Hours Worked and Lifetime Earnings Inequality*

Alexander Bick
Federal Reserve Bank of St. Louis, CEPR

Adam Blandin
Vanderbilt University

Richard Rogerson
Princeton University, NBER

December 20, 2023

Abstract

We document substantial differences between individuals in cumulative hours worked over a lifetime. We assess the extent to which these differences contribute to inequality in lifetime income through the lens of a quantitative structural model of human capital accumulation. We show that the model requires both permanent and transitory differences in preferences for work in order to match lifetime hours patterns. In addition to heterogeneity in leisure preferences, we also allow for heterogeneity in initial human capital and learning ability, as well as idiosyncratic shocks to human capital throughout the life-cycle. The parametrized model implies that one quarter of the variance in lifetime income is driven by heterogeneity in leisure preferences, and that much of this effect works through human capital investment.

*Contact: alexander.bick@stls.frb.org, adam.blandin@vanderbilt.edu, rdr@princeton.edu. We thank seminar/conference participants at the Bureau of Labor Statistics, Deutsche Bundesbank, Federal Reserve Banks of Minneapolis and St. Louis, I-85 Macro Workshop, LAEF Workshop on Labor Markets and Macroeconomic Outcomes, LAEF and St. Louis Human Capital Conference, Stanford Institute for Theoretical Economics Labor Market and Policies Workshop, Universtiy of Conneticut, University of Bonn, University of North Carolina at Charlotte, and Vanderbilt University for helpful commments and suggestions. Kevin Bloodworth II, Elizabeth Harding, Minju Jeong, 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 or the Federal Reserve System.

1 Introduction

A key objective of research on inequality is to isolate the quantitatively important forces that shape inequality. While the literature considers many factors, little attention has been devoted to the role of heterogeneity in labor supply.[1] In this paper we argue that heterogeneity in lifetime hours of work plays a quantitatively important role in shaping inequality, not only via its mechanical effect on earnings but also via its effect on human capital accumulation. The mechanism connecting heterogeneity in lifetime hours of work with heterogeneity in human capital accumulation is intuitive. In the Ben-Porath model of human capital accumulation, higher expected future hours of work increases the incentive for human capital accumulation today, and so will give rise to heterogeneity in human capital accumulation.[2] A key contribution of our paper is to quantify this channel.

Our quantification proceeds in three steps. The first step is to document the extent and nature of differences in lifetime hours of work. Perhaps surprisingly, while cross-sectional differences in hours of work have been studied extensively, relatively little is known about the differences in hours of work over longer horizons. We use the NLSY79 to document differences in lifetime hours of work as well as properties of the stochastic process that generates these differences. Our analysis focuses on a sample of males with at least a high school education and yields four key findings. First, even though our sample represents a group with very high labor force attachment, differences in lifetime hours are large. Mean annualized lifetime hours in our sample are 2154, with a standard deviation of almost 600. The interquartile ratio is 1.26. Second, almost all of the variation captured by the interquartile ratio is due to differences in average weekly hours while employed. Third, cross-sectional differences reflect both a transitory and a persistent component. In particular, hours are not well captured by a simple AR(1) process. Fourth, we document a strong positive correlation between lifetime hours and lifecycle earnings growth.

The second step in our analysis consists of a simple exercise in which we use existing estimates of Ben-Porath models to assess the importance of the heterogeneity in lifetime hours documented in the first step for heterogeneity in life cycle earnings growth.[3] In particular, the parameter values that we use for the Ben-Porath learning technology are the same as have been used in the literature to assess the quantitative importance of other sources of heterogeneity on life cycle earnings growth. In this exercise the only source of heterogeneity among workers is heterogeneity in tastes for work, which will in turn generate differences both in lifetime hours of work and the division of total hours between producing and investing. We calibrate this heterogeneity to match observed

[1] Important contributions include Heckman, Lochner and Taber (1998), Rubinstein and Weiss (2006),Kuruscu (2006), Huggett, Ventura and Yaron (2006, 2011), Bagger et al. (2014), Guvenen, Kuruscu and Ozkan (2014), Badel, Huggett and Luo (2020), Ozkan, Song and Karahan (2023) and Hosseini, Kopecky and Zhao (2022).

[2] Alternatively, if human capital accumulation is the result of learning by doing then heterogeneity in hours worked will directly generate heterogeneity in human capital accumulation. See, for example, Imai and Keane (2004) and Blandin (2018).

[3] One could also carry out this exercise in a model that features learning by doing. We leave this for future work.

heterogeneity in hours and the autocorrelation properties of hours. Our calibrated model accounts for roughly twenty five percent of the variance of the log of lifetime earnings. In the absence of preference heterogeneity, all individuals would have the same lifecycle profile for earnings and human capital. When we match differences in lifetime hours of work our model accounts for almost half of earnings dispersion for workers at age 50, and generates substantial heterogeneity in human capital at age 50. While the permanent component of hours heterogeneity is the most important component for hours variation across workers, the quantitative impact depends signficantly on capturing the autocorrelation function for hours is quantitatively important. In particular, what really matters is the permanent component of hours heterogeneity across workers.

The third step in our analysis enriches the previous model along multiple dimensions. In addition to heterogeneity in tastes for work, we allow for heterogeneity in learning ability and in initial human capital, as well as life-cycle shocks to human capital. We calibrate the parameters of this model so as to match a large set of moments that characterize the distribution of earnings and hours over the life cycle, and then use the calibrated model to assess the contribution of different factors to inequality in both hours and earnings. Two key results emerge. First, even in this model with many sources of heterogeneity, generating the amount of heterogeneity in cross-sectional and lifetime hours found in the data requires heterogeneity in preferences. Second, preference heterogeneity is an important source of heterogeneity in both the increase in earnings inequality over the life cycle and inequality in lifetime earnings. In particular, we find that heterogeneity in preferences accounts for more than 25 percent of the variation in the log of lifetime earnings as well as more than 25 percent of the increase in the variance of log earnings over the life cycle between the ages of 30 and 50.

Our model of hours heterogeneity and inequality has important implications for policy. First, our model suggests that human capital heterogeneity at the time of labor market entry is a less important source of lifetime earnings inequality than found in earlier work. Second, it implies that policies which directly impact the distribution of hours of work will potentially have important effects on both the mean level of earnings as well as dispersion in lifetime earnings. In particular, we use our model to evaluate the consequences of the regulation adopted by France that limits weekly work hours to be no greater than 48 hours. Our model implies that this policy reduces both lifetime inequality in earnings and mean earnings.

Our paper relates to the vast literature on inequality, but is most related to those of Huggett, Ventura and Yaron (2006, 2011) and Guvenen, Kuruscu and Ozkan (2014). Like us, the two papers by Huggett et al focus on the factors that can account for the evolution of inequality over the life cycle. The 2006 paper only allowed for heterogeneity in initial human capital and learning ability, while the 2011 paper also considered shocks to human capital. Importantly, their analyses imposed that the sum of time spent producing and investing was the same for all individuals and constant over the life cycle. Our main contribution relative to them is to introduce a labor-leisure choice and to assess the role of hours heterogeneity as an individual factor. Guvenen, Kuruscu and Ozkan (2014) also studies the life cycle pattern of inequality in a setting that allows for differences in initial human capital and learning ability and shocks to earnings. Their model does include an endogenous labor-leisure choice and they highlight how differences in tax progressivity across countries influences work incentives and human capital accumulation. Our main contribution relative to them is to allow for preference heterogeneity. Importantly, we show that without preference heterogeneity one cannot generate empirically reasonable differences in cross-sectional and lifetime hours.

An outline of the paper follows. In Section 2 we document properties of lifetime hours and earnings using the NLSY79. In Section 3 we show that existing quantitative versions of the Ben-Porath model imply that the magnitude of hours inequality found in the data have quantitatively large effects on inequality via their effect on inequality in human capital accumulation. Section 4 presents and Section 5 calibrates it. Section 6 presents our main results about the role of preference heterogeneity for inequality. Section 7 evaluates the separate role of endogenizing the labor-leisure choice versus adding preference heterogeneity. Section 8 concludes.

2 Empirical Facts on Hours Worked and Earnings Inequality

This section describes how we use the data from the NLSY79 to construct life cycle profiles for hours worked and earnings for a large sample of individuals. Although differences in annual hours worked across individuals display considerable persistence, they also exhibit mean reversion. It follows that cross-sectional dispersion in annual hours worked overstates the dispersion in lifetime hours worked. Nonetheless, we document a large amount of dispersion in lifetime hours of work. We also document a strong positive correlation between lifetime hours and lifecycle earnings growth.

2.1 The Data and Our "Lifetime" Sample

Our empirical analysis is based on the NLSY79, a longitudinal study of 12,686 individuals born between 1957 and 1964. Respondents were recruited and initially interviewed in 1979, when they were between 14 and 22 years old. They were re-interviewed annually until 1994, then biennially afterward. We use data through the 2018 interview, at which point individuals were between 53 and 61 years old. We neither include the econonomically disadvantaged whites (namely, non-Hispanic, non-Black) supplemental sample as it was discontinued in 1991 nor the supplemental military sample since the vast majority of those individuals was not interviewed after 1983.

The model that we study in the following sections is best understood as a model of individuals who are highly attached to the labor market. For this reason our empirical analysis will focus on men with at least a high school degree or equivalent. The excluded groups are much more likely to have extended periods out of the labor force even during prime working ages. Including these workers would likely increase the amount of dispersion in lifetime hours, and indicate an even larger role for hours variation than what we report. This leaves us with a sample of size of 4,412.

Our objective is to create hours and earnings profiles for individuals from age 25 to 52. We set the lower bound to an age at which most men have completed formal schooling; the upper bound is the oldest age that we can observe labor market information for every cohort. Because we are interested in documenting differences in lifetime hours, we further restrict our sample of men to those that we observe over a long time horizon. Our starting sample of men with at least a high school degree or equivalent not in the military oversample consists of 4,412 individuals. Our first restriction is to limit attention to individuals who we observe at least once from age 52 or later, indicating that they remained in the survey for a substantial period of time. This leaves us with 3,052 individuals. Second, we drop any individual who has more than one year of missing hours, where we define a missing year to be one with more than four weeks of missing employment information. This leaves us with 2,516 men. Finally, we drop any individual who has more than two years of missing earnings. (We do not include non-interview years after 1994 as missing for this criteria). This leaves us with a final sample size of 2,082 men. In what follows we will refer to this sample as our "lifetime sample".

Our analysis centers on two variables: annual hours worked and annual earnings. We construct our measure of annual hours as follows. In each interview, the NLSY79 asks individuals about all the jobs they held since their last interview. For each job, the NLSY79 records the weeks that the individual was employed in that job and the usual hours worked per week for that job. This information provides a weekly measure of hours worked, which we then sum over the calendar year to arrive at annual hours worked. Because this measure is based on the jobs that were held since the last interview, we observe hours even for years in which the individual is not interviewed (either because they missed a given interview, or because after 1994 the survey became biennial.)

In each interview an individual reports their total wages and salary and farm/business income in the previous calendar year. We use this information to create a time series for annual earnings. Unfortunately, this creates missing values for earnings once the survey becomes biennial in 1994. For these years t, we impute earnings as the average of the earnings in years $t-1$ and $t+1$. If earnings in year $t-1$ are missing, we use earnings from year $t-2$, and similarly for year $t+1$.

While the original sample of the NLSY79 was chosen to be nationally representative, it is possible that attrition could lead to a biased sample. In a companion paper, we conduct a host of validation checks for labor market data between our lifetime NLSY79 sample, the CPS-ASEC (March Supplement), and data from the Social Security Administration. These comparisons use similar age cohorts. Averaging across ages $25-52$, we find that the employment rate in the NLSY79 is

1.3 percent higher than in the CPS. Conditional on being employed in a given year, men in our NLSY79 sample work 0.7 percent more weeks per year than in the CPS, and work 5.0 percent more hours per week. Mean earnings in our pooled NLSY79 sample are 13.2 percent higher than in the CPS, with the largest gaps occurring after age 45. The 90/10 earnings ratio in our pooled NLSY79 sample is 5.6, as compared with 5.2 in the CPS. The higher level of inequality in the NLSY79 is primarily driven by top earnings: the 50/10 and 75/25 earnings ratios are very comparable between the NLSY79 and CPS. Based on these comparisons, we conclude that our sample is broadly similar to the CPS in terms of labor supply and earnings, though it does have somewhat higher levels of hours worked and top earnings.

