifetime Earnings by Lifetime Hours

(b) Lifetime Earnings by Lifetime Hours Decile

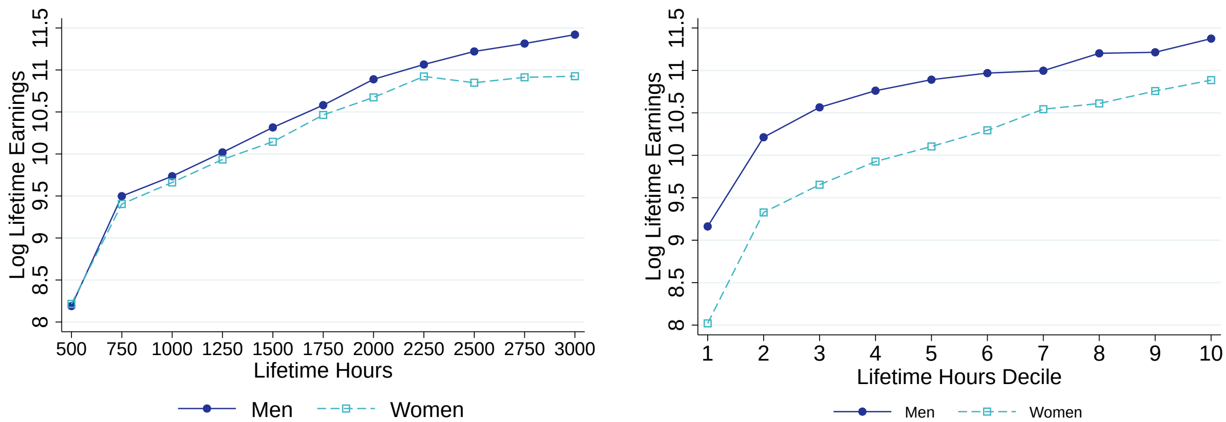

Notes: The 500 hours bin includes anyone with annualized lifetime hours below 750 and the 3000 hours includes anyone with annualized lifetime hours of at least 3000.

over the life-cycle with respect to lifetime hours. Higher human capital accumulation would lead one to expect both higher growth in earnings over the life-cycle and higher lifetime earnings. We emphasize that these reduced-form moments on their own do not provide causal information about the effect of lifetime hours on human capital accumulation. But it is intuitive that human capital responses would influence these moments and in our quantitative work we will compute these moments in our model and the data.

Figure 5a shows the relationship between lifetime hours and lifetime earnings for men and women. Lifetime earnings are strongly increasing in lifetime hours. For example, relative to men in the 1750 lifetime hours bin, men in the 750 hours bin have lifetime earnings that are 108 log points lower, and men in the 2750 hours bin have lifetime earnings that are 73 log points higher. Using the 750 and 2750 hours bins, the implied elasticity of log lifetime earnings with respect to log lifetime hours is 1.5. This elasticity is also above one across all adjacent lifetime hours bins except at the very bottom and top of the hours distribution, where it is approximately one.

The curve for women tracks the curve for men relatively closely up to the 2250 hours bin, but from 2500 hours onward, the curve for women becomes flat and the gap between men and women grows substantially. While this divergence above 2500 hours is potentially interesting, we note that only 4 percent of women lie in this region, compared with 22 percent of men (see Figure 2). As an alternative representation of the data, the right panel of Figure 5 plots log lifetime earnings versus gender-specific deciles of the lifetime hours distribution. This figure shows a roughly constant slope for women beyond the second decile. In particular, the flat portion of the curve for women in the left panel no longer appears.

Next, we examine the relationship between life-cycle earnings growth and lifetime hours. Fig-

Figure 6: Earnings Early and Late in the Life-Cycle and Lifetime Hours Deciles

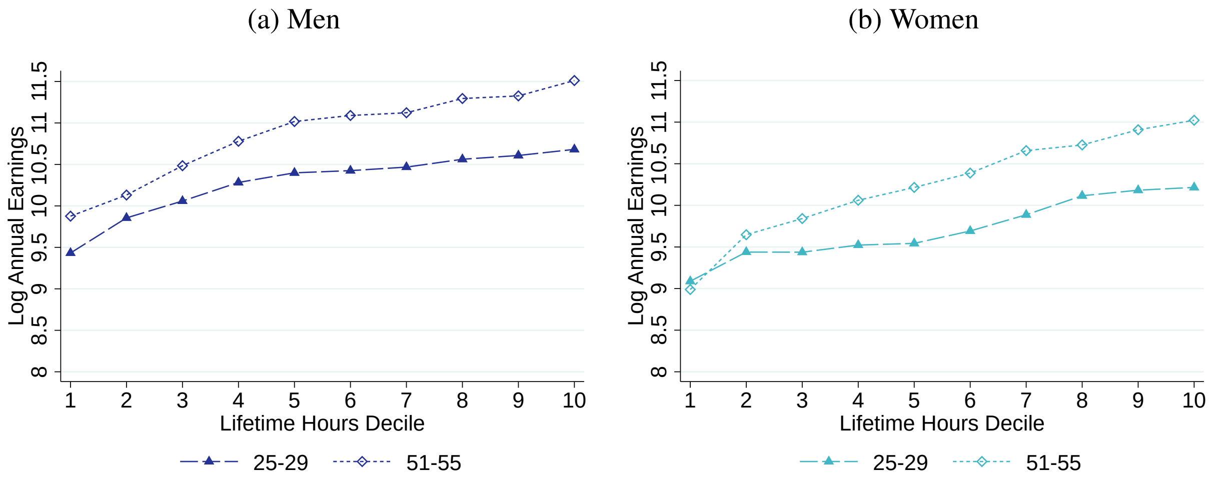

Notes: We average over the cross-sectional earnings between ages 25-29 and 51-55 for all person-year observations with positive hours and then calculate the growth rate.

ure 6 studies how log earnings early in life (ages 25-29) and later in life (ages 51-55) vary with lifetime hours. The gap between these two curves measures how the slope of the age-earnings profile varies with lifetime hours. Two properties emerge. First, earnings at each age are increasing in lifetime hours. In particular, individuals with high lifetime hours have higher average earnings than individuals with low lifetime hours at all ages. Second, the gap between the two curves is also increasing in lifetime hours. These properties hold for both men and women. The second property in particular is consistent with higher lifetime hours being associated with greater human capital accumulation.

The unconditional correlations presented in the previous two figures serve as suggestive evidence for a link between lifetime hours and human capital accumulation. Consistent with much of the literature on life-cycle earnings growth, the model that we develop later in this paper will allow for heterogeneity in learning ability. It is of interest to know if the relationships between lifetime hours and earnings are robust to controlling for learning ability.

To do this we leverage the fact that the NLSY contains information on AFQT scores, which might plausibly be viewed as a proxy for learning ability. We supplement the bin scatter plots in Figures 5 and 6 with individual level regression analyses in which we examine the effect of adding AFQT scores as a control. Results are presented in Table 1.

Two main messages emerge. First, the regressions without AFQT scores confirm the same relationships found in the bin scatter plots: log lifetime earnings and life-cycle earnings growth are both increasing in lifetime hours. In particular, the elasticity of log lifetime earnings with respect to log lifetime hours is greater than one for both men and women. Second, although AFQT scores are statistically significant in three of the four cases, the effects of lifetime hours remain statistically

Table 1: Lifetime Hours and Earnings Regressions

Men

(a) log of Lifetime Earnings

	(1)	(2)
log Lifetime Hours	1.44***	1.30***
	(0.02)	(0.02)
AFQT Percentile		0.01***
		(0.00)
Constant	10.74***	10.74***
	(0.01)	(0.01)
N	3008	2881
R^2	0.59	0.68

(b) Life-Cycle Earnings Growth

	(1)	(2)
log Lifetime Hours	1.44***	1.20***
	(0.20)	(0.20)
AFQT Percentile		0.02***
		(0.00)
Constant	1.17***	1.20***
	(0.05)	(0.05)
N	2291	2200
R^2	0.02	0.08

Women

(c) log of Lifetime Earnings

	(1)	(2)
log Lifetime Hours	1.30***	1.28***
	(0.01)	(0.01)
AFQT Percentile		0.01***
		(0.00)
Constant	10.01***	10.02***
	(0.01)	(0.01)
N	3253	3163
R^2	0.75	0.81

(d) Life-Cycle Earnings Growth

	(1)	(2)
log Lifetime Hours	0.94***	0.96***
	(0.16)	(0.16)
AFQT Percentile		-0.00
		(0.00)
Constant	1.35***	1.35***
	(0.05)	(0.05)
N	1999	1956
R^2	0.02	0.02

Notes: In all regressions, log annualized lifetime hours and AFQT percentiles are demeaned such that the constants are comparable across specifications (1) and (2). Standard errors are in parentheses. *,**, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively. In panels (b) and (d), we first compute, for each individual, average earnings at ages 25–29 and 51–55, considering only years with at least 520 hours worked. We then construct the growth rate between these two averages. The sample is restricted to individuals with at least three years of at least 520 hours worked in both the 25–29 and 51–55 age ranges.

and economically significant, with the elasticity of lifetime earnings with respect to lifetime hours still larger than one.

4 Lifetime Hours and Human Capital Accumulation: A First Look

It is well known in the inequality literature that the cross-sectional dispersion of log earnings (and wages) increases substantially over the life-cycle (Deaton and Paxson, 1994; Storesletten, Telmer, and Yaron, 2004; Huggett et al., 2006). The literature has focused on two proximate causes: the accumulation of persistent shocks and heterogeneity in learning ability. In this section we argue that heterogeneity in hours worked, which the literature has largely neglected, constitutes a third relevant factor. We make this argument in the context of a standard Ben-Porath model. The key economic channel captured by this model is that individuals who expect to work longer hours throughout their lifetime have greater incentives to invest in human capital. This will in turn have effects on both lifetime earnings and life-cycle earnings growth.

In the first subsection we lay out the simple benchmark Ben-Porath model that is at the heart of much recent work on life-cycle models of inequality. In the second subsection we highlight the mechanism linking permanent differences in hours of work and incentives to accumulate human capital. A similar result appears in Neal and Rosen (2000). To facilitate transparency, here we assume that total hours in each period of working life are exogenously fixed at some level and ask how changes in this level affect incentives to invest in human capital. Our quantitative model in Section 5 will endogenize total hours.

4.1 A Simple Ben-Porath Model

We consider an individual who is born in period 1, works until period T_R, and dies at the end of period T. They have h units of time each period prior to retirement and preferences over streams of consumption (c_t) given by:

$$\sum_{t=1}^{T} \beta^t u(c_t)$$

where $0 < \beta < 1$ is a discount rate and $u(c_t)$ is strictly increasing, strictly concave, and twice continuously differentiable.

The individual is endowed with initial human capital x_1. At each age t prior to retirement the individual chooses to divide their h units of time between production (n_t) and investing in human capital (s_t). Because the individual does not value leisure, it follows that $n_t + s_t = h$ during this period of the life-cycle.

Time devoted to human capital investment at age t produces new human capital in the amount of $\alpha \cdot (x_t s_t)^{\phi}$, where $\alpha > 0$ reflects productivity of the human capital production function, $x_t s_t$ reflects efficiency units of time devoted to human capital investment, and $0 < \phi < 1$ governs the extent of diminishing returns in the human capital production function. In what follows we will refer to α as learning ability. The law of motion for an individual's human capital follows:

$$x_{t+1} = (1-\delta)x_t + \alpha(x_t s_t)^{\phi},$$

where $0 < \delta < 1$ reflects depreciation of human capital.

A competitive labor market offers a time-invariant wage rate of w per efficiency unit of production labor services, so the individual receives labor earnings equal to $n_t x_t w$ at age t. While the individual faces the same wage per efficiency unit of production labor services at all ages, the payment per unit of time spent working is influenced by both the individual's level of human capital and how they divide their time between producing and investing. The individual can borrow and lend at the time-invariant interest rate r. The only constraint on borrowing is that the individual cannot die with negative assets.

4.2 The Mechanics of Human Capital Investment

Heterogeneity in human capital accumulation is a potential source of heterogeneity in earnings growth. Because time devoted to investment directly affects accumulation of human capital, it is of interest to examine the forces that shape time devoted to human capital investment.