2.2 Properties of Hours Worked

In this subsection we report several properties of the hours worked distribution for our lifetime sample.

We start by presenting the mean and standard deviation of annual hours worked over the life-cycle in Figure 1. We highlight two patterns. First, over the first ten years of the life cycle there is a substantial increase (13.3 log points) in the mean of log hours and a substantial drop (8.7 log points) in the standard deviation of log hours. Second, from the early thirties onward both profiles are remarkably stable.

The variation in annual hours of work is very large: the values for the standard deviation profile indicate that, for a given age and year, someone working one standard deviation above the mean works roughly 50% more hours than someone working one SD below the mean. Importantly, large cross-sectional differences in hours need not reflect large differences in lifetime hours across individuals. If hours differences are purely transitory over time, two individuals might differ in their hours in a particular year, but have similar average hours over their lifetimes. To investigate the persistence of hours differences over time, we now exploit the panel dimension of the NLSY79.

Figure 1: Cross-section of Hours Worked Over the Life-Cycle, Ages 25-52

(a) Mean Annual Hours (b) Standard Deviation of Annual Hours

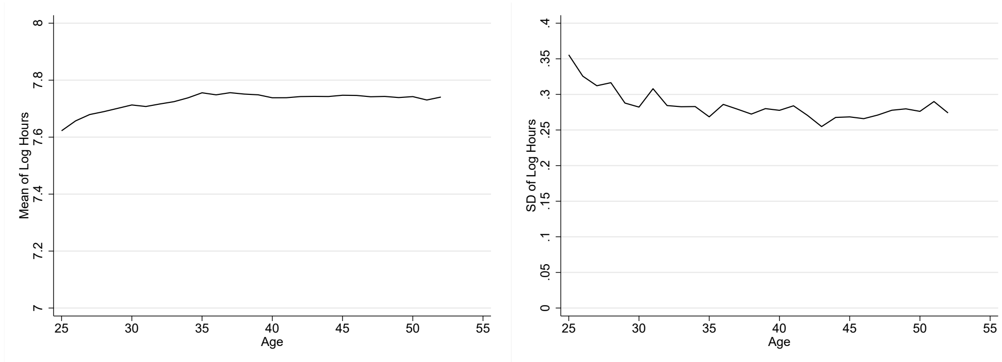

Notes: Moments are conditional on working at least 500 hours per year and an hourly wage of at least $2.

Our measure of persistence is the autocorrelation in annual hours worked. The long panel component present in our lifetime sample allows us to examine the autocorrelation of hours at long lags. To construct autocorrelations at lag length $t > 0$ we collect all length-t pairs of hours $\{(h_{i,a}, h_{i,a+t})\}_{a=25}^{52-t}$ from our sample of individual hours profiles $\{h_{i,a}\}_{a=25}^{52}$ and compute the pairwise correlation of this collection. Figure 2 plots the resulting correlation coefficients for lag lengths $t = 1, ..., 20$.

The key message from Figure 2 is that although there is some tendency for mean reversion, hours display considerable persistence over long horizons. The 1-year autocorrelation of hours is 0.681. The autocorrelation at a lag of five years has fallen by half to 0.341, and as the lag increases the autocorrelation continues to decline, but at a noticeably slower rate. By 20 years the correlation is still substantially above zero, at 0.176. The autocorrelations at one and twenty year horizons are not consistent with a simple AR(1) process. An AR(1) process with an autocorrelation of 0.681 at a lag of one year would have an autocorrelation of effectively zero at a twenty year lag.

Because hours differences persist over time, the differences observed in a given cross-section do not tend to cancel out when cumulated over a lifetime. To analyze these cumulative differences, we construct a measure of cumulative lifetime hours, where we add up annual hours in all years worked. We then divide these sums by 28 years (the number of years covered in our lifetime sample) to get an annualized measure.

The first column of Table 1 summarizes the distribution over annualized lifetime hours. On average, a man in our sample works 2,153.6 hours per year from ages 25-52. The standard deviation is 599.8 hours per year, which is 27.9% of the mean. The interquartile range is 515.1 hours per year, and the interquartile ratio is 1.26.

It is of interest to examine the extent to which the overall variation in lifetime hours stems

Figure 2: Autocorrelation of Hours Worked, Ages 25-52

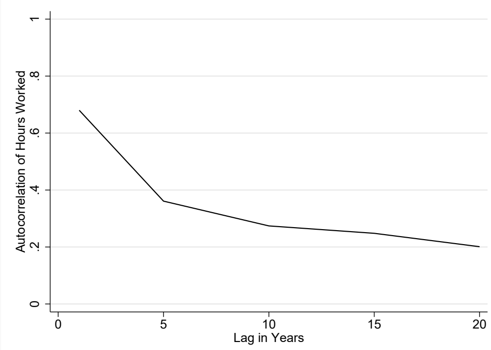

Notes: Moments are conditional on working at least 500 hours per year and an hourly wage of at least $2.

Table 1: The Distribution of Lifetime Hours and Components

Percentile	Annualized lifetime hours	Years worked	Weeks per working year	Hours per working week
5th	860.4	14.4	42.0	41.0
10th	1379.5	21.7	44.6	40.6
25th	1957.9	27.2	49.1	41.1
50th	2196.7	27.2	49.8	45.5
75th	2473.0	28.0	50.1	49.3
90th	2768.2	27.9	50.7	54.8
95th	2985.6	27.7	50.1	60.3
Mean	2153.6	26.0	48.8	46.5
Std	599.8	4.8	4.8	7.7
N		2082		

Notes: Individuals are sorted into percentiles by their lifetime hours and divide them by 28 years, the number of years an in individual is in our selected sample. For each percentile, we report the average of the individual at the percentile and of the 5 individuals just above and below. *Years worked*, *Weeks per working year*, and *Hours per working week* are the average values for all individuals in a given percentile. Note that the product of these three variables is slightly different than the average annualized lifetime hours for each percentile.

primarily from variation in years worked, weeks worked per working year, or hours worked per working week. The first two margins will largely reflect differences in labor market attachment and unemployment outcomes, while the latter is an indicator of work intensity conditional on being employed. The second through the fourth columns of Table 1 report the mean value of each of these margins. We emphasize two messages. First, for individuals in the lowest ten percent of the lifetime hours distribution, there is substantial variation in the "extensive" margins of years worked and weeks worked per working year. Second, above the tenth percentile of lifetime hours, there is little variation in the extensive margins. Instead, in this range almost all the variation is driven by hours worked per workweek. The fact that the intensive margin drives hours variation for the vast majority of individuals in our sample will rationalize our decision to focus on a model that represents the choice of annual hours as a choice along the intensive margin.

Figure 3: Cross-section of Annual Earnings Over the Life-Cycle, Ages 25-52

(a) Mean (b) Standard Deviation

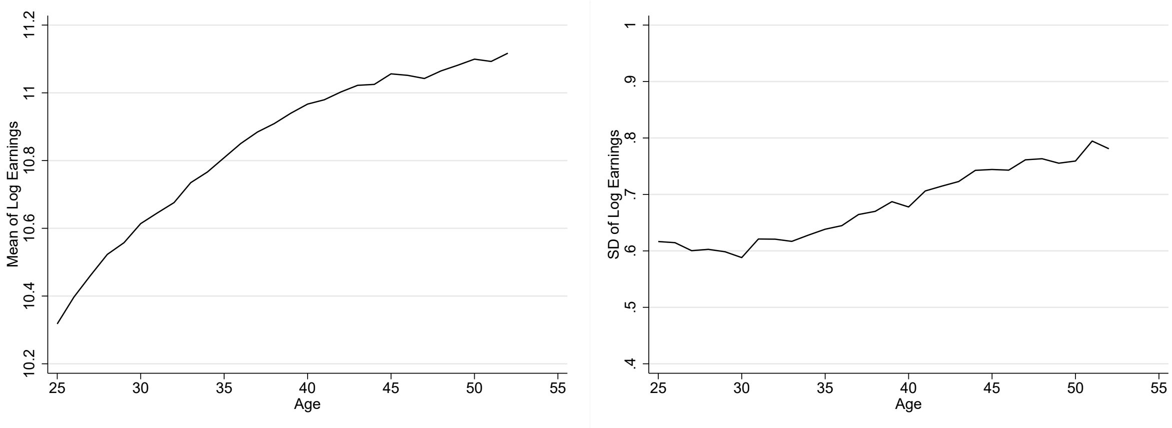

Notes: Moments are conditional on working at least 500 hours per year and an hourly wage of at least $2.

2.3 Properties of Earnings

In this section we study the properties of earnings in our lifetime sample. Whereas our presentation of the facts about lifetime hours is new to the literature, many researchers have documented the properties of earnings over the life cycle. Nonetheless, it is important for us to document the properties of earnings for our sample since we want our facts about hours and earnings to correspond to the same sample of individuals.

Figure 3 displays the mean and standard deviation of the log of annual earnings over the life-cycle. Mean earnings slightly more than double between the ages of 25 and 52, with a noticeable flattening after age 40. The standard deviation of log earnings increases from 0.62 at age 25 to 0.79 by age 52. These patterns are broadly consistent with analogous life-cycle profiles from the PSID (Huggett, Ventura and Yaron (2006, 2011)).

Next we establish that just as was the case for hours, earnings display a substantial degree of persistence over time in our dataset. We compute autocorrelations for earnings using the same method as we used for hours, though we only consider even numbered lag lengths given that the NLSY79 eventually becomes biennial and does not collect earnings information in non-survey years. Figure 4 plots the results. The autocorrelation of earnings is 0.849 at a two year lag, and declines to 0.456 at a twenty year lag. We conclude that earnings are even more persistent than hours. It is also true here that these two values are inconsistent with a simple AR(1) process, as an autocorrelation of 0.849 at a two year lag would imply a correlation of only around 0.20 at a twenty year lag.

Using tax data from the Social Security Administration, Kopczuk, Saez and Song (2010) doc-

Figure 4: Autocorrelation of Earnings, Ages 25-52

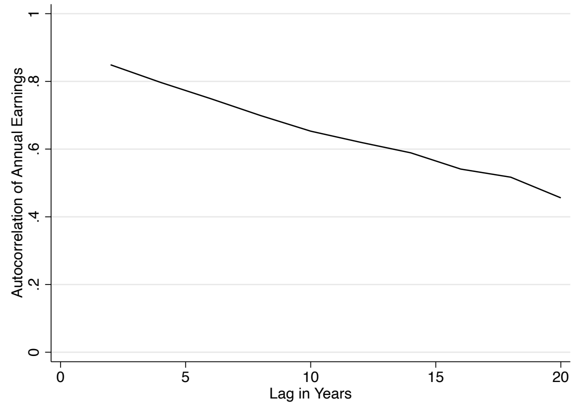

Notes: Moments are conditional on working at least 500 hours per year and an hourly wage of at least $2.

ument a one year autocorrelation coefficient for male earnings of 0.869 in the years 1978-1982 and 0.898 for the period 2000-2004. If one were to assume an AR(1) process, our two year autocorrelation estimate of 0.849 would imply a one year autocorrelation of $0.849^{1/2} = 0.921$, slightly above their estimate. But as we just noted, the AR(1) assumption does not seem warranted. While this differs from the estimate in Kopczuk, Saez and Song (2010), it is important to note that there are important differences in samples and time frames. We view our findings for persistence as being broadly similar to those found elsewhere in the literature.

We construct lifetime earnings for each individual by summing up individual earnings over the ages of 25 to 52. (Recall that we impute earnings for non-survey year using interpolation.) As with hours, persistent earnings differences imply sizable variation in average lifetime earnings. We find that on average, a man in our sample earns $\$60,429$ per year from ages $25-52$. The standard deviation is $\$61,021$ per year, which is roughly the same size as the mean. The standard deviation of log lifetime earnings is 0.74. The interquartile range is $\$44,897$ per year, and the interquartile ratio is 2.61.

Guvenen et al. (2022) use tax data from the Social Security Administration to document properties of the distribution of average lifetime earnings over ages $25-55$. For men born in 1958 (their most recent birth cohort), they report an interquartile earnings ratio of 2.75, compared with 2.61 in our data.[4] Once again, we view our results documented using our lifetime sample from the NLSY79 as broadly consistent with those documented with other samples.