In Appendix A we show that the optimal allocation of time between producing and investing at each age t, assuming an interior solution, satisfies the following equation:

$$wx_t = \alpha\phi x_t^{\phi} s_t^{\phi-1} \sum_{t'=t+1}^{T_R-1} \left[\frac{1}{1+r}\right]^{t'-t} wh(1-\delta)^{t'-(t+1)} \tag{1}$$

This condition can be interpreted as requiring that the marginal value of time allocated to production should equal the marginal value of time allocated to investment. The left-hand side is the effective wage for an individual at age t and thus reflects the value of a marginal increase in production time at age t, holding all else constant. Turning to the right-hand side, higher investment today produces a stream of benefits in all future periods until retirement, and the overall benefit is the present value of this stream. To understand the terms in this sum, note that the term $\alpha\phi x_t^{\phi} s_t^{\phi-1}$ reflects the marginal increase in human capital at age $t+1$ as a result of a marginal increase in time devoted to investment. This investment will also increase human capital in period $t' > t+1$ by the amount $\alpha\phi x_t^{\phi} s_t^{\phi-1}(1-\delta)^{t'-(t+1)}$. The value of this additional human capital at age t' is the product of the total number of hours worked at age t' and the wage per efficiency unit of labor services.[12] An important point is that higher (future) hours worked increase the marginal benefit of additional investment today.

We use this equation to highlight two forces that shape time devoted to investment. The first highlights the key mechanism in the analysis of Huggett et al. (2006). In their framework, heterogeneity in learning ability α across individuals leads to heterogeneity in the growth rate of earnings and thus an increase in the cross-sectional distribution of earnings over the life-cycle. Heterogeneity in α will generate heterogeneity in human capital accumulation and earnings even if there is no impact on time allocation. But importantly, equation (1) shows that an increase in α holding x_t constant raises the right-hand side, thereby requiring an increase in s_t. The effect is intuitive: A higher value of α increases the return to time spent investing in human capital relative to time spent producing output. This reinforces the direct effect of heterogeneity in α.

Next we use Equation (1) to show that higher total hours also creates an incentive to devote additional time to investment in human capital. In particular, consider two individuals with the same human capital x_t and the same learning ability, but assume that future total hours of work h are exogenously higher for one individual. The individual with higher future hours will have a higher value of the right-hand side of Equation (1) and thus will require a higher value of s_t in order to maintain equality of the left- and right-hand sides. This result is also intuitive: Higher future hours increase the marginal value of additional investment today. Differences in human capital investment will lead to differences in human capital stocks at older ages and are thus a source of inequality in life-cycle earnings growth.

5 A Structural Model of Lifetime Hours and Earnings

In this section we develop a generalization of the heterogeneous agent Ben-Porath models in Huggett et al. (2011) and Guvenen et al. (2014) in order to assess the extent to which variation in lifetime hours of work across individuals affects inequality in lifetime earnings. Our model features the three sources of heterogeneity in these papers: heterogeneity in initial human capital endowments, heterogeneity in learning ability, and idiosyncratic shocks to income. Differently from both of these papers, our model also allows for preference heterogeneity in order to account

[12] Note that it is total future hours that enter on the right-hand side, and not future production hours. This is because higher human capital increases productivity for both activities and the marginal value of time is equated across the two activities.

for the salient features of heterogeneity in hours of work.

Our framework shares two features with the frameworks in Huggett et al. (2011) and Guvenen et al. (2014). First, we abstract from fertility. For this reason our model is not well suited to analyzing life-cycle labor supply decisions of women; and therefore, like Huggett et al. (2011) and Guvenen et al. (2014), we will connect our model to data on men. Second, individuals in our model will have positive hours in all time periods. For this reason, we will further restrict our sample to men for whom the years worked margin is not an important margin. In the conclusion, we discuss extending our analysis to a larger sample.

5.1 Households

We study the choices of a single cohort of individuals in partial equilibrium. Each individual i in the cohort lives from age $t = 1$ to $t = T$, retires exogenously at age $t = T_R$, and has $\bar{h}$ units of time each period.

A key element of our analysis is to extend the basic Ben-Porath model from the previous section to account for the differences in lifetime hours across individuals. As documented in Bick, Blandin, and Rogerson (2022), observables explain very little of the variation in hours worked across individuals. Like them, we introduce preference heterogeneity as a parsimonious way to generate the magnitude and nature of hours variation found in the data. As noted in the introduction, we use preference heterogeneity to capture a variety of factors that influence desired labor supply holding wages and non-labor income constant, including such things as the productivity of non-market time, family demands on time, health, and social norms in addition to true differences in preferences. Our specification of preference heterogeneity is in turn motivated by the properties of the autocorrelation of hours documented in Figure 3. The fact that the autocorrelation declines over time suggests a transitory mean reverting component to preference heterogeneity, while the fact that the autocorrelation plateaus above zero at longer lags suggests a permanent component. Previous research has allowed for either permanent preference heterogeneity (e.g., Kaplan (2012), Heathcote et al. (2014), Keane and Wasi (2016) and Bick et al. (2022)) or transitory preference heterogeneity (e.g., Imai and Keane (2004) and Chang, Kim, Kwon, and Rogerson (2020)), but not both. We show below that allowing for both is quantitatively important.

We thus assume that individual i has preferences of the form:

$$\sum_{t=1}^{T} \beta^{t-1} \left[\frac{c_{i,t}^{1-1/\sigma}}{1-1/\sigma} - \psi_i \pi_{i,t} \frac{h_{i,t}^{1+1/\gamma}}{1+1/\gamma} \right]$$

where $c_{i,t}$ is consumption at age t and $h_{i,t}$ is total time devoted to work at age t. The parameters β, σ, and γ are common to all individuals and satisfy $0 < \beta < 1$, $\sigma > 0$, and $\gamma > 0$. The terms ψ_i and $\pi_{i,t}$ are individual-specific preference shifters. ψ_i is time-invariant and assumed to be log-normally distributed with mean μ_ψ and standard deviation σ_ψ. In contrast, $\pi_{i,t}$ is time-varying and follows an AR(1) process:

$$\log \pi_{i,t+1} = \rho_\pi \log \pi_{i,t} + \nu_{i,t+1} \tag{2}$$

$$\nu_{i,t+1} \sim N(0, \sigma_\pi) \tag{3}$$

where the innovations $\nu_{i,t+1}$ are assumed to be iid over time and across individuals.[13]

As in Section 4, total hours of work at age t for individual i, $h_{i,t}$, are allocated between producing ($n_{i,t}$) and investing in human capital ($s_{i,t}$). Following Huggett et al. (2011), the human capital accumulation process is now specified as:

$$x_{i,t+1} = z_{i,t+1}\left[(1-\delta)x_{i,t} + \alpha_i(s_{i,t}x_{i,t})^\phi\right]$$

where α_i is an individual specific learning ability and $z_{i,t+1}$ is a log-normally distributed shock to human capital:

$$\log z_{i,t+1} \sim N(0, \sigma_z)$$

This shock is iid over time and across individuals. Note that while the shock $z_{i,t+1}$ is purely transitory, its effect is persistent because it affects the individual's stock of human capital. That is, these shocks will accumulate over the life-cycle and cause human capital levels to spread out. We take a broad interpretation of these shocks and, in particular, view them as potentially capturing some part of the stochastic movements up and down the job ladder that are explicitly modeled in studies such as Bagger et al. (2014) and Ozkan et al. (2023), in addition to any shocks that affect the value of an individual's human capital.

Each individual is characterized by two fixed effects (ψ_i and α_i) and two initial conditions ($\pi_{i,1}$ and $x_{i,1}$). In our quantitative analysis we assume that $x_{i,1}$ and ψ_i are joint log-normally distributed:

$$\log(x_{i,1}, \psi_i) \sim N(\mu_x, \mu_\psi, \sigma_x, \sigma_\psi, \rho_{x,\psi})$$

We assume that α_i is log-normally distributed, and to economize on the size of our state space we assume that $x_{i,1}$ is perfectly correlated with α_i.[14] While our analysis does not explicitly model

[13] Some of the transitory movements in hours are accounted for by variation in weeks worked, which likely reflect unemployment spells. Our specification implicitly uses preference shocks to account for these unemployment spells. Alternatively, we could add explicit unemployment shocks as an additional source of transitory hours variation. While it is feasible to do this, we have opted not to. This is because transitory changes in hours turn out to not be quantitatively important for lifetime earnings inequality in our sample, thereby reducing the payoff to exploring multiple distinct shocks as sources of transitory fluctuations.

[14] Guvenen et al. (2014) also assumes that learning ability and initial human capital are perfectly correlated. Huggett

occupations and the potential for human capital accumulation opportunities to vary across occupations, we view the heterogeneity in learning ability α_i as reflecting both the innate ability of an individual to learn as well as the learning abilities inherent in the occupational choice that best suits the individual's innate skills. We impose that the initial distribution for $\pi_{i,1}$ is the ergodic distribution for the π_t process. Because the innovations for this process are uncorrelated with all other variables, we assume that $\pi_{i,1}$ is uncorrelated with all other variables.

5.2 Government

The government levies a proportional tax τ_c on consumption and a progressive tax on labor income. Following Benabou (2002) and Heathcote et al. (2014) we assume an individual with pre-tax labor income y receives post-tax labor income of

$$\tau_0 y^{1-\tau_1}$$

The parameter $\tau_0 \in [0,1)$ determines the overall level of the income tax, while the parameter $\tau_1 \in [0,1)$ determines the progressivity of the tax. To minimize the state space, we model a simplified Social Security system. In particular, we assume that all individuals in our model receive a transfer from the government equal to I_{ss} in each period beginning in period T_R. Because our quantitative analysis will focus on the highly attached sample of men described in Section 6.1 that is eligible for very few transfers, we abstract from all transfers other than Social Security. Because our model is partial equilibrium and our quantitative analysis will focus on a subset of the overall population, we do not require that the government budget balances.

5.3 Markets and Prices

As noted earlier, we study the choices of our cohort in partial equilibrium. We assume a stationary economic environment in which the interest rate r is constant and the wage per efficiency unit of labor grows exogenously at constant rate g_w. These properties of prices are consistent with a balanced growth path equilibrium in economies with constant-returns-to-scale production functions and constant productivity growth that is labor augmenting.

We let w_t denote the wage per efficiency unit that the cohort faces at age t. An individual of age t with human capital x that supplies n units of time to production earns pre-tax labor income nxw_t. There are no state contingent securities, but individuals are able to borrow and save at the

et al. (2006) argue that a very high correlation between human capital and learning ability at our starting age of 25 is consistent with a much lower correlation between these two values at an earlier age. They use a value of 0.746 for the correlation between initial human capital and learning ability at age 23.

interest rate r subject to the natural borrowing constraint. An age t individual with human capital x_t and assets k_t who is not retired has the following period budget equation:

$$(1+\tau_c)c_t + k_{t+1} = \tau_0 \cdot (w_t x_t n_t)^{1-\tau_1} + (1+r)k_t$$

A retired individual with assets k_t faces a period budget equation:

$$(1+\tau_c)c_t + k_{t+1} = I_{ss} + (1+r)k_t$$

5.4 The Individual's Problem

At the start of a period, an individual has a six-dimensional state (t,k,x,α,ψ,π) consisting of age (t), assets (k), human capital (x), learning ability (α), permanent work disutility (ψ), and transitory work disutility (π). Taking as given the individual state, the wage rate w_t, the interest rate r, and government policies, a working age individual, $t < T_R$, solves the following recursive problem:

$$V(t,k,x,\alpha,\psi,\pi) = \max_{c,k',s,n} \quad \frac{c^{1-1/\sigma}}{1-1/\sigma} - \psi\pi\frac{(n+s)^{1+1/\gamma}}{1+1/\gamma} + \beta\mathbb{E}_{z',\pi'}V(t+1,k',z'\hat{x},a,\psi,\pi') \tag{4}$$

$$s.t. \quad (1+\tau_c)c + k' = (1+r)k + \tau_0 \cdot (nxw_t)^{1-\tau_1} \tag{5}$$

$$\hat{x} = [(1-\delta)x + \alpha(sx)^{\phi}] \tag{6}$$

$$c,s,n \geq 0 \quad \text{and} \quad n+s \leq \bar{h} \tag{7}$$

A retired individual, $t \geq T_R$, solves an identical problem, except that they receive a Social Security transfer of I_{ss} and face the added constraint that $n = s = 0$. An individual in their last period of life, $t = T$, faces an additional nonnegative savings constraint: $k' \geq 0$.

6 Calibration

This section describes our model calibration in three steps. First, we describe the sample we use to connect our model with the data. Second, we describe our procedure for choosing parameters to match salient features of earnings and hours over the life-cycle. Third, we report on the ability of the calibrated model to match both targeted and untargeted moments.

6.1 Sample Selection

The balanced panel that we created in Section 2 included men and women and did not impose any explicit criterion regarding labor force attachment beyond excluding the very few individuals with annualized lifetime hours less than one. As noted in the previous section, our model does not include fertility or adjustment along the years worked margin. This motivates us to focus on a sample of highly attached men for whom the margin of years worked does not play an important role.

To create our baseline sample of highly attached men we restrict attention to men who have (annualized) lifetime hours of at least 1750. Referring back to panel (b) of Figure 4, this threshold serves to eliminate those individuals for whom differences in years worked are an important contributor to variation in lifetime hours. The sample of men with lifetime hours of at least 1750 includes 80 percent of the men in our balanced panel. While the high attachment sample displays less dispersion in lifetime hours than the overall sample, dispersion in lifetime hours remains large. For example, the interquartile range remains large, falling from 525 to 442, and still represents more than 20 percent of median lifetime hours. A key message is that there is substantial variation in lifetime hours even after removing individuals with low attachment.

We have also carried out our analysis using a different sample selection criterion and found very similar results. Rather than applying a threshold based on lifetime hours, this alternative criterion measured attachment as the number of years in which an individual worked at least 520 hours. The value of 520 was chosen because it was used by both Huggett et al. (2011) and Guvenen et al. (2014) to restrict their samples and also appears in the criterion used by Guvenen et al. (2022).

We explored two different thresholds: (i) at least 520 hours in all 31 years and (ii) at least 520 hours in at least 20 of the 31 years. Among all men in our balanced panel, 47 percent meet threshold (i) and 87 percent meet threshold (ii), compared with 80 percent who work at least 1750 lifetime hours. There is substantial overlap among all three samples, and so it is not surprising that we find similar results when applying the alternative sample selection criteria.[15]

6.2 Calibration Procedure

The model described in Section 5 features 22 parameters: three price parameters (r, w_1, and g_w), three common preference parameters (β, σ, and γ), three technology parameters for the human capital accumulation process (δ, ϕ, and σ_z), two parameters characterizing the distribution of

[15] 99.1 percent of men who meet threshold (i) and 87.9 percent of men who meet threshold (ii) have lifetime hours of at least 1750. Conversely, 68.5 percent of men who work at least 1750 lifetime hours meet threshold (i), and 99.9 percent of them meet threshold (ii).

Table 2: Externally Calibrated Parameter Values

Parameter	Interpretation	Value	Source
β	Patience	0.980	Huggett et al. (2011)
r	Interest rate	0.020	$1/\beta$
σ	CRRA	1.000	–
γ	Frisch elasticity	0.300	–
δ	Human capital depreciation	0.020	Huggett et al. (2011)
τ_0	Tax Rate	0.810	Heathcote et al. (2014)
τ_1	Tax Progressivity	0.181	Heathcote et al. (2014)
τ_c	Consumption Tax	0.070	McDaniel (2007)
μ_x	Mean of $\log x_0$	0.00	Normalization
g_w	Exogenous wage growth rate	0.005	CPS

learning ability (μ_α and σ_α), two parameters characterizing transitory preference heterogeneity (ρ_π and σ_π), four parameters characterizing the tax and transfer system (τ_c, τ_0, τ_1, and I_{ss}), and five parameters characterizing the joint distribution of x_0 and ψ (μ_x, μ_ψ, σ_x, σ_ψ, and $\rho_{x,\psi}$). When connecting our model with the data we will allow for classical measurement error in both log hours and log wages, with standard deviations of σ_{mh} and σ_{me}, respectively. This increases the total number of parameters to 24.

Ten parameters are set externally and summarized in Table 2. We normalize the parameter $\mu_x = 0$ since it is not identified separately from w_1. Setting the period length equal to a year, we set $r = 0.02$ and $\beta = 1/(1+r) = 0.9804$. We adopt commonly used values for the two curvature parameters σ and γ: $\sigma = 1$, and $\gamma = 0.3$. Tax function parameters are set according to the estimates in Heathcote, Storesletten, and Violante (2017): $\tau_0 = 0.81$ and $\tau_1 = 0.181$, and the consumption tax τ_c is set to 0.07. Following Huggett et al. (2011), we set $\delta = 0.02$. We set $g_w = 0.005$ by comparing average real wages of men aged 25-29 in 1981 with average real wages of men aged 25-29 in 2018 using CPS data. Details of this calculation are provided in Appendix B.3.

For the remaining fourteen parameters, we use a simulated method of moments procedure, selecting values so that the solved model matches fourteen target moments.[16] We first list these moments and then provide some intuition about how they relate to the model parameters.

Because the distributions of hours and earnings are the core outcomes of interest in our analysis, thirteen of the targeted moments describe these distributions. Specifically, we target five moments related to hours, six moments related to earnings, and two moments related to joint properties of earnings and hours. The other targeted moment is the ratio of the Social Security transfer to mean

[16] When solving the model, we approximate distributions using gridpoints: 9 for ψ_i, 7 for α_i, 21 for $x_{i,0}$, and 5 each for the shocks $z_{i,t}$ and $\pi_{i,t}$.

annual earnings.

The five hours moments that we target are the overall mean and standard deviation of log hours for the pooled sample of all individuals between ages 25 and 55, and values for the autocorrelation of individual hours at lags of 1, 10, and 20 years. The six earnings moments that we target are mean log earnings at ages 30 and 50, the standard deviation of log earnings at 30 and 50, and the autocorrelation of log earnings at lags of 1 and 20 years. The two joint hours and earnings moments that we target are the correlation between hours and earnings at age 30, and the slope of the relationship between lifetime hours and life-cycle earnings growth, as depicted in Figure 5a.

Because several of our moments involve taking the log of hours, we drop any annual observations with a zero value. With a relatively small sample size, standard deviations of logs are quite sensitive to the presence of a few very low values. For this reason, we follow Huggett et al. (2011) and Guvenen et al. (2014) and also exclude annual observations with hours less than 520 when computing our empirical moments. We note that relatively few observations are dropped because of these choices: 1.6 percent of the annual person-year observations in our sample of men with lifetime hours of at least 1750 have zero annual hours, and 0.9 percent have positive annual hours below the 520 hours threshold. We apply the same criterion when computing moments in the model simulated data.

In what follows we provide some intuition about the connection between the targeted moments and fourteen parameters that we calibrate. This discussion should be understood as purely heuristic since all fourteen parameters influence all fourteen moments. Nonetheless, we think it provides some useful insight into the mechanics of the calibration procedure.

Given values for all of the other parameters, the five moments of the hours distribution can be used to pin down values for μ_ψ, σ_ψ, ρ_π, σ_π, and σ_{mh}. Intuitively, the value of μ_ψ is tightly linked to the cross-sectional mean of annual hours. Each of the other four parameters influences both the variance of log annual hours as well as the shape of the hours autocorrelation profile, so matching the variance of log annual hours and the value of the autocorrelation at three different lag lengths will determine their values.

Holding other parameters fixed, w_1, σ_x, μ_α, σ_α, σ_z and σ_{me} will impact properties of the earnings distribution. The values of w_1 and σ_x will impact the level and variance of earnings for young workers, and μ_α will influence mean life-cycle earnings growth. The values of σ_α, σ_z, and σ_{me} will influence the extent to which earnings become more dispersed with age as well as the autocorrelation of earnings. Because initial human capital impacts earnings of young individuals and the value of ψ impacts hours, the correlation parameter $\rho_{x,\psi}$ will influence the correlation of earnings and hours for young workers. We note that because learning ability and initial human capital are perfectly correlated, the value of $\rho_{x,\psi}$ also controls the correlation between learning ability and tastes for work. The values σ_z and σ_{me} have different effects on the shape of the autocorrelation function for earnings, so targeting values of the autocorrelation function for log earnings will help to determine the values of these parameters.

The final parameter is ϕ, which determines the elasticity of human capital with respect to investment. All else equal, a higher value of ϕ will increase the slope of the relationship between life-cycle earnings growth and lifetime hours worked. We emphasize that the targeted slope is not an estimate of the causal effect of lifetime hours on life-cycle wage growth. In particular, if learning ability varies across the lifetime hours distribution, then the observed relationship between lifetime hours and life-cycle earnings growth will also include the effect of learning ability. We also note that our procedure targets a low-frequency relationship to help pin down the value of ϕ. Frictional models of wage setting such as Cahuc, Postel-Vinay, and Robin (2006) imply that wages may respond to productivity with a lag, in which case focusing on short-term variation in hours and wages may be misleading. Non-linearities in the hours-wage profile that workers face may also create issues when using short-term variation in hours to estimate the effect of hours on human capital accumulation, see, e.g., Bick et al. (2022).

Before presenting the calibrated parameter values, it is important to be explicit about how we connect hours in the data with hours in the model. The potential issue is the extent to which reported hours worked in the data include time devoted to investment in human capital. The standard convention in the literature is to include time spent in investment as part of reported work hours for employed individuals.[17] We adopt this convention when connecting our model to the data. To the extent that some of the time devoted to human capital investment is not included in reported working hours in the data, total work hours are underestimated. We have experimented with other specifications, allowing for some fraction of investment time to not be counted as reported work hours in the data. Modest departures from our benchmark were found to have only minor effects on our quantitative results both in this section and in later sections.

Table 3 displays the calibrated parameter values. We draw attention to four properties of these values. First, although we use a novel moment to help identify the parameter ϕ, the calibrated value of 0.6 is within the broad range of estimates found in the literature, though toward the lower end.[18] Second, despite using different information in our calibration than Huggett et al. (2011), our procedure yields a very similar value for σ_z.

Third, we find a modest negative correlation between permanent tastes for work and initial human capital, with $\rho_{x,\psi}$ equal to -0.30. Because we impose that learning ability α and initial human capital are perfectly correlated, our calibration thus also implies a negative correlation between permanent disutility for work and learning ability. We previously argued that high learning

[17] Guvenen et al. (2014) added a constraint that limited the amount of time that could be devoted to investment when production time is positive.

[18] See, for example the handbook chapter by Browning, Hansen, and Heckman (1999) as well as the discussion in Heckman, Lochner, and Taber (1998) and Huggett et al. (2006).

Table 3: Internally Calibrated Parameter Values

Parameter	Interpretation	Value	Moment
w_1	Wage in first period	24.00	Mean log earnings, age 30
σ_x	SD of $\log x_1$	0.35	SD log earnings, age 30
μ_α	Mean of $\log \alpha$	-2.30	Mean log earnings, age 55
σ_α	SD of $\log \alpha$	0.18	SD log earnings, age 55
μ_ψ	Mean of $\log \psi$	4.22	Mean log annual hours, age 25-55
σ_ψ	SD of $\log \psi$	0.60	SD log annual hours, age 25-55
$\rho_{x,\psi}$	Corr. of $(\log x_1, \log \psi)$	-0.30	Correlation of hours and earnings, age 30
σ_π	SD of $\log \pi$	0.49	Hours autocorrelation profile
ρ_π	Autocorrelation of $\log \pi$	0.87	Hours autocorrelation profile
σ_{mh}	SD measurement error	0.10	Hours autocorrelation profile
σ_{me}	SD measurement error	0.17	Earnings autocorrelation profile
σ_z	SD human capital shock	0.11	Earnings autocorrelation profile
ϕ	HC elasticity wrt investment	0.60	Lifetime hours, earnings growth
I_{ss}	Social Security benefit	0.19	$0.4\bar{e}$

ability and high hours are both sources of higher life-cycle earnings growth. Given that the correlation between ψ and α is only -0.30, it follows that these two channels operate somewhat independently of each other.

Fourth, the transitory component of preference heterogeneity is very persistent and contributes substantially to the cross-sectional variation in tastes for work. Specifically, the variance of $\log \pi$ in the ergodic distribution of the transitory process is equal to 0.99, which is almost three times larger than the variance in the log of the permanent component ψ, which is equal to 0.36. Our calibration procedure also implies a substantial amount of measurement error in hours and earnings, with $\sigma_{mh} = 0.10$ and $\sigma_{me} = 0.17$, respectively. These values imply that measurement error in hours and earnings accounts for 13 and 5 percent of the cross-sectional variance in annual hours and annual earnings, respectively.

6.3 Fit of the Calibrated Model

In this subsection we report on the ability of the model to match both targeted and untargeted moments. In what follows we will show several figures that report life-cycle profiles for outcomes related to hours and earnings. In the literature it is standard to control for time or cohort effects when documenting life-cycle profiles. Because our sample consists of individuals born in a relatively narrow range of years, we view our sample as effectively representing a single cohort and so do not control for cohort effects. Rather than extracting time effects from the data, we deal with

Figure 7: Model Fit for Earnings

time effects by introducing time-varying wages in the model (captured by g_w).