Figure 5: The Correlation Between Lifetime Hours and Earnings, Ages 25-52

(a) Lifetime Earnings & Lifetime Hours

(b) Lifetime Earnings **Growth** & Lifetime Hours

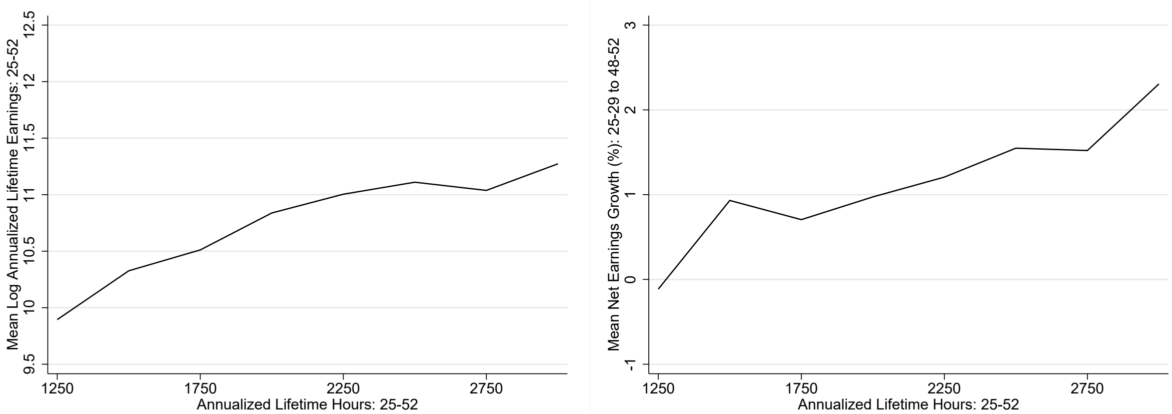

Notes: In (b), we restrict the sample to individuals with who work at least 500 hours and earn at least $ 2 per hour at each age 25-29 and 48-52.

2.4 Correlation Between Lifetime Hours and Earnings

In this subsection we document some aspects of the relationship between hours worked and earnings. We begin by examining the relationship between our measures of lifetime hours and lifetime earnings. We do not apply any discounting when computing our measures of lifetime hours or earnings, but note that we obtain similar results when we use a two percent annual discount rate for earnings.[5]

Figure 5a shows that individuals who work more hours over a lifetime earn substantially more. The positive correlation between lifetime hours and lifetime earnings is perhaps not surprising given that higher hours will increase earnings holding wages constant. More interesting is the fact that the percent change in earnings across hours bins is significantly larger than the percent change in hours. The horizontal axis consists of 250-hour bins for lifetime hours. Comparing the two datapoints near the edge of the interquartile hours range, the average earnings of individuals who work $2,250-2,499$ hours per year are 56 log points, or 74%, higher than the earnings of individuals who work $1,750-1,999$ hours. Averaging across the full range of $1,250-3,000$ hours, an increase of 500 hours per year is associated with an increase in average earnings of 65 log points (91%).

As discussed in the introduction, we are also interested in the relationship between lifetime hours and the growth rate of earnings over the life cycle. We construct a measure of life-cycle earnings growth as follows. We average an individual's earnings between ages $25-29$ to get a

[4]We note that these statistics do refer to different samples; our sample only includes males with at least a high school degree.

[5]Results for the case where we include a discount factor are reported in the Appendix.

measure of initial earnings, and average their earnings between ages $48 - 52$ to get a measure of final earnings. Earnings growth is then just the percent change between initial and final earnings.

Figure 5b plots average earnings growth separately by lifetime hours bins. Comparing the two datapoints near the edge of the interquartile hours range, the average earnings growth of individuals who work $2,250 - 2,499$ hours per year are 26 percentage points higher than the earnings growth of individuals who work $1,750 - 1,999$ hours. Averaging across the full range of $1,250 - 3,000$ hours, an increase of 500 hours per year is associated with an increase in average earnings growth of 40 percentage points.

3 Hours Heterogeneity and Human Capital Accumulation: A First Look

A key feature of inequality is that the cross-sectional dispersion of log earnings and wages increases substantially over the life cycle. One strand of literature uses the Ben-Porath model to understand this empirical regularity. This literature has focused on two factors that can generate an increasing dispersion of human capital over the life cycle: the accumulation of persistent shocks and heterogeneity in learning ability. In this section we argue that the literature has neglected a third quantitatively relevant factor. Specifically, we carry out a simple exercise to show that the large dispersion in lifetime hours documented in the previous section also generates a quantitatively important increase in the dispersion of human capital over the life cycle. Importantly, we show that this result holds when using the same parameterization of the Ben-Porath technology that the existing literature has used to establish the importance of the other factors.

In the first subsection we lay out the simple benchmark Ben-Porath model that is at the heart of recent work on life cycle models of inequality and highlight the mechanism through which it implies a positive relationship between total work hours and human capital accumulation holding other factors fixed. In the second subsection we provide a quantitative assessment of this effect. To do this we extend a benchmark Ben-Porath model to allow several sources of preference heterogeneity and calibrate the model so as to match the salient features of hours heterogeneity documented in the previous section. Our specification of preference heterogeneity features permanent and transitory components and is motivated by the properties of the autocorrelation of hours. The fact that the autocorrelation declines over time suggests a transitory component to preference heterogeneity, while the fact that the autocorrelation plateaus above zero suggests a permanent component. We show that this structure, combined with measurement error in hours, can match both the cross-sectional variance of log hours in the overall workforce as well as the autocorrelation of hours at the individual level.

Our quantitative analysis has two key results. First, we show that the heterogeneity in hours of work across individuals documented in the previous section generates substantial dispersion in human capital accumulation and earnings growth across otherwise identical individuals. In particular, our model generates a standard deviation of log earnings among 50 year olds that is almost fifty percent of its value in the data. More than half of this dispersion is accounted for by differences in human capital. We find that heterogeneity in hours accounts for more than twenty-five percent of the variance in the log of lifetime earnings.[6]

The second key result concerns the importance of the underlying process that generates dispersion in lifetime hours of work. Our specification features both transitory and permanent preference heterogeneity. In principle, one can generate a given dispersion in lifetime hours assuming only permanent preference heterogeneity or only transitory preference heterogeneity. We show that the effect of dispersion in lifetime hours varies dramatically across these two specifications. In particular, a given level of dispersion in lifetime hours induces much greater dispersion in human capital accumulation than when this dispersion is generated by permanent preference heterogeneity as opposed to transitory preference heterogeneity. This result is qualitatively intuitive: our theoretical analysis of the Ben-Porath model shows that high future hours influence current incentives to accumulate human capital, and expected differences in future hours are much larger when preference heterogeneity is permanent. This finding implies that it is important to model the details of variation in hours at the individual level.

3.1 A Simple Ben-Porath Model

The starting point for our analysis is a simple homogeneous agent version of the Ben-Porath model. A unit mass of identical individuals live for J periods. Each individual has preferences over streams of consumption (c_j) and total work hours (h_j) given by:

$$\sum_{j=1}^{J} \beta^t [\frac{\sigma}{\sigma-1} c_j^{1-\frac{1}{\sigma}} - \psi \frac{\gamma}{\gamma+1} h_j^{1+\frac{1}{\gamma}}]$$

where $0 < \beta < 1$ is a discount rate, $\psi > 0$ is a parameter controlling the disutility of work, and $\sigma > 0$ and $\gamma > 0$ are elasticity parameters.

Each individual is endowed with the same initial level of human capital x_0. At each age j the individual chooses how much time to devote to production (n_j) and how much time to devote to investing in human capital (s_j). Total work hours at age j are denoted by $h_j = n_j + s_j$. Time devoted to human capital investment at age j serves to augment the amount of human capital that

[6] We note that in all that follows we use the term "lifetime earnings" to refer to earnings between the ages of 25 and 52, which is the age range covered by our NLSY79 sample.

the individual will have at age $j+1$:

$$x_{j+1} = (1-\delta)x_j + \alpha \cdot (x_j s_j)^{\phi}.$$

The parameter $0 < \delta < 1$ reflects depreciation of human capital, and $\alpha \cdot (x_j s_j)^{\phi}$ reflects the production of new human capital. In particular, $x_j s_j$ reflects efficiency units of time devoted to human capital investment, ϕ governs the extent of decreasing returns in the production of human capital, and α reflects productivity of the human capital production function. In what follows we will refer to α as learning ability.

There is a competitive labor market that offers a time invariant wage rate of w per efficiency unit of labor services. If an individual of age j with human capital x_j devotes n_j units of time to production they will receive labor earnings e_j given by:

$$e_j = x_j n_j w.$$

Individuals can borrow and lend at the time invariant gross interest rate R. The only constraint on borrowing is that the individual cannot die with negative assets.

We assume that each individual exogenously retires at age J_R. We capture this with the constraint:

$$n_j = s_j = 0 \text{ if } j \geq J_R$$

Time allocation is at the core of this model: each period an individual must choose how much time to devote to production and how much time to devote to investment. Factors that generate heterogeneity in time devoted to investment are intuitively important in understanding the increasing dispersion in wages and earnings over the life cycle. Holding all else constant, heterogeneity in time devoted to investment will lead to heterogeneity in human capital accumulation. Because time invested in learning tends to be very small at older ages, cross-sectional dispersion in wages at older ages is closely related to cross-sectional dispersion in human capital at older ages. For this reason, dispersion in human capital accumulation is a potentially important source of dispersion in earnings and wages at older ages.

In the Appendix we show that the optimal allocation of time between producing and investing at each age j takes the following form:[7]

[7] Our discussion here assumes interior solutions purely to simplify the analysis.

$$wx_j = \alpha\phi x_j^{\phi} s_j^{\phi-1} \sum_{j'=j+1}^{J_R-1} [\frac{1}{1+R}]^{j'-j} wh_{j'}(1-\delta)^{j'-(j-1)} \quad (1)$$

This condition can be interpreted as requiring that the marginal value of time spent in production should equal the marginal value of time in investment. The left hand side reflects the value of a marginal increase in production time at age j. The marginal benefit of additional production time is just the wage per unit of production time, which is the product of the individual's human capital (x_j) and the wage per efficiency unit of labor services (w).

The right hand side reflects the value of a marginal increase in investment time at age j. 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 $\alpha\phi x_j^{\phi} s_j^{\phi-1}$ reflects the marginal increase in human capital at age $j+1$ as a result of a marginal increase in time devoted to investment. This investment will also increase human capital in period $j' > j+1$ by the amount $\alpha\phi x_j^{\phi} s_j^{\phi-1}(1-\delta)^{j'-j+1}$. The value of this additional human capital at age j' is the product of three terms: the increase in human capital at age j', the total number of hours worked at age j', and the wage per efficiency unit of labor services. An important point is that higher future hours worked increase the marginal benefit of additional investment today.

This equation highlights the key mechanism in the analysis of Huggett, Ventura and Yaron (2006). They used this framework to assess the role of heterogeneity in learning ability α across individuals for the evolution of the cross-sectional distribution of earnings over the life cycle. Heterogeneity in α will generate heterogeneity in human capital even if there is no heterogeneity in time devoted to investment. But Equation (1) shows that an increase in α holding all else constant raises the right hand side, thereby leading to an increase in s_j. The effect is intuitive: a higher value of α increases the return to time spent investing in human capital. This reinforces the direct effect of heterogeneity in α.

Equation (1) can also be used to explain why differences in total hours of work across individuals generate differences in human capital accumulation. Higher total hours imply that at least one of production or investment time must increase. The only way for human capital to not be affected would be if all of the additional hours were devoted to production. But looking at Equation (1), if all of the additional hours were devoted to production, the right hand side would increase, implying that the marginal benefit of investment increases, which is inconsistent with investment remaining unchanged. The key logic here is the point highlighted earlier: higher production hours in the future serve to increase the marginal value of additional investment today.

In the next subsection we quantify this effect. An important point that Equation (1) highlights is that the quantitative effect on time devoted to investment will likely depend heavily on the value of ϕ since this parameter dictates the extent of decreasing returns to scale in the investment technology. The greater the extent of decreasing returns, the faster the marginal benefit of investment declines in response to increases in time devoted to investment. Importantly, this is also true for assessing the effect of heterogeneity in learning ability. By using the same value of ϕ that has been used in the literature that studies heterogeneity in α, we ensure that the magnitude of the effects we find can be directly compared to those of the existing literature.