We begin by examining the model's ability to fit moments related to earnings. Results are shown in the three panels of Figure 7. Overall, the model closely tracks the evolution of the mean and standard deviation of log earnings over the life-cycle, as well as the persistence of earnings.[19] As a reminder, for the mean and standard deviation, we targeted the moments at ages 30 and 50; and for the autocorrelation, we targeted lags 1 and 20. One small discrepancy to note is that the model implies a slightly concave profile for the standard deviation, whereas the profile in the data is slightly convex. This same pattern is present in the calibrated model of Huggett et al. (2011); see their Figure 2(b).[20] We also note that the mean of log earnings in the model for ages 25-30 is higher than in the data, which – as we will discuss momentarily – stems from a similar discrepancy between hours in the model and the data.

Next we consider the moments for hours. Results are shown in the three panels of Figure 8. Panel (a) shows the age profile for log mean hours, and Panel (b) shows the age profile for the standard deviation of log hours. Recall that our calibration procedure targeted the overall cross-sectional mean and variance of log hours, but not the values for any particular age. In both the model and the data these profiles are relatively flat between the ages of 30 and 50, so matching the overall sample value leads to a reasonable fit to the life-cycle profiles.

As was the case for earnings, we see some discrepancy for hours between the model and data over the 25-30 age range. In particular, while both profiles are also relatively flat in the model over

[19] Appendix Figure B.3 shows that as for hours the autocorrelation profile of earnings excluding imputed values for annual earnings is basically the same as the one shown in Figure 7c.

[20] In principle, a model like ours could generate a flat profile for the variance in earnings in the early part of the life-cycle. This would happen if individuals with low initial human capital have high growth of human capital, or if individuals with high initial human capital spend so much time investing that they have lower initial earnings. But as the figure shows, these effects are not sufficiently powerful in our calibrated model.

Figure 8: Model Fit for Hours

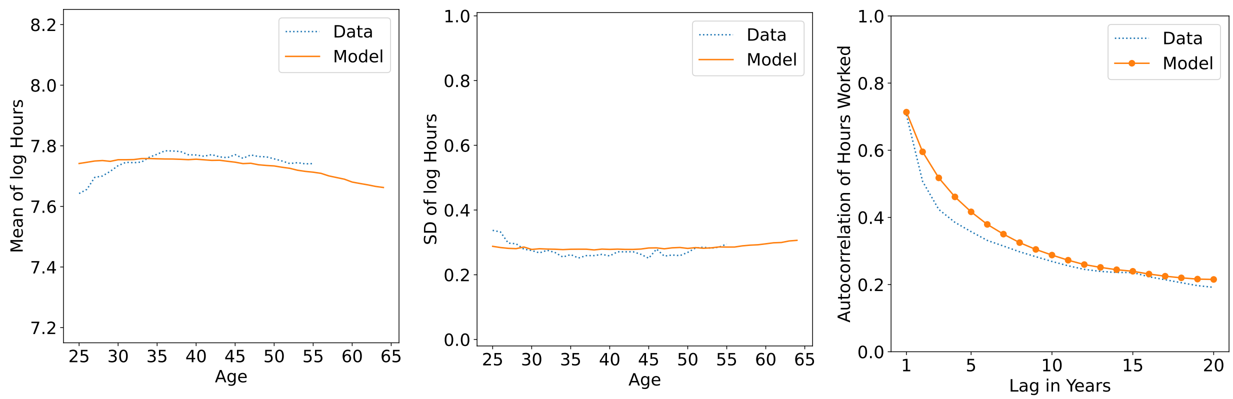

the 25-30 age range, mean log hours are increasing and the variance of log hours is decreasing in the data over this age range. These two properties are intimately related: the increased prevalence of individuals with relatively low annual hours of work in this age range tends to both decrease the mean and increase the variance. We offer two potential rationalizations for this discrepancy. The first is that individuals in this age range are more likely to have spells of unemployment as they search for a good match, a feature that our model abstracts from. (See Kaplan (2012) for an analysis incorporating this feature.) The second is our assumption that the data on hours reflect both production and investment time. This age range is the period of highest investment in human capital, so if some investment time is not included in reported work hours in the data, our model would be expected to overestimate work hours at young ages.[21]

Panel (c) shows the autocorrelation function for hours. We targeted values at lags of one, ten, and twenty, and Panel (c) shows that the model does a good job of capturing the entire profile, though the model profile is somewhat less steep than the data at low values of the lag.

Next we consider the two moments that involve relationships between earnings and hours. Figure 9a shows the age profile for the correlation between earnings and hours in the model and in the data. The only targeted value was the correlation at age 30. Both in the data and in the model the age profile for this correlation displays a very modest downward drift, though in the data there is an additional dip between ages 35 and 45.

Figure 9b shows the relationship between lifetime hours and life-cycle earnings growth in the model and in the data. The curve from our calibrated model tracks the curve from the data very closely except for the 3000 hours bin. Whereas our model shows that the growth rate of earnings continues to increase as we move from the 2750 hours bin to the 3000 hours bin, in the data this

[21] In fact, if we assume that a fixed fraction of time devoted to investment is not included in reported total work hours then our model provides a better match to these patterns in the data.

Figure 9: Model Fit for the Relationship between Earnings and Hours

(a) Correlation of
log Earnings-log Hours

(b) Lifetime Hours
and Earnings Growth

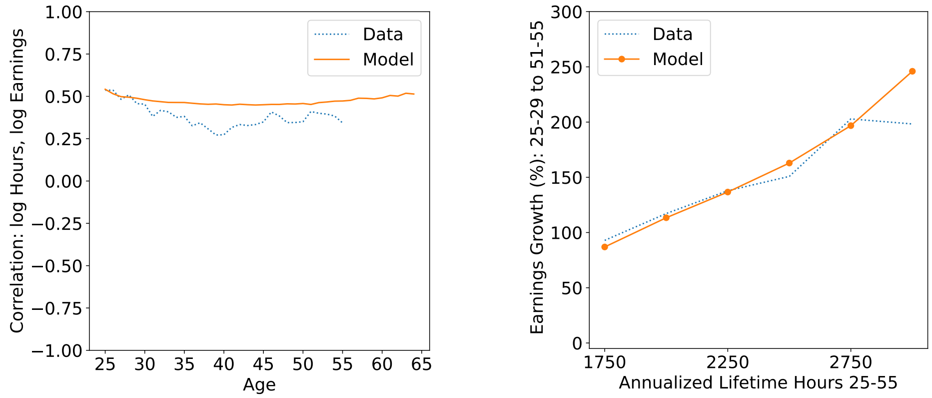

relationship is flat. One plausible rationale for this discrepancy is that in reality the returns to investment in human capital may diminish at very long hours. The significance of this discrepancy is somewhat limited, since both in the model and in the data there are very few individuals in the 3000 hours plus bin.

We now turn to moments that were not explicitly targeted by our calibration procedure. We did not explicitly target any moments based on wages, but given that we capture the patterns for earnings and hours individually and do a reasonable job of matching the profile for the correlation between earnings and hours, it is not surprising that our model does a reasonable job of matching the properties of wages. The three panels of Figure 10 display the results.

Panel (a) shows that the model closely tracks the evolution of mean log wages over the life-

Figure 10: Model Fit for Wages

(a) Mean of log Wages

(b) S.D. of log Wages

(c) Correlation of
log Earnings-log Hours

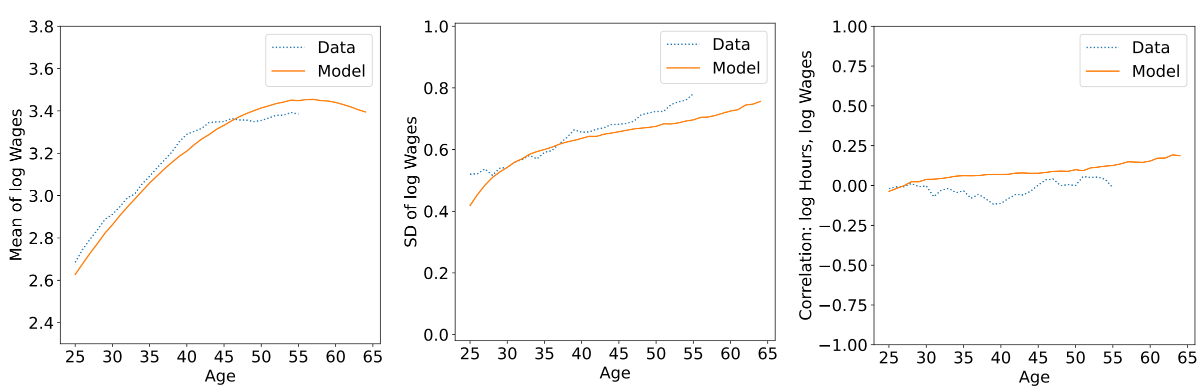

cycle. Panel (b) shows the model slightly understates the standard deviation of wages between ages 25 and 30 but fits the profile quite well between ages 30 and 55. The understatement of the standard deviation at young ages results from the fact that the model understates the standard deviation of hours over this age range. Panel (c) displays the correlation between hours and wages over the life-cycle. Consistent with our results for earnings, the model captures the fact that this correlation is near zero, but the model generates a slightly larger positive trend over the life-cycle than is present in the data.

It is noteworthy that our model closely captures the properties of both earnings and wages, as these measures are of independent interest. When studying consumption and wealth inequality, inequality of earnings is of primary importance; but when assessing heterogeneity in opportunities at a point in time, it is inequality in wage rates that is of particular interest. Notably, Huggett et al. (2011) studied only properties of earnings, and Guvenen et al. (2014) studied only properties of wages.

Figure 11 displays the distributions of annualized lifetime hours and annualized lifetime earnings in both the model and the data. Neither of these lifetime distributions were explicitly targeted. For earnings, we targeted the cross-sectional variance at ages 30 and 50 and two values of the autocorrelation function. For hours, we targeted the mean and variance for log annual hours in the overall cross-section and three values of the autocorrelation function.

Panel (a) shows that the distribution of lifetime hours in the model and data are broadly similar. However, the model does not generate sufficient concentration in the 2000 lifetime hours bin. This reflects the well-known issue that models with log-normally distributed heterogeneity cannot generate the high concentration found in the distribution of annual hours worked. Simply put, the large spike of individuals who report 2000 annual hours (40 hours a week for 50 weeks) cannot be well approximated with a normal distribution. Although the concentration in annualized lifetime

Figure 11: Model Fit of Lifetime Earnings and Hours

hours is less severe, the issue is still present. Consistent with this, the variance of log lifetime hours in the data is 0.025, which is modestly larger than the 0.023 value in the data. Although we could generate additional concentration in the hours distribution by introducing non-linearities into the mapping from hours to earnings, as in Bick et al. (2022), we have opted not to do so in order to better focus on the relationship between preference heterogeneity, lifetime hours and human capital accumulation.

Turning to lifetime earnings in Panel (b), we highlight two discrepancies between the model and the data. First, it is again true that the model does not generate the level of concentration found in the data, though the extent of the discrepancy is more modest for earnings than for hours. This is intuitive: Because there is substantial heterogeneity in wages for workers in the hours bin containing 2000, the concentration in hours yields much less concentration in earnings. Second, our model does not generate sufficient mass in the right tail of the distribution. This is a well-known issue in the literature for models like ours with log-normal shocks. Panel (b) confirms that we do not have sufficient mass in the final bin, which corresponds to individuals with annualized lifetime earnings greater than $150,000. When comparing the variance of log lifetime earnings in the model and the data these two effects partly offset each other, so that these variances are very similar: 0.348 in the model versus 0.346 in the data.

Panel (c) of Figure 11 shows how average annualized lifetime earnings vary across bins of the annualized lifetime hours distribution. Although not targeted, the model and data profiles track each other very closely, except for the highest hours bin. The discrepancy arises because the model-based relationship is fairly linear, whereas the relationship in the data is slightly concave. This suggests that it might be reasonable to introduce some additional non-linearities to the investment production function at both low and high hours.

Lastly, we examine whether the model generates a reasonable joint relationship between lifetime earnings, lifetime hours, and learning ability. In Section 3, we estimated regressions of log lifetime earnings and life-cycle earnings growth on log lifetime hours and AFQT percentile, which we considered a reasonable proxy for learning ability in our model. The results in Table 1 were for our overall balanced panel. We now re-run these regressions in our baseline sample of highly attached men and in our model-simulated data. Table 4 shows that in both regressions we get similar coefficients on log lifetime hours and AFQT / ability as well as similar R-squareds. In particular, the estimated lifetime hours coefficients in the model and data regressions are not significantly different from one another at the 5% level.