3.2 A Quantitative Exercise

In this subsection we use estimates of the Ben-Porath technology from Huggett, Ventura and Yaron (2011) to assess the quantitative effect of heterogeneity in lifetime hours on earnings inequality. As documented in Bick, Blandin and Rogerson (2022), observables explain very little of the variation in hours worked. Motivated by this we assume heterogeneity in preferences as the source of hours differences. As noted earlier, we allow for two sources of preference heterogeneity in order to match the autocorrelation properties of hours and so now consider preferences of the form:

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

where c_{ij} is consumption at age j, h_{ij} is total time devoted to work at age j, ψ_i is an individual specific time invariant preference shifter and $\pi_{i,j}$ is an idiosyncratic shock to preferences. We assume that ψ is distributed according to a log normal distribution with mean μ_ψ and standard deviation σ_ψ. The idiosyncratic shock $\pi_{i,j}$ is assumed to follow:

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

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

where the innovations $\nu_{i,j+1}$ are assumed to be iid over time and across individuals. We assume that initial values for the transitory shock, $\pi_{i,0}$, are drawn from the invariant distribution for the $\pi_{i,j}$ process. Because the draws for $\nu_{i,j}$ are iid we assume that the initial values $\pi_{i,0}$ are uncorrelated with the ψ_i. The parameters β, σ, and γ satisfy $0 < \beta < 1$, $\sigma > 0$ and $\gamma > 0$ and are common to all individuals. We also allow for classical measurement error in hours. Specifically, we assume that measured log hours are equal to true log hours plus an iid term that is log normally distributed with mean zero and standard deviation σ_{mh}.

Implementing our quantitative exercise requires that we assign values for 15 parameters: demographic parameters J_R and J, elasticity parameters σ and γ, the preference heterogeneity parameters μ_ψ, σ_ψ, ρ_π, and σ_π, measurement error parameter σ_{mh}, human capital investment parameters α and ϕ, the depreciation rate δ, initial human capital x_0, the discount rate β and the interest rate R.

As noted in the previous subsection, the returns to scale parameter ϕ in the human capital investment function is likely to be an important parameter. We set it equal to the value calibrated by Huggett, Ventura and Yaron (2011) so that we are using the same value that existing studies have relied on when using the Ben-Porath model to study earnings inequality.

Several other parameters are also set similar to Huggett, Ventura and Yaron (2011). In particular, we will interpret our model as representing individuals from age 25, with retirement occurring at age 65 ($J_R = 40$) and death at age 80 ($J = 55$). We set $R = 1.02$ and choose $\beta = 1/(1+R) = .98$. The two elasticity parameters are set according to $\sigma = 1$ and $\gamma = .3$. We set α equal to its mean value in Huggett, Ventura and Yaron (2011). The depreciation rate for human capital is set to 2.0% per year. The value of x_0 affects the average life cycle earnings profile but has little effect on the heterogeneity across profiles. We set its value equal to the mean value in Huggett, Ventura and Yaron (2011).

There are five remaining parameters: μ_ψ, σ_ψ, ρ_π, σ_π and σ_{mh}. The parameter μ_ψ serves to shift mean hours. Huggett, Ventura and Yaron (2011) assume that total time spent producing and investing is the same for all individuals and normalized this value to equal unity, so we choose μ_ψ so that the mean of lifetime hours is equal to one. We choose the remaining four parameters to target four moments: the cross-sectional variance of log hours in the overall population between the ages of 25 and 52, and the values of the autocorrelation function at lags of one, ten and twenty years. When solving the model we approximate the log normal distribution for ψ_i with a discrete distribution over 8 grid points.

An important issue when connecting a Ben-Porath model to the data concerns the mapping between hours of work in the model and hours of work in the data. The issue is the extent to which hours of work as reported in the data necessarily include all hours spent accumulating human capital. The standard convention in the literature is to count time spent in investment as reported work hours for individuals with positive production time, and to count time spent in investment as education time for individuals with zero time spent in production.[8] We adopt this standard 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 work time in the data, total working time is underestimated.[9]

Table 2 summarizes the parameter values. We list along with each parameter the moment most closely affected by the parameter with the understanding that each parameter affects each moment. We note that this procedure generates a relatively high persistence for the transitory component π_{it}, with an AR(1) coefficient of 0.88.

[8] See, for example, Manuelli, Seshardi and Shin (2012). Guvenen, Kuruscu and Ozkan (2014) added a constraint that limited the total amount of time that could be devoted to training when production time is positive.

[9] We have experimented with other specifications, allowing for some fraction of investment time to be not 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 2: Jointly Calibrated Parameter Values

Parameter	Interpretation	Value	Moment / Source
J_R	Retirement age	65	–
J	Final age	80	–
R	Gross interest rate	1.02	Huggett, Ventura and Yaron (2011)
β	Patience	0.9804	Huggett, Ventura and Yaron (2011)
x_0	Initial human capital	105	Huggett, Ventura and Yaron (2011)
α	Learning ability	0.33	Huggett, Ventura and Yaron (2011)
ϕ	HC elasticity wrt investment	0.7	Huggett, Ventura and Yaron (2011)
δ	Human capital depreciation	0.02	Huggett, Ventura and Yaron (2011)
σ	CRRA	1.0	Balanced growth
γ	Frisch elasticity	0.3	–
μ_ψ	Mean of log ψ	1.001	Mean lifetime hours, age 25-52
σ_ψ	SD of log ψ	0.691	SD lifetime hours, age 25-52
σ_π	SD of log π	0.47	Hours autocorrelation profile
ρ_π	Autocorrelation of log π	0.88	Hours autocorrelation profile
σ_{mh}	SD measurement error	0.12	Hours autocorrelation profile

Absent preference heterogeneity, our model would imply that all individuals have the same lifecycle profile for hours, human capital and earnings. Our main objective is to assess the extent to which the preference heterogeneity that we add to the model generates inequality in lifetime earnings. We find that the standard deviation of the log of lifetime earnings is equal to 0.19. In our NLSY79 sample the corresponding value is 0.74, so our model with only preference heterogeneity accounts for more than twenty five percent of overall inequality in lifetime earnings.

It is also of interest to document the divergence in life cycle earnings, and in particular the extent to which these are driven by differences in human capital accumulation. To do this we focus on the dispersion in earnings and human capital at older ages. At age 50, the standard deviation of log earnings in our economy is 0.35. In our NLSY79 sample of males the standard deviation of log earnings at age 50 is 0.76. We conclude that taking existing estimates of the Ben-Porath investment technology as given, heterogeneity in lifetime hours of work is an important contributor to inequality in earnings.

The dispersion in log earnings in our model reflects dispersion in both wages and hours. In a textbook static model of labor supply, dispersion in log hours across otherwise identical individuals will create an identical amount of dispersion in log earnings. The novel feature of our analysis concerns the extent to which heterogeneity in hours leads to heterogeneity in human capital accumulation and hence wages. For this reason it is of interest to understand the relative importance of dispersion in wages and hours in generating the dispersion in log earnings in our model. At age 50 the standard deviation of log human capital is 0.12, and the standard deviation of log hours is 0.23. In our economy with only preference heterogeneity, hours and human capital at age 50 are highly correlated, so the two dimensions reinforce each other in terms of generating inequality in earnings. More generally this need not be the case, and our more general analysis in the next section will allow for this. The key message from this exercise is that differences in lifetime hours of the magnitude found in the data have quantitatively important implications for inequality via their effect on wages. These wage effects are the result of how differences in the desire to work generate differences in time devoted to investment.

3.3 The Importance of the Autocorrelation in Hours

The nature of preference heterogeneity adopted in the previous subsection was motivated by features of the autocorrelation of hours at various lags. As noted earlier, if one was only interested in generating a given level of dispersion in lifetime hours without matching features of the autocorrelation of hours, one could do this by assuming preference heterogeneity to be either purely permanent or purely transitory. In this subsection we show that the implications for inequality are dramatically different across these two specifications. It follows that it is important to model the nature of hours heterogeneity at a more micro level.

In our first specification we set $\rho_\pi = 1$, $\sigma_\pi = 0$ and σ_{mh} equal to its value in our previous calibration. As before we set μ_ψ so that mean hours are equal to one, and we set σ_ψ so that the variance of log lifetime hours is the same value as implied by our previous calibration. In the second specification we set $\sigma_\psi = 0$, set ρ_π and σ_{mh} equal to their values from the previous calibration, choose μ_ψ so that mean hours are equal to unity and choose σ_π so that the variance of log lifetime hours is the same value as implied by our previous calibration. By construction, all three of our economies have the same variance for the log of lifetime hours, and the same measurement error. Where they differ is in the extent to which the preference heterogeneity that drives hours heterogeneity is permanent versus transitory.

Here we focus on the implications for human capital accumulation. As before, absent any heterogeneity in preferences, all individuals would have the same human capital at age 50. Comparing the standard deviation of log human capital at age 50 across specifications is thus a measure of the inequality induced by differences in hours. In our benchmark specification the standard deviation of log human capital at age 50 was equal to 0.119. In the model with only permanent heterogeneity this value is equal to 0.167, while in the model with only transitory heterogeneity this value is equal to 0.015. The message is very stark: although transitory preference heterogeneity can generate large differences in lifetime hours, even with a substantial amount of persistence (recall that $\rho_\pi = 0.88$) the impact of this heterogeneity on human capital accumulation is of second order importance.[10] What really matters is the amount of variation in the permanent component of preference heterogeneity. If one assumed that permanent heterogeneity was the only source of differential hours, the effects on dispersion in human capital would be roughly fifty percent larger than in our baseline calibration.

As noted earlier, it is intuitive that transitory preference differences are less powerful in generating inequality in human capital accumulation. This exercise shows that this effect is very significant quantitatively.

3.4 Summary

Existing analyses of inequality in the context of Ben-Porath models have abstracted from the potential role of heterogeneity in lifetime hours. In this section we have argued that this abstraction reflects a substantively important omission. In particular, using the same investment technology specifications as existing analyses, we show that the dispersion in lifetime hours documented in the previous section has a large impact on earnings inequality via its effect on dispersion in human capital levels.

The quantitative analysis in this section abstracted from several features. In particular, the only source of hours heterogeneity is preference heterogeneity. Other sources of heterogeneity may also produce heterogeneity in hours and it is of interest to understand how much of the variation in lifetime hours is due to preference heterogeneity versus the interaction between these other sources of heterogeneity and the endogenous labor supply margin. Third, we have abstracted from tax and transfer systems, which may also moderate the effects that we document. The remainder of the paper is devoted to addressing these issues.

4 A Structural Model of Lifetime Hours and Earnings

In this section we develop a generalization of the model studied in the previous section. We frame our model as extending the heterogeneous agent Ben-Porath model in Huggett, Ventura and Yaron (2011), though modulo a few details it can also be viewed as extending the model of Guvenen, Kuruscu and Ozkan (2014). In addition to the three sources of heterogeneity in Huggett, Ventura and Yaron (2011) (heterogeneity in initial human capital endowments, heterogeneity in learning ability, and idiosyncratic shocks to the evolution of human capital over time) our model features

[10] We have repeated this exercise assuming higher values for ρ_π. Even when this value gets close to unity we find that the effect on human capital dispersion is still very small. For example, repeating the exercise with $\rho_\pi = .98$ yields a standard deviation of log human capital at age 50 of only 0.07.

two dimensions of preference heterogeneity that allow us to account for the salient features of heterogeneity in hours worked over the life cycle. We also include a progressive tax system.

4.1 Households

Time is discrete and at each date a new generation of unit mass is born. Each new generation is identical, so in what follows we focus on a generic generation and abstract from calendar time. Each individual lives from age $j = 1$ to $j = J$, and retires exogenously at age $j = J_R$. Individuals are endowed with one unit of time in each period. Lifetime utility of individual i is the same as in Section 3:

$$E_1 \sum_{j=1}^{J} \beta^t [\frac{c_{i,j}^{1-1/\sigma}}{1-1/\sigma} - \psi_i \pi_{j,i} \frac{h_{i,j}^{1+1/\gamma}}{1+1/\gamma}]$$

As in the previous section, total hours of work at age j for individual i, $h_{i,j}$, are allocated between producing ($n_{i,j}$) and investing in human capital ($s_{i,t}$). But following Huggett, Ventura and Yaron (2011), the human capital accumulation process from the previous section is generalized along two dimensions. Letting $x_{i,j}$ denote the human capital of individual i at age j, human capital for this individual at age $j+1$ is now given by:

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

where $s_{i,j}$ is time spent investing in human capital in at age j, α_i is an individual specific learning ability, and z_{it} is a log normally distributed shock to human capital:

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

This shock is iid over time and across individuals. Note that while the shock $z_{i,j}$ is purely transitory, its effect is persistent because it affects the individual's stock of human capital. The parameters δ and ϕ are as before.