Table 4: Model Fit of Regression Analysis

(a) log of Lifetime Earnings

	Data	Model
log Lifetime Hours	1.43***	1.36
	(0.07)	
AFQT Percentile	0.01***	0.01
	(0.00)	
R^2	0.36	0.47

(b) Life-Cycle Earnings Growth

	Data	Model
log Lifetime Hours	1.82***	2.34
	(0.30)	
AFQT Percentile	0.02***	0.01
	(0.00)	
R^2	0.08	0.08

In all regressions, log annualized lifetime hours and AFQT percentiles are demeaned. Standard errors are in parentheses. *,** , and *** indicate significance at the 10%, 5%, and 1% levels, respectively. The model based regressions are based on a large sample so standard errors are effectively zero. In panel (b), we first compute, for each individual, average earnings at ages 25–29 and 51–55, considering only years with at least 520 hours worked. We then construct the growth rate between these two averages. The sample is restricted to individuals with at least three years of at least 520 hours worked in both the 25–29 and 51–55 age ranges. Appendix Table B.1 presents the full regression results for the data and also includes the regressions for women for the 1750+ lifetime hours sample.

7 Sources of Lifetime Earnings Inequality

In this section we use our calibrated model to shed light on the sources of lifetime earnings inequality. We note that while our model generates hours and earnings profiles from age 25 to 65, we compute all of our lifetime and annual measures in this section using only data from age 25 to 55 in order to be consistent with the age range covered in our empirical analysis. Including outcomes for the final ten years (i.e., between 56 and 65) has little effect on our conclusions.

7.1 Heterogeneity in Hours of Work and Earnings Inequality

The novel feature of our analysis is its focus on matching the extent and nature of heterogeneity in life-cycle hours profiles across individuals. In this subsection we highlight the contribution of heterogeneity in life-cycle hours profiles to lifetime earnings inequality. To do this, we consider a counterfactual in which individuals are not free to choose their total hours of work, $n+s$. Instead, we assume that all individuals must choose the same total hours of work in each period of their working life; i.e., $n_{i,t}+s_{i,t}=\hat{h}$ for all i and t. This counterfactual corresponds to the specification in Huggett et al. (2011), as they assumed total hours were constant over time and the same for all individuals. We implement this by choosing $\hat{h}$ equal to mean total hours for non-retired individuals in our calibrated model. We re-solve the model with this additional constraint on total working time, assuming that each individual faces the exact same human capital shocks and draws of measurement error as in the original simulation.

Table 5: The Role of Hours Heterogeneity

	log Lifetime Earnings		log Earnings at 55 & 35		log Earnings at 55	
	Variance	% of BM	Δ Variance	% of BM	Variance	% of BM
Benchmark (BM)	0.348	100%	0.234	100%	0.620	100%
Counterfactual: $n_{i,t} + s_{i,t} = \hat{h}$	0.284	81.7%	0.138	59.2%	0.444	71.6%

Notes: The first two columns report the variance of log lifetime earnings. The middle two columns report the increase in the variance of log earnings between ages 35 and 55. The last two columns report the variance of log earnings at age 55. The first row reports statistics from our benchmark model over the age range 25-55. The second row reports the same statistics for the counterfactual economy where total hours worked are restricted to be the same across all individuals and ages. Columns labeled percent of BM report the percent of the benchmark values that remain in the counterfactual.

Results for this counterfactual are reported in Table 5. The first row repeats statistics for our benchmark model, and the results in the second row report results for the counterfactual in which total hours are assumed to be constant over the life-cycle and across individuals.

Eliminating hours heterogeneity reduces the variance of log lifetime earnings by 18.3 percent, or nearly one fifth.[22] Notably, the drop in the variance of log lifetime earnings, 0.064, is much larger than the drop in the variance of log lifetime hours, 0.025. In a static model in which wages and hours are uncorrelated, reducing the variance of log hours will reduce the variance of log earnings by the same amount. This calculation does not translate directly to our Ben-Porath model, but it is suggestive that effects on human capital accumulation may be significant.

Two additional results in Table 5 also suggest an important role for human capital. Heterogeneity in human capital accumulation intuitively leads to heterogeneity in earnings growth and in earnings at age 55. When we remove heterogeneity in total hours, the increase in the variance of log earnings from age 30 to 55 falls from 0.234 to 0.138—a reduction of more than 40 percent. The variance of log earnings at age 55 also drops significantly, from 0.620 to 0.444, a decrease of nearly 30 percent.

To provide a quantitative assessment of the role of human capital we offer the following decomposition of the changes in lifetime earnings from eliminating differences in total hours of work. Let $e_{i,t} = n_{i,t} \cdot x_{i,t}$ be the earnings for individual i at age t in our benchmark model, and denote the three analogous series from the counterfactual as $\hat{e}_{i,t}$, $\hat{n}_{i,t}$, and $\hat{x}_{i,t}$. Lifetime earnings in the benchmark

[22]This exercise fixes total hours to be the same across all individuals but still allows individuals to vary the allocation of time between producing and investing. If we instead impose that each of the profiles for production time and investment time are constant across time and individuals, the variance of log of log lifetime earnings decreases by more than 23 percent. This larger value reflects the overall importance of time allocation.

model and the counterfactual are defined as:

$$\bar{e}_i^{BM} = \frac{\sum_{t=1}^{31} n_{it} \cdot x_{it}}{31} \tag{8}$$

$$\bar{e}_i^{CF} = \frac{\sum_{t=1}^{31} \hat{n}_{it} \cdot \hat{x}_{it}}{31} \tag{9}$$

Lifetime earnings in the counterfactual reflect changes in both the n_t and x_t sequences. We define two additional measures of lifetime earnings to capture the separate effects of changes in each of these two sequences:

$$\bar{e}_i^{N} = \frac{\sum_{t=1}^{31} \hat{n}_{it} \cdot x_{it}}{31} \tag{10}$$

$$\bar{e}_i^{X} = \frac{\sum_{t=1}^{31} n_{it} \cdot \hat{x}_{it}}{31} \tag{11}$$

The measure $\bar{e}_i^N$ holds the human capital profile fixed at its level in the benchmark model and considers only the effect of changes in the profile of production time. The measure $\bar{e}_i^X$ holds the profile for production time fixed and considers only the effects of changes in the human capital profile.

We then compute the variance of log lifetime earnings using each of these measures and compute the percent change between the benchmark model and the other three measures. We will refer to the change between $\bar{e}_i^{BM}$ and $\bar{e}_i^{CF}$ as the total effect, the change between $\bar{e}_i^{BM}$ and $\bar{e}_i^{N}$ as the direct channel, and the change between $\bar{e}_i^{BM}$ and $\bar{e}_i^{X}$ as the human capital channel. We note that the direct and human capital channels will not necessarily sum to the total change, since production and investment profiles can be correlated.

Results of this decomposition exercise are reported in Table 6. The direct channel delivers an 8.8 percent decrease in the variance of log lifetime earnings, while the human capital channel delivers a 7.9 percent decrease. The sum of these two channels is slightly smaller than the total effect.

Assessing the contribution of the human capital channel requires that one take a stand on how to assign the interaction effects. For our headline number we assign the interaction effects proportionately. If we do this, the share of the overall decline in earnings inequality from removing hours heterogeneity that is due to the human capital channel is 47 percent $\left(= \frac{100-92.1}{(100-92.1)+(100-91.2)}\right)$. We can also generate an interval of values by considering the two extremes in which we assign all of the interaction effects to either the direct or human capital channels. If we assign all of the interaction effects to the human capital channel, then the human capital channel is assigned everything not accounted for by the direct channel, which implies a 52 percent $\left(= \frac{91.2-81.7}{100-81.7}\right)$ share for the human capital channel. If we assign all of the interaction effects to the direct channel, this figure

Table 6: The Role of Human Capital

	log Lifetime Earnings	
	Variance	% of BM
Benchmark (BM)	0.348	100%
Counterfactual: $n_{i,t} + s_{i,t} = \bar{h}$	0.284	81.7%
Direct Channel	0.317	91.2%
Human Capital Channel	0.320	92.1%

Notes: The first row reports the variance of log lifetime earnings over the age range 25-55. In the subsequent rows, we report the same statistic under various counterfactuals. The second column reports the percent of the benchmark values that remain in each counterfactual.

becomes 43 percent $\left(= \frac{92.1-81.7}{100-81.7}\right)$. Focusing on either our headline number of 47 percent or the range of 43 to 52 percent, we conclude that the human capital channel is quantitatively important, accounting for roughly half of the total effect.

7.2 The Role of Preference Heterogeneity

The previous subsection showed that differences in hours worked across individuals are an important contributor to lifetime earnings inequality. We now show that differences in lifetime hours are almost entirely driven by heterogeneity in preferences (as opposed to heterogeneity in initial human capital, learning ability, and human capital shock realizations). While permanent and transitory preference heterogeneity both contribute to differences in lifetime hours, they have very different implications for lifetime earnings inequality.

We begin by considering a counterfactual that eliminates all preference heterogeneity in our calibrated model by setting $\psi_i = \mu_\psi$ and $\pi_{i,t} = 1$ for all i and t. This is similar to the specification in Guvenen et al. (2014), in which hours were endogenous but preferences were homogeneous and time invariant. We then re-solve the model, assuming that each individual experiences the same sequence of shocks to human capital and measurement error as in the benchmark economy. The second row of Table 7 shows the results from this exercise. To facilitate comparison, the first row repeats the results for our benchmark model.

Eliminating preference heterogeneity removes most of the variation in cross-sectional hours, and virtually all the variation in lifetime hours. Pooling all observations for individuals between ages 25 and 55, the variance of log annual hours falls by 85 percent, from 0.079 to 0.010. The variance of log lifetime hours falls by 96 percent, from 0.025 to 0.001.[23]

[23]The variation in annual hours in our homogenous preference counterfactual aligns with the variation in Guvenen et al. (2014), which, as noted earlier, assumes no preference heterogeneity and is calibrated to a sample of male

Table 7: The Role of Preference Heterogeneity

	Variance of log Hours		Variance of log Earnings
	Annual	Lifetime	Lifetime
Benchmark (BM)	0.079	0.025	0.348
$\sigma_\psi = \sigma_\pi = 0$	0.010	0.001	0.288
$\sigma_\psi = 0, \sigma_\psi > 0$	0.066	0.009	0.294
$\sigma_\psi > 0, \sigma_\pi = 0$	0.022	0.012	0.311

Notes: The first row reports statistics from our benchmark model over the age range 25-55. The subsequent rows report the same statistics under various counterfactuals that eliminate one or both sources of preference heterogeneity.

Eliminating preference heterogeneity reduces the variance of log lifetime earnings by 17.2 percent, from 0.348 to 0.288. Recalling that the overall impact of hours heterogeneity on lifetime earnings was a reduction of 18.3 percent, it follows that 94 percent of the overall impact of hours on lifetime earnings is driven by preference heterogeneity. Together with the findings in the previous paragraph, we conclude that preference heterogeneity is essential for the model to capture the empirical variation in hours and the role of hours for lifetime earnings inequality.

Our model features both permanent and transitory heterogeneity in work preferences. Which source of heterogeneity has a larger impact on lifetime inequality in hours and earnings? To answer this question, we start from the specification with no preference heterogeneity (i.e., the second row of Table 7) and compare it with versions of the model with only one source of heterogeneity (rows three and four of Table 7).

We find that the transitory component of preference heterogeneity has a larger impact on cross-sectional hours inequality than the permanent component, while the reverse is true for inequality in lifetime hours. Starting from the case with no preference heterogeneity (row 2), adding transitory preference heterogeneity (row 3) increases the variance of log annual hours by 0.056. This increase is nearly five times larger than when we add permanent preference heterogeneity (an increase of 0.012 from row 2 to row 4). By contrast, transitory heterogeneity increases the variance of log lifetime hours by 0.008 versus an increase of 0.011 for permanent heterogeneity. In other words, transitory preference heterogeneity has a much larger effect on cross-sectional hours inequality than lifetime hours inequality, while permanent heterogeneity has a nearly identical effect on both measures.

This pattern is even more stark when looking at earnings. Starting from the case with no preference heterogeneity, adding transitory heterogeneity increases the variance of log lifetime

workers aged 25 to 55 in the PSID. Specifically, they report that their data sample produces a standard deviation of log hours of 0.369, while their model produces a standard deviation of log hours of 0.112 (corresponding to a variance of 0.013). Since they do not examine the properties of lifetime hours, our finding that the variance of lifetime hours is almost entirely eliminated is novel.

earnings by 0.006 versus an increase of 0.023 when adding permanent heterogeneity. That is, permanent preference heterogeneity increases lifetime earnings inequality by nearly four times as much as transitory heterogeneity. This is despite the fact that, as shown in the previous paragraph, both sources of preference heterogeneity produce a similar amount of inequality in lifetime hours. Intuitively, workers with permanently low disutility from work will invest more in human capital because they expect to work long hours throughout their career, which raises the return to human capital investment early in the life-cycle. By contrast, workers with temporarily low disutility from work do not expect to work much more 10 or 20 years in the future, and so will not increase their investment as much.