Each individual is characterized by two fixed effects (their permanent disutility of work parameter ψ_i and their learning ability α_i) and two initial conditions (their initial value of the idiosyncratic preference shock $\pi_{i,0}$ and their initial human capital $x_{i,0}$). In our quantitative analysis we assume that $x_{i,0}$ and ψ_i are joint lognormally distributed:

$$\log(x_{i,0}, \psi_i) \sim N(\mu_x, \mu_\psi, \sigma_x, \sigma_\psi, \rho_{x,\psi})$$

One noteworthy effect of the correlation parameter $\rho_{x,\psi}$ is the potential for human capital profiles of agents to cross each other, which will in turn influence the behavior of the variance of earnings early in the life cycle.

To economize on the size of our state space we assume that α_i takes on only eight values, and that an individual's value for $x_{i,0}$ is perfectly correlated with α_i.[11] We assume that the initial distribution for $\pi_{i,0}$ is the ergodic distribution for the π_j process. Because the idiosyncratic shocks for this process are uncorrelated with all other variables, we assume that $\pi_{i,0}$ is uncorrelated with all other variables.

The government levies a linear tax on consumption, τ_c, and a progressive tax on labor income. We assume the labor income tax function is of the form proposed by Benabou (2002) and Heathcote, Storesletten and Violante (2014): 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. Each period the government returns tax revenues to households in the form of an equal lump sum transfer T, with the size of the transfer dictated by budget balance for the government.

We will focus on a situation in which the aggregate state of the economy is stationary, so that the wage per efficiency unit of labor and the interest rate are both constant over time, denoted by w and r respectively. It follows that an individual with human capital x that supplies n units of time to production will earn labor income equal to xnw. An individual of age j who is not retired and has human capital x_j and capital k_j will thus face a period budget equation of:

$$(1+\tau_j)c_j + k_{j+1} = \tau_0 \cdot (wx_j n_j)^{1-\tau_1} + (1+r)k_j + T$$

An individual of retirement age with capital k_j will face a period budget equation of:

$$c_j + k_{j+1} = T + (1+r)k_j$$

We impose the natural borrowing constraint, so the only restriction we impose on borrowing is that individuals cannot die with negative wealth

[11] Huggett, Ventura and Yaron (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. For this reason our assumption does not imply that "true" initial values for these two variables are perfectly correlated.

4.2 The Individual's Problem

At the start of a period, an individual has a six-dimensional state $e = (j,k,x,\alpha,\psi,\pi)$ consisting of age (j), physical assets (k), human capital (x), learning ability (α), permanent work disutility (ψ), and persistent work disutility (π). Taking as given the individual state, the wage rate w, the gross interest rate R, and government policies, a working age individual, $j < J_R$, solves the following recursive problem:

$$V(j,k,x,\alpha,\psi,\pi) = \max_{c,k',s,n} \quad \frac{c^{1-1/\sigma}}{1-1/\sigma} - \psi\pi\frac{(n+s)^{1+1/\gamma}}{1+1/\gamma} + \beta\mathbb{E}_{z',\pi'}V(j+1,k',x',a,\psi,\pi') \tag{4}$$

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

$$x' = x_{i,j+1} = z'[(1-\delta)x + \alpha(sx)^{\phi}] \tag{6}$$

Equation (5) is the individual's budget constraint, which states that after-tax expenditures on consumption and capital must equal the sum of income from physical assets rk, after-tax labor income, and the government transfer T.

A retired individual, $j \geq J_R$, solves an identical problem except with the added constraint that $n = s = 0$. An individual in their last period of life, $j = J$, faces an additional nonnegative savings constraint: $k' \geq 0$. That is, the only borrowing constraint that we impose is the natural borrowing constraint.

4.3 Stationary Equilibrium

We close the model broadly following Huggett, Ventura and Yaron (2011). As noted earlier, we assume that in each period there is a new generation of unit mass born, and that each generation is identical. We assume a Cobb-Douglas aggregate production function $F(K_t,X_t) = K_t^{\theta}X_t^{1-\theta}$ with constant returns in aggregate capital K_t and labor efficiency units X_t. We assume an open economy with an exogenous global interest rate. Given that the aggregate production function displays constant returns to scale, this implicitly fixes the wage per efficiency unit of labor.[12] Given this assumption, the only endogenous equilibrium object that enters into the life-cycle problem above is the government transfer T.

Define Λ to be a measure over individual states $e = (j,k,x,\alpha,\psi,\pi)$. We then have the following definition for a stationary equilibrium:

Definition: Given a world interest rate r, a stationary equilibrium is a wage rate w, transfers

[12] Starting from the first order conditions for capital and labor, it is easy to show that $w = (1-\theta)(\theta/r)^{\theta/(1-\theta)}$.

T, a value function $V(e)$ with associated decision rules $\{c(e), k'(e), n(e), s(e)\}$, and a distribution over individual states Λ such that:

1. Taking as given w, r and T, the value function satisfies the Bellman equations and decision rules are optimal.
2. Competitive factor prices: $w = (1-\theta)\left(\frac{\theta}{r}\right)^{\frac{\theta}{1-\theta}}$
3. Government budget balances.

4.4 Measurement

In this subsection we discuss two issues related to mapping hours in the model to those in the data. The first issue concerns measurement error. A large volume of work has documented substantial measurement error in both hours and earnings. The extent and nature of measurement error can have important effects on some of the variables that we highlight, such as the correlation between hours and wages, and the autocorrelation of hours. For this reason we will assume the presence of measurement error in both hours and earnings. In both cases we assume classical measurement error, with standard deviations of σ_{me} and σ_{mh}.

The second issue was mentioned earlier in Section 3 and concerns the extent to which hours spent accumulating human capital in the model should be considered hours of work as reported in the data. We continue to follow the standard convention in the literature and count time spent in investment as work hours for individuals with positive production time, and to count time spent in investment as education time for individuals with zero time spent in production.

5 Calibration

In this section we calibrate our model to match salient features of earnings and hours over the life cycle.

5.1 Calibration Procedure

Our model features 20 parameters: one price (r), three common preference parameters (β, σ, and γ), three technology parameters for the human capital accumulation process (δ, ϕ, and σ_z), two measurement error parameters (σ_{me} and σ_{mh}), two parameters characterizing the distribution of learning ability (μ_α and σ_α), two parameters characterizing transitory preference heterogeneity (ρ_π and σ_π), two parameters characterizing the tax function (τ_0 and τ_1) and five parameters characterizing the joint distribution of x_0 and ψ (μ_x, μ_ψ, σ_x, σ_ψ, and $\rho_{x,\psi}$). Without loss of generality we normalize μ_x to equal unity, leaving us with 19 parameters.

Four of these parameters (r, β, σ, and γ) appear frequently in quantitative models and for these we follow the literature and assign standard values. Setting the period length equal to a year, we set $r = 1.02$, $\beta = 1/r = .9804$, $\sigma = 1$, and $\gamma = 0.3$.

Five parameter values are assigned based on estimates in the existing literature. In particular, following Huggett, Ventura and Yaron (2011) we set $\sigma_z = 0.11$ and $\sigma_{me} = 0.15$, and following Guvenen, Kuruscu and Ozkan (2014) we set $\delta = 0.015$. Parameters of the tax function are set according to the estimates in Heathcote, Storesletten and Violante (2017), with $\tau_0 = 0.81$ and $\tau_1 = 0.181$.

Because our model features an endogenous margin for total hours of work, the effect of z shocks on earnings dispersion may differ from what Huggett, Ventura and Yaron (2011) found. They found z shocks to be an important source of inequality in lifetime earnings. Because we use their value σ_z, it is likely that we will also find z shocks to be an important source of variation in lifetime earnings. Our richer model is most likely to impact the the relative importance of other factors, such as heterogeneity in learning ability.

For the remaining ten parameters we choose values using a simulated method of moments procedure. We first list the ten moments that we target and later provide some intuition about the connection between model parameters and these moments.

Dispersion in earnings and hours are the core outcomes of interest in our analysis, so accordingly we target ten moments related to these outcomes. We target three moments related to earnings: the standard deviation of log earnings at 30 and 50, and the increase in mean log earnings that occurs between these ages.[13] We target 5 moments related to hours: the overall mean and standard deviation of hours for the pooled sample of all individuals between the ages of 25 and 52, and values for the autocorrelation of individual hours at lags of 1, 10 and 20 years, as depicted in Figure 2. We also target the correlation between hours and earnings at age 25.

We also want our model to be consistent with evidence that relates lifetime hours of work to life cycle earnings growth. For this reason we target one moment that jointly relates earnings and hours: the slope of the relationship between lifetime hours and life cycle earnings growth, as depicted in Figure 5b.

[13] Recall that we normalize the mean of initial human capital to equal one. For this reason the units of earnings at age 25 have no particular meaning, which is why we only target the increase in mean log earnings.

In what follows we provide some intuition about the connection between these moments and the ten parameters that we are calibrating. This discussion should be understood as purely heuristic, since all ten parameters influence all ten moments. Nonetheless, we think it is useful for understanding 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}. This mirrors the calibration exercise that we carried out in Section 3. Intuitively, the value of μ_ψ directly influences the cross-sectional mean of hours. Each of the other four will affect the cross-sectional standard deviation of hours and the shape of the autocorrelation function. In particular, the standard deviation of measurement error is intimately related to the autocorrelation of hours at one lag.

The parameter ϕ has a direct influence on the relationship between lifetime hours and life cycle wage growth, so targeting the empirical slope of this relationship helps to pin down the value for ϕ. It is noteworthy that we target a low frequency statistic rather than a high frequency statistic. 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 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, Blandin and Rogerson (2022).)

The parameters μ_a and σ_a will directly influence mean life cycle earnings growth and the extent to which earnings become more dispersed with age. For this reason targeting mean life cycle earnings growth and the dispersion in earnings at age 50 help to identify these two parameters. The parameter σ_x has a direct effect on the dispersion in earnings for young workers and so the standard deviation of earnings at age 30 is related to this value.

The final parameter is $\rho_{x,\psi}$ and the final moment is the correlation of earnings and hours for workers aged 25. This parameter is important for two reasons. First, as previously noted, it will influence the extent to which heterogeneous preferences serve to partially offset differences in initial human capital. The value of $\rho_{x,\psi}$ affects the relationship between initial human capital and time spent in investment, which will impact the correlation between wages and hours in the cross-section for young workers. For this reason the correlation between hours and wages for young workers helps to determine this parameter. This parameter also controls the correlation between learning ability and tastes for work. This is significant for interpreting the slope of the relationship between lifetime hours and lifecycle earnings growth, which is one of the moments that we target. This slope is not an estimate of the causal effect of lifetime hours on lifecycle wage growth: if learning ability and lifetime hours are highly correlated, then this moment necessarily reflects the effects due to heterogeneity in learning ability.

Table 3: Independently Calibrated Parameter Values

Parameter	Interpretation	Value	Source
R	Gross interest rate	1.02	Huggett et al. (2011)
β	Patience	0.98	$1/\beta$
σ	CRRA	1.00	—
γ	Frisch elasticity	0.30	—
δ	Human capital depreciation	0.015	Guvenen et al. (2014)
σ_z	SD human capital shock	0.11	Huggett et al. (2011)
σ_{me}	SD measurement error	0.10	Huggett et al. (2011)
τ_0	Labor income tax, level	0.81	Heathcote et al. (2017)
τ_1	Labor income tax, progressivity	0.181	Heathcote et al. (2017)
μ_x	Mean of $\log x_0$	0.00	Normalization

5.2 Calibrated Parameter Values

Tables 3 and 4 display the calibrated parameter values. We note several features of the estimates. First, our calibrated value of ϕ is 0.67. Notably, this value is very much in the range of estimates in the literature despite the fact that we use a novel moment to help identify it.[14] It is perhaps reassuring that different procedures imply similar values. We find a strong negative correlation between permanent tastes for work and initial human capital, with $\rho_{x,\psi}$ equal to -0.85. Recalling that we impose a deterministic positive relationship between learning ability α and initial human capital, our calibration implies that individuals who begin with high human capital at age 25 also have high learning ability and have less distaste for work. Recalling that we start our model at age 25, this pattern seems desirable, as individuals who work hard and have higher learning ability would be expected to have high human capital at age 25.