The results in Table 7 also show that there are important interaction effects between the two components of preference heterogeneity. When adding both components of preference heterogeneity together, the effect on the variance of log lifetime earnings is roughly twice as large as the sum of the individual effects (0.060 versus 0.029). This implies that even though transitory preference heterogeneity by itself is not an important source of lifetime earnings inequality, it is important to include it in the analysis.

7.3 Other Sources of Inequality

Huggett et al. (2011) quantified the role of heterogeneity in initial human capital, learning ability, and human capital shocks for lifetime earnings inequality. Their model imposed that all individuals worked the same total hours. How do their conclusions change in a model that matches the empirical dispersion in hours worked?[24]

Results for lifetime measures are reported in Table 8. (Table C.2 in the Appendix provides results for lifetime and annual measures.) The entries in this table represent the share of overall variance remaining after the indicated channels are eliminated. For comparison, the row repeats the earlier results of eliminating preference heterogeneity, as shown in Table 7.

The central exercise in Huggett et al. (2011) is to quantify the roles of initial conditions (initial human capital and learning ability) versus human capital shocks for lifetime earnings inequality. Their primary finding is that eliminating differences in initial human capital and learning ability reduces the variance of lifetime earnings by 61.5 percent. The same exercise in our model reduces the variance of log lifetime earnings by only 52.1 percent (see row 3, column 1). Their model is calibrated to match earnings inequality over the life-cycle but features no hours inequality. Intuitively, abstracting from hours heterogeneity requires larger differences in initial human capital

[24]Some caution is in order when interpreting these comparisons. For example, their sample is men in the PSID, while ours is men in the NLSY79 with at least 1750 lifetime hours. Also, in their model individuals work from ages 23-60, while in our model individuals work from 25-65, though we measure lifetime earnings based on ages 25-55.

Table 8: Sources of Variation in Lifetime Earnings and Hours

	Variance of log Lifetime Earnings	Hours
$\sigma_\psi = \sigma_\pi = 0$	82.8%	2.0%
$\sigma_z = 0$	60.0%	99.2%
$\sigma_x = \sigma_\alpha = 0$	48.9%	100.0%

Notes: Each row reports the percent of variance in the benchmark model that remains when eliminating different sources of heterogeneity.

in order to generate the same distribution of earnings. Their calibrated standard deviation of log initial human capital is 0.462, compared with 0.353 in our model.

When we eliminate human capital shocks in our model, the variance of log lifetime earnings decreases by 40.0 percent (see row 4, column 1). This is roughly the same as in Huggett et al. (2011). This is perhaps not too surprising given that our shock process has a variance that is quite close to theirs. An important implication implicit in this finding is that an endogenous hours margin does little to affect the propagation of human capital shocks to lifetime earnings, see column 2.

7.4 Hours Worked and Upward Mobility

A central component of the "American dream" is the idea that one can overcome poor initial conditions through hard work. The results presented in the last subsection indicate that both high lifetime hours and initial human capital are important determinants of lifetime earnings. In this subsection we take a closer look at the prevalence of the American dream in our calibrated model. In particular, we examine the extent to which higher lifetime hours compensates for lower initial human capital in our calibrated model. We will use the term upward mobility to describe individuals whose position in the lifetime earnings distribution is higher than their position in the initial human capital distribution.

To do this, we place individuals at age 25 in one of 25 bins. These bins are created using a two-step procedure. In the first step we sort individuals into five equally size bins based on their quintile in the distribution of initial human capital. In the second step, we divide each of these bins into five equally sized bins based on their quintile in the (bin-specific) distribution of ψ values.[25] By construction, each bin represents 4 percent of the overall population. Recall that in our model,

[25] Recall that initial human capital and ψ are negatively correlated, which implies that the distribution of ψ is bin-specific. An alternative approach is to construct quintiles separately for ψ and initial human capital, which implies that the mass of individuals varies across the joint distribution of quintiles. The qualitative conclusions we draw in this section remain very similar under this alternative grouping.

Figure 12: Hours Worked, Human Capital, and Earnings Mobility

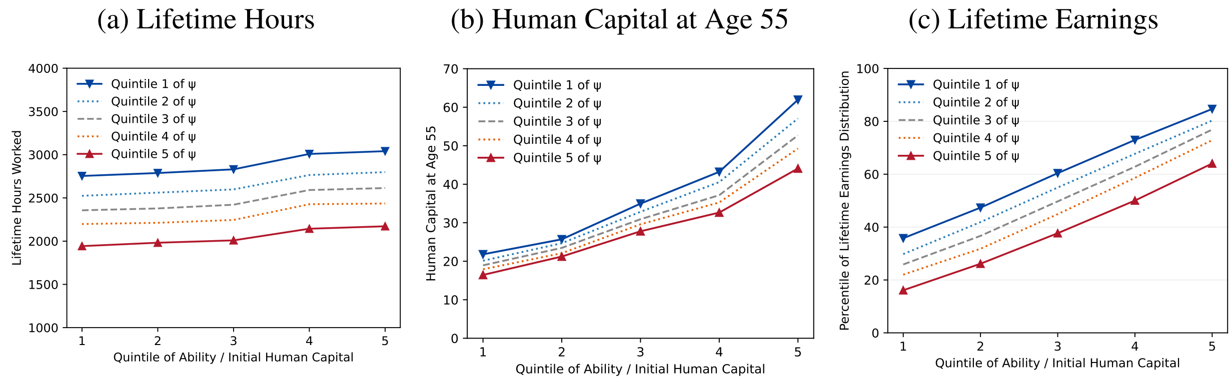

initial human capital and learning ability are perfectly correlated, so that variation in initial human capital is bundled with variation in learning ability.

Figure 12a shows the variation in mean lifetime hours worked across these 25 bins. Each line in the figure shows the relationship between initial human capital and mean lifetime hours worked holding the value of ψ constant. The relationship between mean lifetime hours of work and ψ holding initial human capital constant is revealed by comparing the relative height of the five lines at a given level of initial human capital. Two messages appear. First, mean lifetime hours do not substantially vary over the initial human capital distribution. Second, holding initial human capital constant, variation in ψ is associated with large differences in mean lifetime hours. (Because we are computing mean lifetime hours for each bin, variation due to shocks that occur over the life-cycle is being averaged out.)

Our analysis has emphasized the relationship between permanent preference heterogeneity, lifetime hours, and human capital accumulation. Figure 12b illustrates this by plotting mean human capital at age 55 for each of the 25 bins. The fact that each line slopes upward reflects the intuitive result that, on average, higher initial human capital leads to higher human capital at age 55. But the vertical distance between the lines shows that a lower value of ψ holding initial human capital constant also increases human capital at age 55. It follows that a lower ψ value can partially offset the effects of a lower initial human capital, and thus can be a source of upward mobility in terms of human capital. For example, moving from the highest to the lowest quintile of the (bin specific) ψ distribution has a slightly larger positive impact on human capital at age 55 than does moving an individual with a given value of ψ to the next highest quintile of initial human capital.

Figure 12c shows how mean lifetime earnings varies across the 25 bins. Once again, because we are reporting mean lifetime earnings for each bin we are averaging out the effects of shocks that happen over the life-cycle. To better illustrate the position of a given bin within the lifetime earnings distribution, the y-axis in this figure reflects mean earnings for each bin by its corresponding percentile in the overall lifetime earnings distribution.

The key point to notice in Figure 12c is that the distributions of (mean) lifetime earnings corresponding to different quintiles of the initial human capital distribution have very large areas of overlap. Specifically, individuals who are in the lowest quintiles of both initial human capital and ψ at age 25 have mean lifetime earnings comparable to those in the third ψ quintile among individuals in the second quintile of initial human capital. It follows that these individuals with both human capital and ψ in the lowest quintile will "leapfrog" many individuals with initial human capital in the second quintile in terms of lifetime earnings. This holds when considering other adjacent quintiles of the initial human capital distribution.

In summary, while previous sections establish that heterogeneity in lifetime hours increases lifetime earnings inequality, this section shows that high lifetime hours are also a source of upward mobility for many individuals who start with low initial human capital. Higher hours increase earnings both directly, through higher production time, but also indirectly through greater human capital accumulation. Because permanent preference heterogeneity is only mildly correlated with initial human capital ($\rho_{x,\psi} = -0.3$), many high-hours workers are born with low initial human capital and are therefore upwardly mobile.

8 Conclusion

A key goal of the literature on inequality is to understand the quantitatively important driving forces and mechanisms that generate inequality. In this paper, we use the NLSY79 to document large differences in lifetime hours of work across individuals and then use a heterogeneous agent model of labor supply and human capital accumulation to argue that these differences in lifetime hours play a quantitatively important role in shaping lifetime earnings inequality. In particular, we find that heterogeneity in hours of work over the life-cycle accounts for almost 20 percent of the variance of log lifetime earnings. Over 90 percent of this effect is due to heterogeneity in preferences, and roughly one-half of this effect reflects human capital accumulation. A key message from our analysis is that it is important to include heterogeneity in hours of work in analyses of inequality.

We close by noting three important areas for future research. First, we have relied on preference heterogeneity as a parsimonious way to generate the salient features of hours heterogeneity found in the micro data. Our specification of preference heterogeneity essentially amounts to a wedge in the first-order conditions involving hours. This wedge can capture the effect of factors beyond true preference heterogeneity. For some specific issues, like the design of disability insurance, it is important to have a deeper analysis of these factors. More generally, this will be important for assessing issues that involve welfare comparisons.

The second is to extend our quantitative analysis to groups beyond the sample of highly attached men that we studied. Differences in lifetime hours of work are even larger if we consider broader groups, raising the possibility that effects for the overall population will be even larger than we find. Extending the analysis to men with low levels of attachment may require a richer specification beyond the one we have used here. In addition to incorporating an explicit extensive margin, it may be necessary to allow for correlation between human capital shocks and preference shocks, to generalize the human capital accumulation process to allow for differential rates of depreciation during extended spells of non-participation, and to explicitly model features of the tax and transfer system. The recent paper by Hosseini et al. (2025) includes these features in their analysis of health shocks. Extending the analysis to women will also require explicit modeling of fertility. In view of our results about the importance of permanent heterogeneity, it will potentially be critical to correctly capture expectations of future work hours. This will require a richer model that explicitly accounts for spells of non-participation.

A third area is to consider a range of policies that directly affect the distribution of hours of work. Several European countries have enacted legislation to reduce the standard workweek. Another class of policies aims to compress the distribution of hours worked from above. For example, all advanced economies have legislation that stipulates a threshold for overtime hours and the level of the overtime premium. France offers a more direct example of this class of policies. Specifically, in a pair of laws adopted between 1998-2000 and rolled out between 2000-2002 ("Aubry I and Aubry II"), France imposed a 35-hour workweek and adopted a regulation imposing that most workers could work no more than 48 hours per week.[26] Relatedly, in 2024 Greece enacted legislation to make it easier for firms to have a six day workweek that would imply a 48-hour workweek. Our partial equilibrium model implies that these policies have large effects on both the mean and dispersion of lifetime earnings, holding prices fixed. A more complete examination of the effects of these policies on both outcomes and welfare will require a general equilibrium analysis that introduces a richer specification of firm decisions.

References

Badel, Alejandro, Mark Huggett, and Wenlan Luo. 2020. "Taxing top earners: a human capital perspective." *The Economic Journal* 130 (629): 1200–1225.

Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. "Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics." *American Economic Review* 104 (6): 1551–1596.

[26] In the opposite direction, France also recently enacted a regulation imposing a minimum workweek for part-time workers. Carry (2024) develops a model to study the effect of this regulation.

Benabou, Roland. 2002. “Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?” *Econometrica* 70 (2): 481–517.

Bick, Alexander, Adam Blandin, and Richard Rogerson. 2022. “Hours and Wages.” *The Quarterly Journal of Economics* 137 (3): 1901–1962.

Bick, Alexander, Adam Blandin, and Richard Rogerson. 2025. “After 40 Years, How Representative Are Labor Market Outcomes in the NLSY79?.”

Blundell, Richard, Monica Costa-Dias, David Goll, and Costas Meghir. 2021. “Wages, Experience, and Training of Women over the Life Cycle.” *Journal of Labor Economics* 39 (S1): S275–S315.

Blundell, Richard, Monica Costa Dias, Costas Meghir, and Jonathan Shaw. 2016. “Female Labor Supply, Human Capital, and Welfare Reform.” *Econometrica* 84 (5): 1705–1753.

Browning, Martin, Lars Peter Hansen, and James J Heckman. 1999. “Micro data and general equilibrium models.” *Handbook of macroeconomics* 1 543–633.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. “Wage bargaining with on-the-job search: Theory and evidence.” *Econometrica* 74 (2): 323–364.

Carry, Pauline. 2024. “The Effects of the Legal Minimum Working Time on Workers, Firms and the Labor Market.” Working Paper, Princeton University.

Chang, Yongsung, Sun-Bin Kim, Kyooho Kwon, and Richard Rogerson. 2020. “Cross-sectional and aggregate labor supply.” *European Economic Review* 126 103457.

Deaton, Angus, and Christina Paxson. 1994. “Intertemporal choice and inequality.” *Journal of Political Economy* 102 (3): 437–467.

Erosa, Andres, Luisa Fuster, and Diego Restuccia. 2016. “A quantitative theory of the gender gap in wages.” *European Economic Review* 85 165–187.

Fan, Xiaodong, Ananth Seshadri, and Christopher Taber. 2024. “Estimation of a Life-Cycle Model with Human Capital, Labor Supply, and Retirement.” *Journal of Political Economy* 132 (1): 48–95.

Farr, William. 1853. “The Income and Property Tax.” *Journal of the Statistical Society of London* 16 (1): 1–44.

Fillmore, Ian, and Trevor S Gallen. 2023. “Is Earnings Inequality Driven by Talent or Tastes?” working paper, Washington University in St. Louis and Purdue University.

Glover, Andrew, José Mustre-del Río, and Emily Pollard. 2022. “Lifetime Earnings Differences across Black and White Individuals: Years Worked Matter..” *Federal Bank of Kansas Economic Review* 108 (1): 5–22.

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

Guvenen, Fatih, Burhanettin Kuruscu, and Serdar Ozkan. 2014. "Taxation of human capital and wage inequality: A cross-country analysis." *The Review of Economic Studies* 81 (2): 818–850.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2014. "Consumption and labor supply with partial insurance: An analytical framework." *American Economic Review* 104 (7): 2075–2126.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2017. "Optimal tax progressivity: An analytical framework." *The Quarterly Journal of Economics* 132 (4): 1693–1754.

Heckman, James J. 1976a. "Estimates of a Human Capital Production Function Embedded in a Life-Cycle Model of Labor Supply." In *Household Production and Consumption*, 225–264, National Bureau of Economic Research, Inc.

Heckman, James J. 1976b. "A Life-Cycle Model of Earnings, Learning, and Consumption." *Journal of Political Economy* 84 (4): S9–S44.

Heckman, James J, Lance Lochner, and Christopher Taber. 1998. "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents." *Review of Economic Dynamics* 1 (1): 1–58.

Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon M Phillips. 2024. "Production and learning in teams." *Econometrica* 92 (2): 467–504.

Hosseini, Roozbeh, Karen A Kopecky, and Kai Zhao. 2025. "How important is health inequality for lifetime earnings inequality?" *Review of Economic Studies* (forthcoming).

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2006. "Human capital and earnings distribution dynamics." *Journal of Monetary Economics* 53 (2): 265–290.

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

Imai, Susumu, and Michael P Keane. 2004. "Intertemporal labor supply and human capital accumulation." *International Economic Review* 45 (2): 601–641.

Jarosch, Gregor. 2023. "Searching for job security and the consequences of job loss." *Econometrica* 91 (3): 903–942.

Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg. 2021. "Learning from coworkers." *Econometrica* 89 (2): 647–676.

Kaplan, Greg. 2012. “Inequality and the life cycle.” *Quantitative Economics* 3 (3): 471–525.

Keane, Michael P, and Nada Wasi. 2016. “Labour supply: the roles of human capital and the extensive margin.” *The Economic Journal* 126 (592): 578–617.

Keane, Michael P., and Kenneth I. Wolpin. 1997. “The Career Decisions of Young Men.” *Journal of Political Economy* 105 (3): 473–522.

Lise, Jeremy, and Fabien Postel-Vinay. 2020. “Multidimensional skills, sorting, and human capital accumulation.” *American Economic Review* 110 (8): 2328–2376.

Manuelli, Rodolfo, Ananth Seshadri, and Yongseok Shin. 2012. “Lifetime Labor Supply and Human Capital Investment.”Technical report, National Bureau of Economic Research.

McDaniel, Cara. 2007. “Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003.” *Manuscript, Arizona State University* 19602004.

Morchio, Iacopo. 2020. “Work Histories and Lifetime Unemployment.” *International Economic Review* 61 (1): 321–350.

Neal, Derek, and Sherwin Rosen. 2000. “Chapter 7 Theories of the distribution of earnings.” In *Handbook of Income Distribution*, Volume 1. 379–427, Elsevier.

Ozkan, Serdar, Jae Song, and Fatih Karahan. 2023. “Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job-Ladder Risk versus Human Capital.” *Journal of Political Economy Macroeconomics* 1 (3): 506–550.

Papageorgiou, Theodore. 2014. “Learning your comparative advantages.” *Review of Economic Studies* 81 (3): 1263–1295.

Rauh, Christopher, and Arnau Valladares-Esteban. 2023. “On the black-white gaps in labor supply and earnings over the lifecycle in the US.” *Review of Economic Dynamics* 51 424–449.

Shaw, Kathryn L. 1989. “Life-Cycle Labor Supply with Human Capital Accumulation.” *International Economic Review* 30 (2): 431–456.

Stantcheva, Stefanie. 2017. “Optimal Taxation and Human Capital Policies over the Life Cycle.” *Journal of Political Economy* 125 (6): 1931–1990.

Storesletten, Kjetil, Christopher I Telmer, and Amir Yaron. 2004. “Consumption and risk sharing over the life cycle.” *Journal of monetary Economics* 51 (3): 609–633.

Wallenius, Johanna. 2011. “Human capital accumulation and the intertemporal elasticity of substitution of labor: How large is the bias?” *Review of Economic Dynamics* 14 (4): 577–591.

Wallenius, Johanna. 2013. “Social security and cross-country differences in hours: A general equilibrium analysis.” *Journal of Economic Dynamics and Control* 37 (12): 2466–2482.

ONLINE APPENDIX

A Proof for Optimal Human Capital Investment Condition (1)

The sequential formulation of the individual's utility maximization problem in Section 3.1 is:

$$\max_{\{c_t,n_t,s_t\}_{t=1}^T} \quad \sum_{t=1}^{T} \beta^t u(c_t, n_t + s_t) \tag{A.1}$$

$$s.t. \quad \sum_{t=1}^{T_R-1} \frac{c_t - wx_t n_t}{(1+r)^t} = 0 \tag{A.2}$$

$$x_{t+1} = (1-\delta)x_t + \alpha(x_t s_t)^{\phi} \ \forall t \tag{A.3}$$

$$n_t, s_t \geq 0 \, \forall t, \text{ with equality if } t \geq T_R \tag{A.4}$$

This problem can be written recursively as:

$$V_t(x,k) = \max_{n,s,k'} \quad u((1+r)k + wxn - k', n+s) + \beta V_{t+1}(x', k') \tag{A.5}$$

$$s.t. \quad x' = (1-\delta)x + \alpha(xs)^{\phi} \tag{A.6}$$

$$n, s \geq 0, \text{ with equality if } t \geq T_R \tag{A.7}$$

Following Guvenen et al. (2014), it is helpful to rewrite this problem in terms of new human capital produced rather than in terms of investment. Define the following variables: total hours, $h_t = n_t + s_t$, investment share of total hours, $i_t = s_t/h_t$, newly produced human capital as $Q_t = \alpha(x_t h_t i_t)^{\phi}$, and the opportunity cost of newly produced human capital $C(Q_t) = w(Q_t/\alpha)^{1/\phi} = wx_t h_t i_t$. With these variables defined, we can rewrite the individual's recursive problem as

$$\begin{aligned} V_t(x,k) = \max_{h,Q,k'} \quad & u((1+r)k + whx - C(Q) - k', h) + \beta V_{t+1}(x', k') \\ s.t. \quad & x' = (1-\delta)x + Q \\ & h, Q \geq 0, \text{ with equality if } t \geq T_R \end{aligned}$$

We can characterize the optimal investment choice (assuming an interior solution) with a FOC and an envelope condition:

$$FOC(Q_t): \quad C'(Q)u_1(c,h) = \beta V_{t+1,1}(x',k') \tag{A.8}$$

$$Env(x_t): \quad V_{t,1}(x,k) = u_1(c,h)wh + \beta V_{t+1,1}(x',k')(1-\delta) \tag{A.9}$$

This formulation has two useful results. First, holding h fixed, Q only affects utility via consumption, not via leisure. Second, an individual's human capital choice yesterday does not affect the opportunity cost of Q today. For intuition on this latter point, note that if today a worker wants to produce Q, they need to set the product $xhi = (Q/\alpha)^{1/\phi}$—perhaps surprisingly, this product is not affected by the worker's current level of human capital, x. For example, if x is high, then to produce Q the necessary hi is low, but the opportunity cost of each unit of time is high because x is high. Alternatively, if x is low then the necessary hi is high, but the opportunity cost of each unit of time is low because x is low.

Combining (A.8), (A.9) and iterating forward in time yields:

$$C'(Q_t) = \sum_{t'=t+1}^{T_R-1} \beta^{t'-t} w(1-\delta)^{t'-t-1} h_{t'} \left(\frac{u_1(c_{t'},h_{t'})}{u_1(c_t,h_t)} \right) \tag{A.10}$$

where $c_{t'} = (1+r)k_{t'} + wh_{t'}x_{t'} - C(Q_{t'}) - k_{t'+1}$. From the Euler equation, we know that

$$\beta^{t'-t}(1+R)^{t'-t} = \frac{u_1(c_t,h_t)}{u_1(c_{t'},h_{t'})} \tag{A.11}$$

and substituting this into the previous equation yields

$$C'(Q_t) = \sum_{t'=t+1}^{T_R-1} \frac{w(1-\delta)^{t'-t-1} h_{t'}}{(1+R)^{t'-t}} \tag{A.12}$$

In words, this says that the optimal investment choice equates the static marginal cost of investing (foregone earnings) to the sum of remaining total work hours, scaled by the wage rate and discounted by both present value and the human capital depreciation rate. Since $C'(Q) = \left(\frac{w}{\alpha\phi}\right)\left(\frac{Q}{\alpha}\right)^{\frac{1}{\phi}-1}$, the marginal cost is increasing in Q. Therefore, we can conclude that investment is higher for individuals who will work more hours in the future (where future hours are discounted by the depreciation rate and the interest rate), as in (A.12).