The parameters that determine preference heterogeneity are very similar to those that we found for the quantitative exercise in Section 3. The main difference is a slightly smaller value for the standard deviation of the permanent component. The standard deviation of $\log \pi$ in the ergodic distribution of thie transitory process is equal to 0.97, which is roughly three times as large as the dispersion in the permanent component ψ. We also find a substantial amount of measurement error in hours, with $\sigma_{mh} = 0.12$.

The standard deviation of $\log \alpha$ is estimated to be 0.06. This value is relatively small compared to values estimated in Huggett, Ventura and Yaron (2006, 2011) and is related to the fact that

[14] 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, Ventura and Yaron (2006). Both Huggett, Ventura and Yaron (2011) and Guvenen, Kuruscu and Ozkan (2014) use a value for ϕ that is quite close to our calibrated value.

Table 4: Jointly Calibrated Parameter Values

Parameter	Interpretation	Value	Moment
μ_α	Mean of $\log\alpha$	-2.26	Mean earnings, age 50
σ_x	SD of $\log x_0$	0.29	SD earnings, age 30
μ_ψ	Mean of $\log\psi$	3.85	Mean annual hours, age 25-52
σ_ψ	SD of $\log\psi$	0.38	SD annual hours, age 25-52
σ_α	SD of $\log\alpha$	0.06	SD earnings, age 50
σ_π	SD of $\log\pi$	0.46	Hours autocorrelation profile
ρ_π	Autocorrelation of $\log\pi$	0.88	Hours autocorrelation profile
σ_{mh}	SD measurement error	0.12	Hours autocorrelation profile
ϕ	HC elasticity wrt investment	0.67	Lifetime hours, earnings growth
$\rho_{x,\psi}$	Corr. of $(\log x_0, \log\psi)$	-0.85	Correlation of (hours, earnings), age 25

Figure 6: Model Fit: Mean and Standard Deviation of log Earnings Over the Life-Cycle

(a) Mean

(b) Standard Deviation

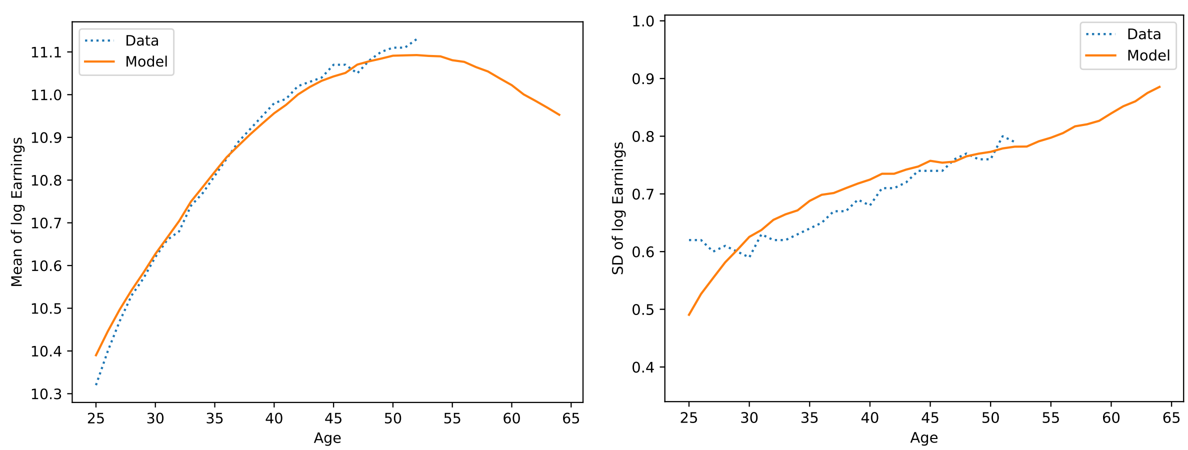

dispersion in lifetime hours in our model has similar effects to dispersion in learning ability.

5.3 Fit of the Model

In this subsection we report on the ability of the model to match both targeted and untargeted moments in the data. We begin by examining properties of earnings. Figure 6 shows the model's ability to match the behavior of the mean and standard deviation of earnings over the life cycle.

The growth in mean earnings between age 30 and 50 was explicitly targeted in the calibration procedure, and Panel (a) shows that the model does a good job of capturing the entire profile

Figure 7: Model Fit: Mean and Standard Deviation of log Hours Over the Life-Cycle

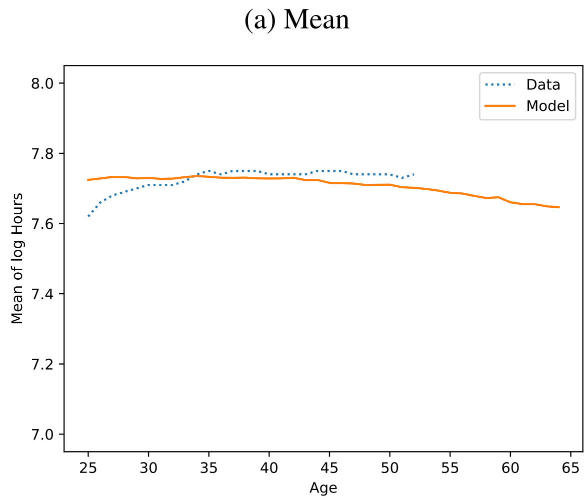

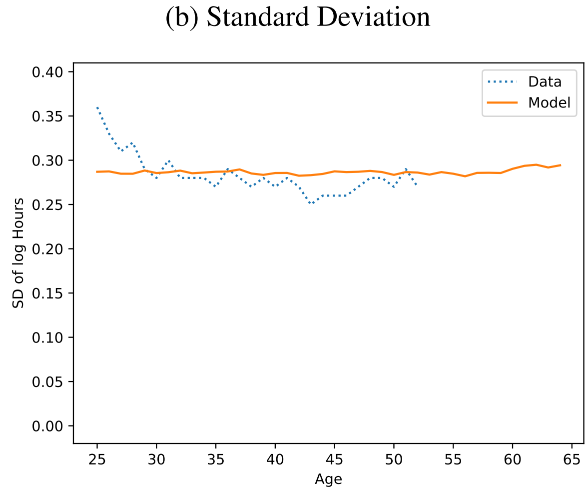

between these two ages. Similarly, the standard deviation of earnings at ages 30 and 50 were also targeted, and panel (b) shows that the model also does a good job of tracking the profile between these two ages, though it does not capture the properties of the profile between 25 and 30. The profile is essentially flat in the data over this range while the model predicts a modest increase.[15]

Figure 7 shows the profiles for mean hours and the standard deviation of hours over the life cycle. Recall that our calibration procedure targeted the overall cross-sectional mean and dispersion in hours worked, 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 necessarily implies a reasonable fit to the life cycle profiles. Here again we see some discrepancy between model and data over the 25-30 age range. In particular, while both profiles are also relatively flat in the model over the 25-30 age range, this is not the case in the data. In the data, mean hours are lower and the standard deviation of log hours is higher for 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 standard deviation. 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, and our model abstracts from this element of life cycle labor supply. The second is our assumption that the data reflects both production and learning time. This age range is the period of highest investment in human capital, so that if some of this is not included in work hours in the data our model would be expected to overestimate work hours at young ages.[16]

[15] In principle, a model like ours can generate a flat profile for the standard deviation in earnings in the early part of the life cycle. One possibility is that individuals with low initial human capital have high growth of human capital. Another possibility is that individuals with high initial human capital spend sufficiently more time investing that they have lower initial earnings. But as the figure shows, these effects are not sufficiently powerful in our calibrated model.

[16] If we assume that a fraction of time devoted to investment is not included in reported total work hours then we are able to match these patterns in the data more closely.

Figure 8: Model Fit: Distribution of Lifetime Earnings and Hours

(a) Hours (b) Earnings

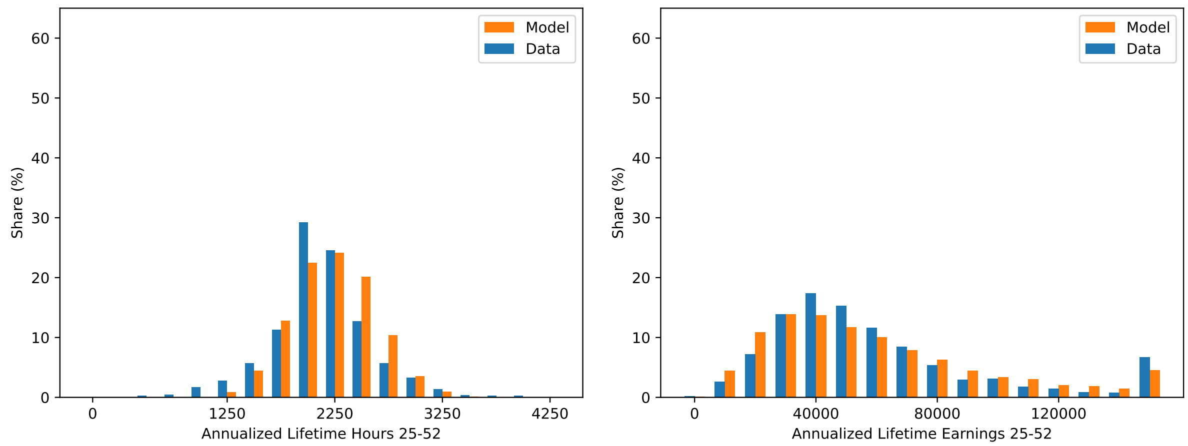

Figure 8 displays the distribution of (annualized) lifetime earnings and lifetime hours in both the model and the data. While our calibration matches the targeted values for the mean and standard deviation for annual hours in the overall cross-section, Panel (a) shows that our model does not generate sufficient concentration in the annualized lifetime hours bin that contains 2000 hours. This bin contains workers who report working the full year at 40 hours a week. While the empirical concentration in this bin is not as strong for annualized lifetime hours as it is for annual hours in a particular year, it is still quite concentrated relative to what one obtains with a normal distribution, which is approximately what our model generates given our assumption of log normal distributions for heterogeneity. While we could add some non-linearities to our model to better capture this concentration, we have opted not to do so in order to better focus on other features. Panel (b) shows that we do a better job of matching the earnings distribution. It remains the case that the model does not generate as much concentration as is found in the data, but the extent of the discrepancy is much more modest compared to the case of hours. This is intuitive; because there is substantial heterogeneity in wages for workers in the hours bin containing 2000, the concentration in hours implies much less concentration in earnings.

The two panels of Figure 9 feature two key relationships that we target in our calibration. Panel (a) shows that our model does a good job of tracking the autocorrelation properties in annual hours at the individual level. Panel (b) shows the relationship between lifetime hours and lifecycle earnings growth. While the profile in the model lies slightly above the profile in the data, the slope of the model profile does a good job of tracking the slope of the empirical profile.