To derive Equation (1) specifically, substitute the above expression for $C'(Q_t)$:

$$\left(\frac{w}{\alpha\phi}\right)\left(\frac{Q_t}{\alpha}\right)^{\frac{1}{\phi}-1} = \sum_{t'=t+1}^{T_R-1} \frac{w(1-\delta)^{t'-t-1}h_{t'}}{(1+R)^{t'-t}} \tag{A.13}$$

Finally, substitute $Q_t = \alpha(x_t s_t)^{\phi}$ and rearrange terms to arrive at:

$$wx_t = \alpha\phi x_t^{\phi} s_t^{\phi-1} \sum_{t'=t+1}^{T_R-1} \frac{w(1-\delta)^{t'-t-1}h_{t'}}{(1+R)^{t'-t}} \tag{A.14}$$

B Data

FIGURE B.1: Cross-Sectional and Lifetime Distribution of Hours Worked

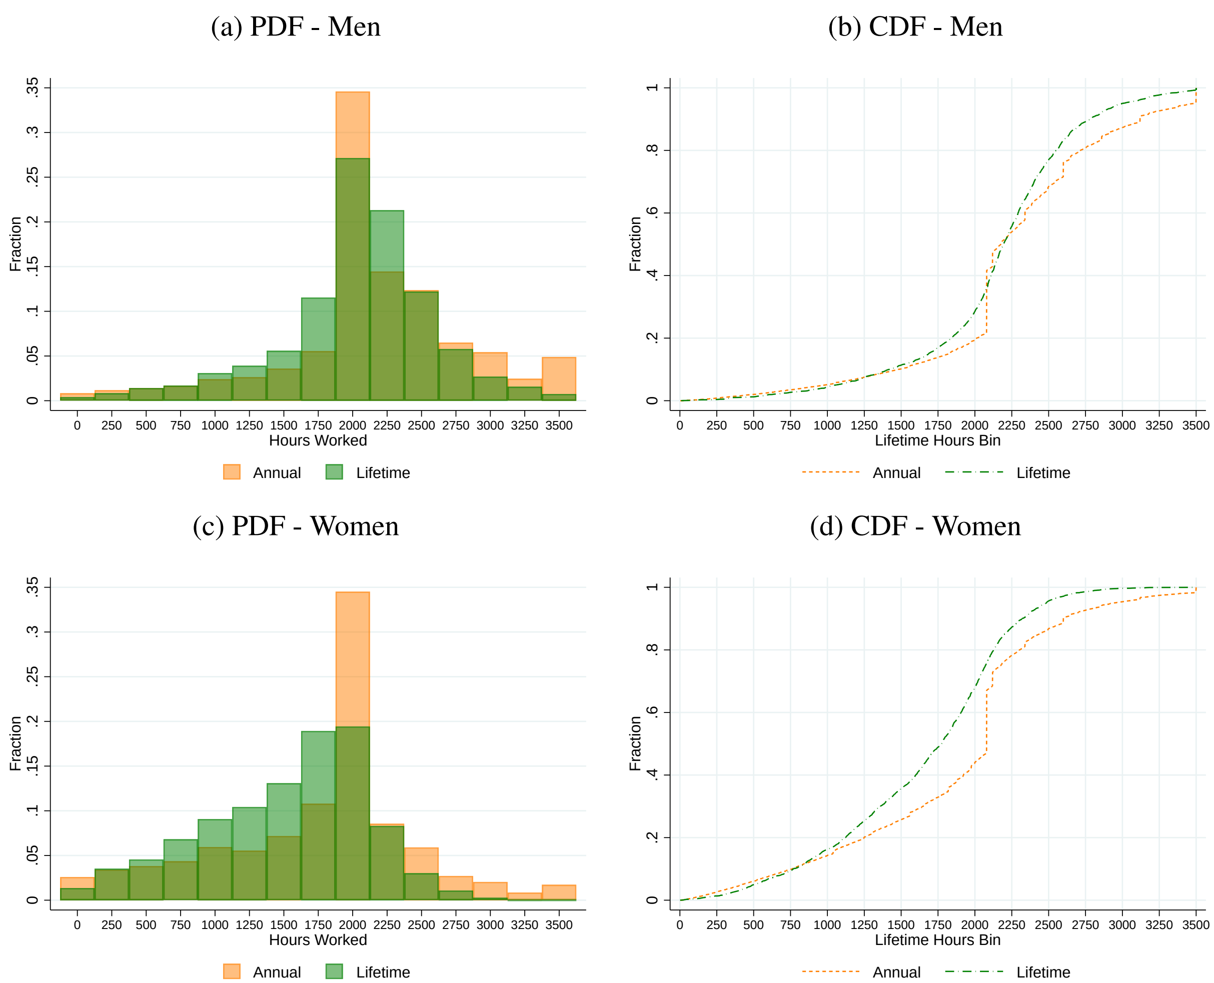

Notes: In the annual sample we only include person-year observations with positive hours worked. For men 7.4 percent and for women 15.5 percent of person-year observations feature zero hours worked.

B.1 Comparison of Autocorrelation Profiles

One may worry that the survey design and our imputation procedure creates too much persistence in annual hours worked and annual earnings.

Figure B.2 shows that the autocorrelation profile for anual hours excluding imputed values for annual hours is virtually identical to the profile shown in Figure 3, see two left panels. The two middle and two right panels distinguish between the period when the NLSY79 was conducted annually and when it switched to being conducted every other years. Note that even during this period we still have observations for hours for every year. In each case, the autocorrelation profile using only direct reports or both, the direct reports and imputed values, are almost identical.

Figure B.3 performs the same comparison for earnings, focusing only the sample of men used in our quantitative analysis (lifetime hours are at least 1750). Before 1994, earnings were collected annually, and during this period the two profiles are virtually identical (middle panel). From 1994 onward, earnings are only reported every other year. The two-year auto-correlation in the directly reported data is slightly lower than when including imputed values. From lag 4 onward, the two profiles again lie almost on top of each other.

FIGURE B.2: Autocorrelation of Annual Hours and Annual Earnings

Men

(a) Annual Hours (b) Annual Earnings 1978-1993 (c) Annual Earnings 1994-2019

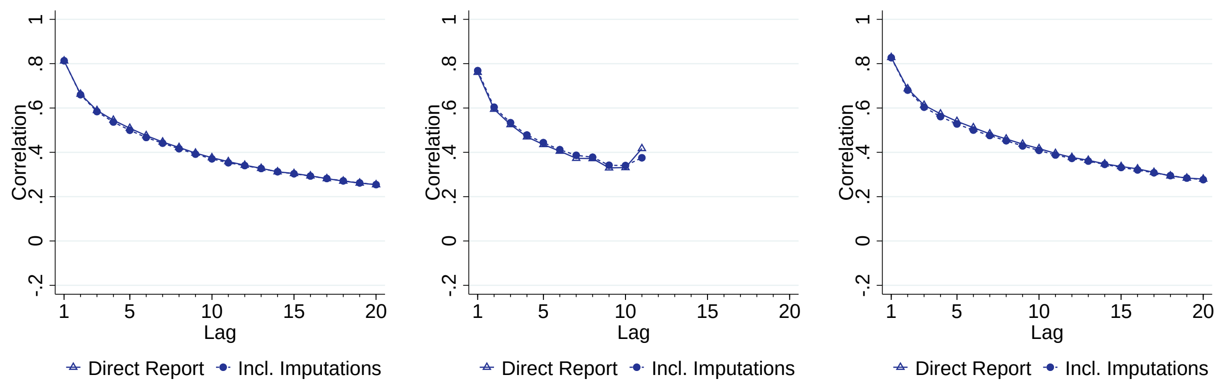

Women

(d) Annual Hours (e) Annual Earnings 1978-1993 (f) Annual Earnings 1994-2019

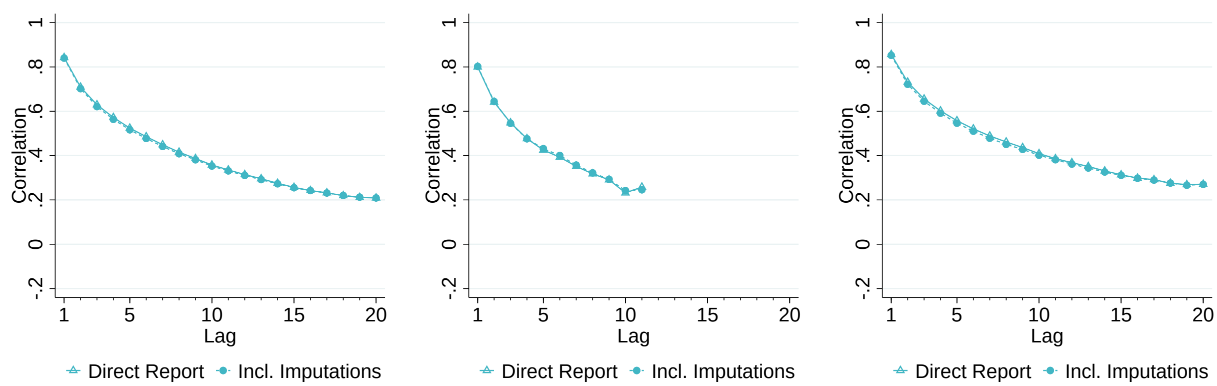

FIGURE B.3: Autocorrelation of Annual Earnings for Men with Lifetime Hours ≥ 1750

(a) All Years (b) 1978-1993 (c) 1994-2019

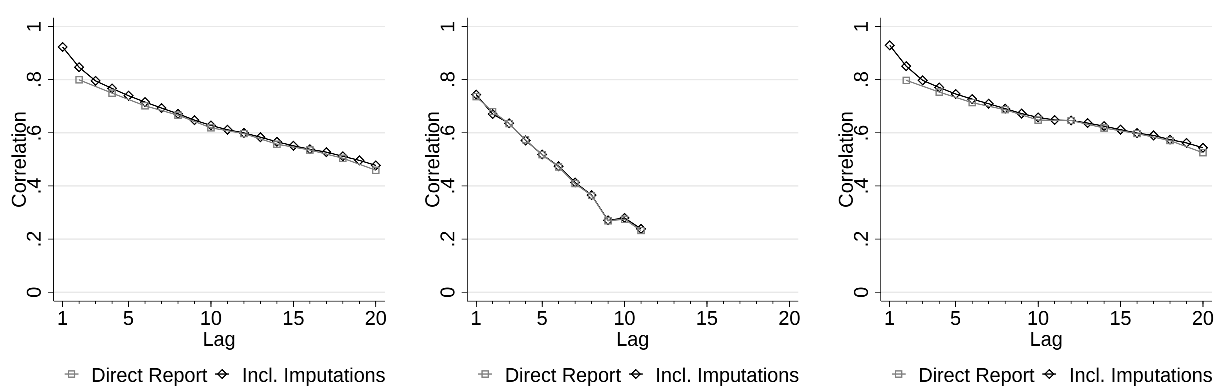

B.2 Regressions for the Lifetime Hours $\geq$ 1750 Sample

Table B.1: Lifetime Hours and Earnings Regressions for Lifetime Hours $\geq$ 1750 Sample

Men

(a) log of Lifetime Earnings

	(1)	(2)
log Lifetime Hours	1.57***	1.43***
	(0.08)	(0.07)
AFQT Percentile		0.01***
		(0.00)
Constant	11.00***	11.00***
	(0.01)	(0.01)
N	2218	2128
R^2	0.16	0.36

(b) Life-Cycle Earnings Growth

	(1)	(2)
log Lifetime Hours	2.02***	1.82***
	(0.30)	(0.30)
AFQT Percentile		0.02***
		(0.00)
Constant	1.08***	1.09***
	(0.06)	(0.06)
N	2054	1969
R^2	0.02	0.08

Women

(c) log of Lifetime Earnings

	(1)	(2)
log Lifetime Hours	1.42***	1.22***
	(0.11)	(0.10)
AFQT Percentile		0.01***
		(0.00)
Constant	10.65***	10.65***
	(0.01)	(0.01)
N	1483	1457
R^2	0.10	0.25

(d) Life-Cycle Earnings Growth

	(1)	(2)
log Lifetime Hours	0.86*	0.82*
	(0.47)	(0.48)
AFQT Percentile		0.00
		(0.00)
Constant	1.33***	1.34***
	(0.06)	(0.06)
N	1375	1351
R^2	0.00	0.00

In all regressions, log annualized lifetime hours and AFQT percentiles are demeaned such that the constants are comparable across specifications (1) and (2). Standard errors are in parentheses. *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. In panels (b) and (d), we first compute, for each individual, average earnings at ages 25–29 and 51–55, considering only years with at least 520 hours worked. We then construct the growth rate between these two averages. The sample is restricted to individuals with at least three years of at least 520 hours worked in both the 25–29 and 51–55 age ranges.

B.3 Average Wage Growth in the CPS

We compute average annual growth of male wages in the CPS-ASEC by educational attainment and weight the values by the educational attainment shares for our NLSY79 sample of highly attached males. We first compute mean hourly wages for 25- to 29-year-old males separately by education in 1982, which corresponds to age 25 for our oldest cohort, $w_{25}^{old,e}$. We then compute mean hourly wages for 25- to 29-year-old males separately by education in 2012, which corresponds to age 55 for our oldest cohort, $w_{55}^{old,e}$. We compute average annual wage growth by education over this period as $g_w^{old,e} = \left(w_{55}^{old,e}/w_{25}^{old,e}\right)^{1/30} - 1$. We then compute the average annual growth rate across education groups using education weights from our NLSY79 sample, resulting in an average growth rate g_w^{old}. Next, we repeat this analysis using the years 1989-2019 (which correspond to ages 25-55 for our youngest cohort) to obtain g_w^{young}. The average of g_w^{young} and g_w^{old} is 0.5% per year.

C Quantitative Results

Table C.2: Sources of Variation in Earnings and Hours

	Variance of log Hours		Variance of log Earnings	
	Annual	Lifetime	Annual	Lifetime
$\sigma_\psi = \sigma_\pi = 0$	79.1%	82.8%	13.1%	2.0%
$\sigma_z = 0$	61.8%	60.0%	99.6%	99.2%
$\sigma_x = \sigma_\alpha = 0$	71.0%	48.9%	100.1%	100.0%

Notes: Each row reports the percent of variance in the benchmark that remains when eliminating different sources of heterogeneity.