Figure 10 studies the lifetime hours and earnings relationship in a bit more detail. Panel (a) shows the relationship between lifetime earnings and lifetime hours, which was not explicitly targeted. Consistent with the previous figure, we see that the model and data profiles track each other quite closely until the highest hours bins. The discrepancy arises because the model based

Figure 9: Model Fit: Additional Targeted Moments

(a) Persistence of Hours Worked

(b) Lifetime Hours and Earnings Growth

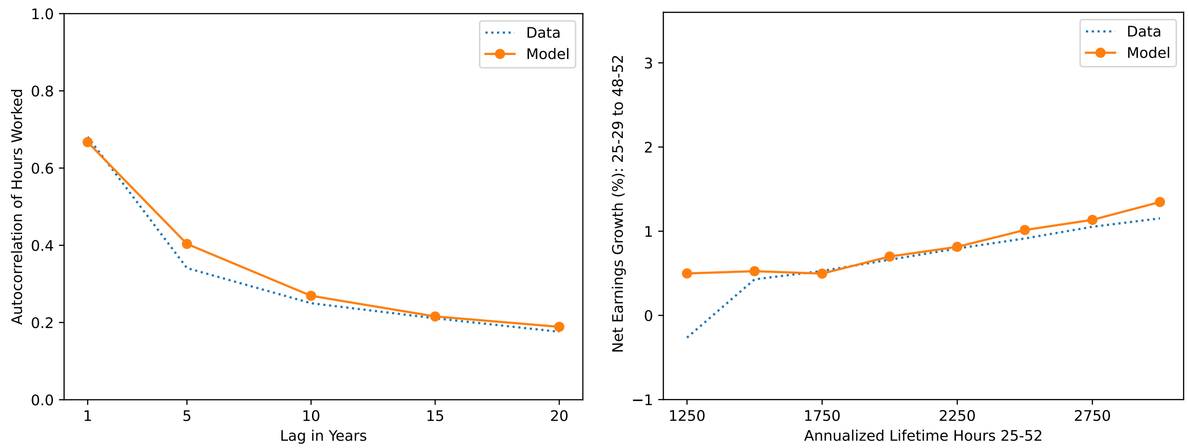

Figure 10: Model Fit: Untargeted Moments

(a) Lifetime Hours and Earnings

(b) Persistence of Earnings

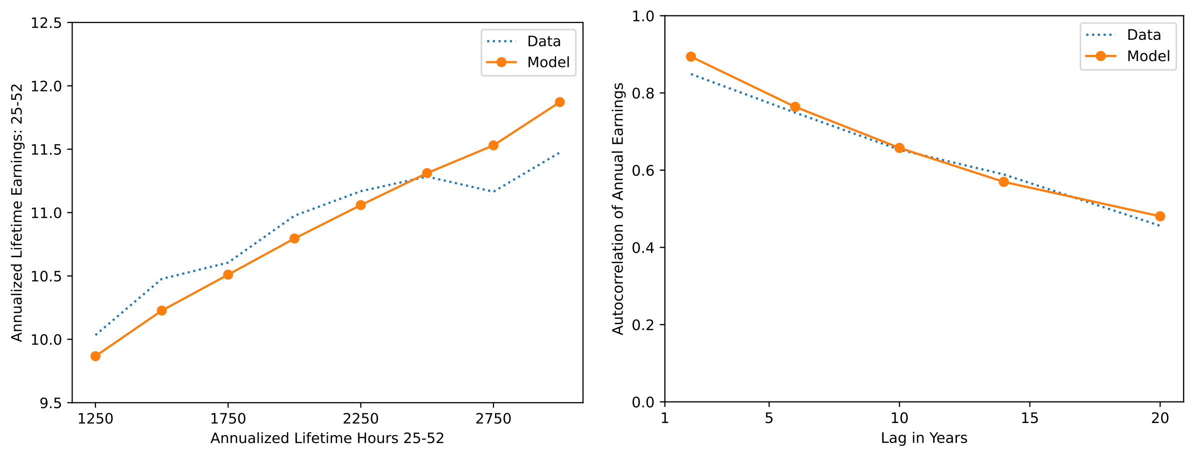

relationship extends fairly linearly into this region whereas the relationship in the data flattens out. This suggests that it might be reasonable to introduce some additional curvature in the investment production function at either high levels of human capital or high levels of hours.

While our calibration procedure targeted the autocorrelation properties of hours, it did not explicitly target the autocorrelation properties of earnings. But panel (b) shows that it nonetheless does a good job of accounting for the these properties of earnings as well.

6 Sources of Earnings Inequality

In this section we use our calibrated benchmark model from Section 5 to assess the contribution of different factors to both the evolution of earnings inequality over the life cycle and inequality in lifetime earnings. We present three main results. The first result shows that heterogeneity in time allocations accounts for a large part of overall earnings inequality. The second result focuses on the increase in cross-sectional inequality over the life cycle and shows that preference heterogeneity accounts for a substantial share of the increase in both earnings and human capital inequality over the life cycle. The third result focuses on the sources of inequality in lifetime earnings. While initial human capital differences and shocks to human capital are the most important sources, we find that preference heterogeneity also contributes significantly.

6.1 Heterogeneity in Time Allocation and Earnings Inequality

In this section we use our model to highlight the role of heterogeneity in time allocations as an important contributor to earnings inequality. To do this we carry out two exercises.

In the first exercise, we recompute outcomes in our calibrated model assuming a common life cycle profile for both time devoted to producing and time devoted to investing. That is, we compute the life cycle profile for mean values of $n_{i,j}$ and $s_{i,j}$ for each j in our calibrated model, denoted by $\bar{n}_j$ and $\bar{s}_j$ respectively and then compute outcomes assuming that each individual follows exactly this time allocation over their life cycle. In particular, we do not let an individual's time allocation respond to shocks.

We compute the variance of log earnings in the cross-section in both our benchmark calibrated economy and this counterfactual economy. We find that both the cross-sectional variance of log earnings and the variance of the log of lifetime earnings decrease by forty percent!

In the above counterfactual, we eliminate cross-sectional heterogeneity in both total work hours and the allocation of these work hours to production and investment for each age. It is also of interest to consider the effect of each of these dimensions of heterogeneity. To do this, we next consider a second counterfactual in which we set the profile for total hours at the individual level $h_{i,j}$ equal to its profile in the original equilibrium, but assume that everyone allocates total hours between producing and investing according to the average life cycle profile. That is, for each age we compute the mean value of $n_{i,j}/(n_{i,j}+s_{i,j})$, which we denote by f_j. We then compute individual level outcomes assuming each individual i works $h_{i,j}$ total hours at age j but allocates the same fraction f_j of these hours to investing, and a fraction $1-f_j$ to producing. When we do this we find that the overall variance of log earnings decreases from 0.530 to 0.491, a decrease of

Table 5: Increase in Cross-Sectional Variance Over the Life Cycle Between Ages 25 and 50

	Benchmark	$\sigma_z = 0$	$\sigma_\alpha = 0$	$\sigma_\psi = 0$	$\sigma_\pi = 0$
log Earnings	0.21	0.02	0.13	0.14	0.23
log Human Capital	0.28	0.02	0.25	0.25	0.28

less than seven percent. The variance of the log of lifetime earnings decreases by a bit more than four percent.[17] We conclude that heterogeneity in total hours of work plays an important role in generating inequality in earnings.

In summary, the key message from these counterfactual exercises is that heterogeneity in time allocation plays a key role in generating dispersion in lifetime earnings.

6.2 Increase in Inequality Over the Life Cycle

A key feature of the data is the increase in the dispersion of log earnings over the life cycle. Our model features four driving forces which serve to increase the variance of log earnings over the life cycle: the accumulation of persistent shocks, heterogeneity in learning ability, permanent heterogeneity in preferences and transitory heterogeneity in preferences. We assess the relative importance of each of these driving forces in accounting for the increase in the variance of both log earnings and log human capital over the life cycle. To do this we consider four counterfactuals in which we shut down each of these driving forces in our benchmark calibrated economy and resolve for the model's equilibrium. Table 5 presents results.

We begin with the results for human capital. The largest contributor to the increase in the variance of earnings over the life cycle is shocks to human capital. Eliminating these shocks reduces the increase in the variance of log human capital by over 90 percent. Permanent preference heterogeneity and learning ability have similar effects, with each reducing the increase in log human capital over the life cycle by a bit more than twelve percent. The transitory component of preferences has essentially no effect. The fact that these contributions collectively add up to more than 100 percent reflects the fact that if all of these driving forces are shut down we would actually observe a decrease in the variance of human capital over the life cycle. In the absence of any of these sources of heterogeneity, there is nothing to maintain the initial heterogeneity in human capital and the variance will shrink over time.

[17] If we alternatively give everyone the average profile for total hours but let each individual allocate total hours between producing and investing according to the fractions in the original calibration, the drop in the variance of log earnings is roughly 34 percent and the drop in the variance of the log of lifetime earnings is a bit less than 36 percent, indicating that the interaction between the two is not substantial.

Table 6: Variance of log Lifetime Earnings and Hours

	Benchmark	$\sigma_z = 0$	$\sigma_x = 0$	$\sigma_\alpha = 0$	$\sigma_\psi = 0$	$\sigma_\pi = 0$
Log Lifetime Earnings	0.345	0.185	0.192	0.334	0.267	0.324
Percent of Benchmark	100%	53.6%	55.7%	96.8%	77.4%	93.9%
Log Lifetime Hours	0.0276	0.0209	0.0257	0.0263	0.001	0.0147
Percent of Benchmark	100%	75.7%	93.1%	95.2%	3.6%	53.3%

The results for earnings are similar, though we now find a larger impact of permanent preference heterogeneity. This larger effect does not reflect an increase in the variance of hours over the life cycle. As noted earlier, the variance of log hours is relatively flat over the life cycle and our calibrated model replicates this pattern. The reason that permanent heterogeneity has a larger contribution to the increasing variance of earnings than to the increasing variance of human capital is because of heterogeneity in the allocation of hours between production and investment at young ages. High hours individuals increasingly allocate hours to production over the life cycle, thereby magnifying the changes in human capital and creating greater dispersion in earnings.

6.3 Inequality in Lifetime Earnings

In this subsection we repeat the exercises from the previous subsection but focusing on the inequality in the log of lifetime earnings. Relative to the previous subsection, we now consider one additional factor, namely the initial variance of human capital. Table 6 presents the results of shutting down each channel individually. This Table also reports the log of lifetime hours, which we measure as the simple sum of hours between ages 25 and 50.

It remains the case that shutting down shocks to human capital has the single largest effect, leading to a reduction of almost 50 percent. Initial differences in human capital contributes roughly 45 percent and permanent preference heterogeneity contributes a bit more than 20 percent to the overall variance of log lifetime earnings. The other two factors have relatively small effects, with heterogeneity in learning ability contributing only around three percent and transitory preference heterogeneity contributing about six percent. If we shut down both sources of preference heterogeneity the drop in the variance of the log of lifetime earnings is a bit larger than 26 percent.

In a textbook static model of labor supply, changing the variance of log hours will have a one-for-one impact on the variance of earnings among workers with the same wage. To gauge the extent to which this effect is driving our results it is instructive to also look at the results for the variance in the log of lifetime hours. As the table shows, shutting down permanent preference heterogeneity reduces the variance of log hours by roughly 2.75 log points while reducing the variance of log lifetime earnings by almost seven log points, an effect that is almost three times larger. So consistent with results presented earlier in this paper, permanent preference heterogeneity generates significant effects on human capital accumulation. Also consistent with earlier results, transitory preference heterogeneity generates very small effects on human capital.

It is useful to compare these results with those in Huggett, Ventura and Yaron (2011). Their analysis abstracted from the choice of total hours of work and so cannot speak to the importance of hours. They instead focused on a decomposition between shocks to human capital versus initial conditions, where the latter includes both initial human capital and heterogeneity in learning ability. They found that shocks contributed 38.5% and initial conditions contributed 61.5%, so that initial human capital and learning ability are more than one and a half times as important as shocks during the lifetime. Our results are quite different: in our model the contribution of initial human capital and learning ability are roughly the same as that of human capital shocks.[18]

The fact that initial human capital is less important in our model reflects the fact that we have an additional margin that helps us account for the dispersion of earnings in the initial period. Preference heterogeneity generates heterogeneity in total hours of work and this generates additional heterogeneity in earnings. Huggett, Ventura and Yaron (2011) chose parameters so as to match the initial dispersion in earnings without allowing for heterogeneity in total hours, and as a result required greater dispersion in initial human capital.

6.4 A Policy Application

In this subsection we discuss two policy implications that result from our model. The first concerns messages for policy makers thinking about ways to reduce inequality. In particular, the greater is the contribution of initial human capital to lifetime inequality, the greater is the potential role for policies that will affect the accumulation of human capital prior to labor market entry. While initial human capital continues to play an important role in our calibrated model, it plays a much less important role than in Huggett, Ventura and Yaron (2011). Another way to quantify this is to compute the average impact on lifetime earnings of a one standard deviation increase in initial human capital. Huggett, Ventura and Yaron (2011) do this calculation for a group of individuals with median values of all initial conditions and find that it increases lifetime income by 47.5 percent. If we do this same calculation in our model, the increase is only 33.6 percent. In our model we can also compute the effect on lifetime earnings of a one standard deviation increase in the value of ψ for a group of individuals with median values of all initial conditions. We find that this reduces lifetime earnings by 17.9 percent.

[18] These measures are not strictly comparable, as their model starts individuals at age 23 whereas we start individuals at age 25. But if anything, this should make initial conditions even more important in our analysis.

A second message from our analysis is that policies which directly impact total hours may have important effects on human capital, which may affect both the level of human capital and dispersion of human capital. In this subsection we consider one such policy, a regulation in France that restricts individuals to work no more than 48 hours per week.

Although the law does allow for some exceptions, in what follows we will interpret it as a hard constraint, though to be conservative, we implement this restriction in our model as a restriction on hours spent in production. This implicitly assumes that investment in human capital might feasibly take place away from work and that workers and firms might work around the hours restriction by reallocating investment time. We thus take our full calibrated model, impose a restriction that annual hours of production time are constrained to be less than 48 hours * 52 weeks = 2496 and then resolve the model. We find that this restriction has large effects on both the level and dispersion in lifetime earnings. In particular, the mean of log lifetime earnings falls by almost thirteen percent and the variance of log lifetime earnings falls by twenty two percent. Although the drop in mean lifetime earnings is partly explained by a drop in mean hours, there is an important amplification effect that occurs through the effect of this policy on human capital accumulation. The drop in lifetime hours is 7.6 percent and the drop in average human capital is 4.1 percent.

Our earlier result about the importance of permanent preference heterogeneity help us understand the reason for these large effects. Specifically, this regulation is effectively not binding for young individuals. Because time devoted to investment is high for young individuals, even those with very low values for ψ will not find the constraint on production time to be binding. However, although the constraint does not bind when young, the fact that time spent in investment is very low at older ages implies that the constraint will bind at old ages. But our earlier results implied that expected hours at older ages play a particularly important role in incentivizing young individuals to invest in human capital. Thus, the fact that the constraint will bind when old greatly dampens the human capital investment of some individuals.

Motivated by the analysis in Guvenen, Kuruscu and Ozkan (2014), we have also carried out a tax experiment in our model. Guvenen, Kuruscu and Ozkan (2014) argued that the Ben-Porath model created a connection between progressive income taxation and inequality in wages. In their calibrated model they find that this effect is empirically relevant, finding that differences in progressive income taxation across OECD countries can account for roughy half of the differences in wage inequality between the US and Europe. As noted previously, Guvenen, Kuruscu and Ozkan (2014) do not allow for preference heterogeneity and so do not generate as much dispersion in hours as in the data. We have redone their exercise to evaluate the effects of moving from a progressive tax on labor income to a flat tax, but found that incorporating preference heterogeneity does not have a first order effect on their finding.

7 The Role of Preference Heterogeneity

Huggett, Ventura and Yaron (2011) assumed that work hours are the same for all individuals in all periods. Our model offers two extensions relative to them. First, we allow for an endogenous labor-leisure choice, and second, we assume heterogeneous tastes for work. Even in the absence of heterogeneity in preferences, idiosyncratic shocks to human capital accumulation, heterogeneity in learning ability and heterogeneity in initial human capital will create cross-sectional heterogeneity in hours. In view of this we think it is also of interest to explore the effects of adding an endogenous labor leisure choice in the absence of preference heterogeneity.

To answer this question we analyze a version of our model that abstracts from preference heterogeneity, i.e., we consider our model with the restrictions that $\sigma_\psi = \sigma_\pi = 0$ and $\rho_\pi = 1$. While the resulting model is similar in spirit to the model in Guvenen, Kuruscu and Ozkan (2014), they did not assess the extent to which their model could account for salient features of the hours distribution, nor did they assess the contribution of various factors to inequality. We proceed in two steps. In the first step we calibrate this version of the model to match the same moments of the earnings distribution as our earlier calibration, and then assess the extent to which it matches moments of the hours distribution. In the second step we use this calibrated model to assess the contribution of various factors to inequality.

7.1 Calibration

In calibrating this model we proceed as before with the following modifications. First, we no longer target the standard deviation of hours in the cross-section or any of the autocorrelation properties of annual hours worked at the individual level. These properties played an important role in determining the values of σ_ψ, σ_π, ρ_π, and σ_{mh}. As before, we do calibrate the value of μ_ψ to match the mean value of hours. Second, while the values for σ_ψ, σ_π and ρ_π are now set exogenously, we still need a value for σ_{mh}. We exogenously set its value equal to its value in the benchmark calibration. Third, because we no longer target properties of the hours distribution, we no longer target the relationship between lifetime hours and lifetime earnings growth, and instead exogenously set the value of ϕ equal to its value from the benchmark calibration.

This model has only 4 parameters that are calibrated internally: μ_ψ, μ_α, σ_α, and σ_x. Table 7 reports the calibrated values for these parameters along with the values from our earlier calibration in the full model.

The differences in calibrated parameters are intuitive. Recall that the mean and standard deviation of learning ability are essentially chosen so as to match mean growth of log earnings and the

Table 7: Jointly Calibrated Parameter Values: Homogeneous Preference Model

Parameter	Interpretation	Value	Moment
μ_α	Mean of $\log\alpha$	-2.314	Mean earnings, age 50
σ_a	SD of $\log\alpha$	0.197	SD earnings, age 50
σ_x	SD of $\log x_0$	0.465	SD earnings, age 30
μ_ψ	Mean of $\log\psi$	3.296	Mean annual hours, age 25-52

Figure 11: Fit of Model with Homogeneous Preferences: Earnings

(a) Mean

(b) Standard Deviation

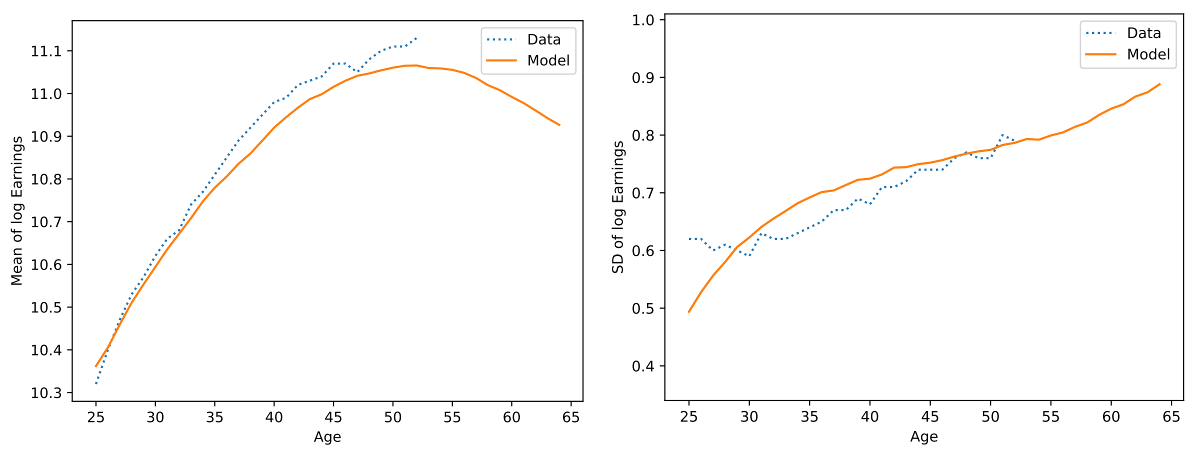

increase in the variance of log earnings given initial conditions and the shocks to human capital accumulation. If we eliminate preference heterogeneity then there is a greater role for learning, and as a result we see that both the mean and standard deviation of α have increased. Similarly, the distribution of initial human capital is important in generating the initial distribution in earnings, and given that we eliminated preference heterogeneity there is a greater role for heterogeneity in initial human capital to generate the initial dispersion in earnings. Consistent with this, we see a larger value for σ_x.

Figure 11 shows the profiles for both the mean and standard deviation of log earnings over the life cycle in both the model and the data.

As can be seen, this restricted version of our model accounts for these moments equally well as our benchmark model. This is probably not surprising given that Huggett, Ventura and Yaron (2011) established that their model does a good job of tracking properties of the earnings distribution over the life cycle, and our model is essentially their model extended to have an endogenous labor-leisure choice.

Figure 12: Fit of Model with Homogeneous Preferences: Hours

(a) Standard Deviation

(b) Autocorrelation

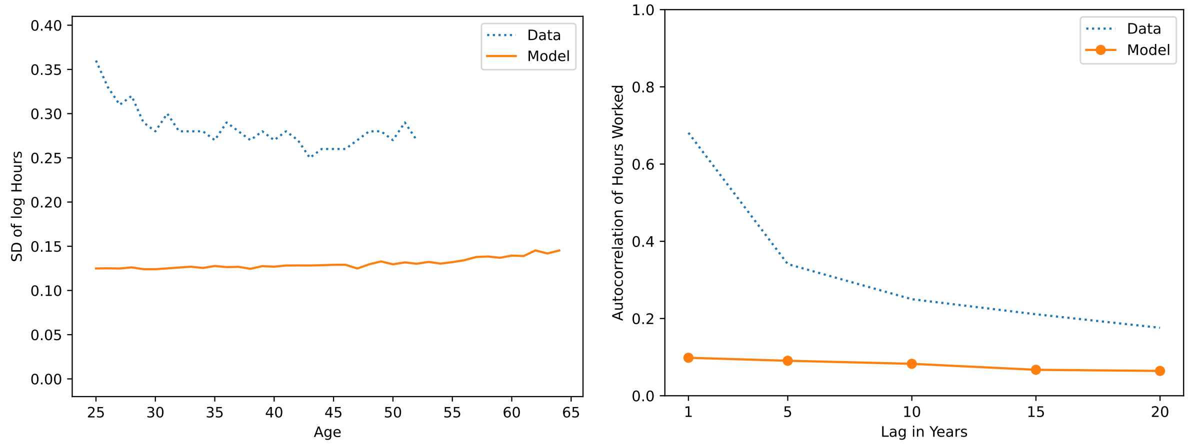

Importantly, however, this model fails very significantly to account for the salient features of the hours distribution. Although it matches the mean level of annual hours in the data, Figure 12 displays two dimensions along which it fails to match the data.

Panel (a) shows the standard deviation of hours over the life cycle in the model and the data. The key message from this figure is that the model can account for less than half of the dispersion in hours found in the data.[19] Panel (b) shows the autocorrelation properties of hours in the model and the data. The key message here is that the model fails to match the persistence found in individual hours. This is particularly significant because our previous analysis indicated that the extent to which variation in hours matters for inequality depends critically on the extent to which the differences persist at long horizons. So even though this model can account for roughly half of the cross-sectional variation in log hours, the variation that it misses is of much greater significance as a factor that shapes inequality. The preceding results suggest that the model will fail to generate sufficient variation in lifetime hours. Figure 13 confirms this. This figure clearly shows that annualized lifetime hours in the model are far too concentrated relative to what is found in the data.

Although the model fails to generate sufficient dispersion in lifetime hours, it still does a reasonable job of capturing the variation in lifetime earnings. This is demonstrated in the left panel of Figure 14. To understand why the failure to match the dispersion in lifetime hours does not translate into a failure to generate sufficient dispersion in earnings, recall that our calibration procedure (and the calibration procedure used by Huggett, Ventura and Yaron (2011)) chooses the variation in learning ability in order to generate the observed increase in the cross-sectional variation in earn-

[19]Recall that we are including the same amount of measurement error in hours in this model as in our benchmark model. About ten percent of the standard deviation in the log of annual hours in this model is the result of measurement error. If we instead look at the distribution of log true annual hours the drop in the standard deviation is even larger.

Figure 13: Fit of Model with Homogeneous Preferences: Lifetime Hours

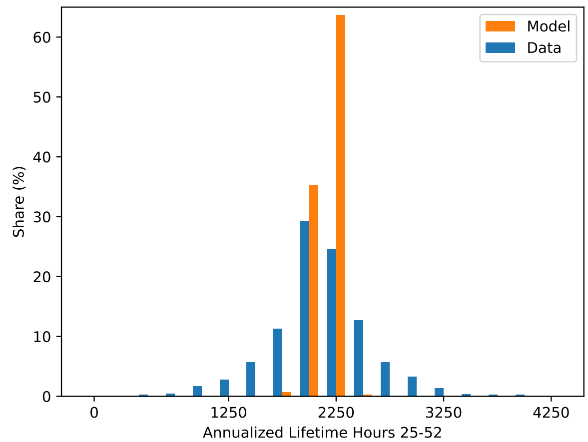

Figure 14: Model Fit: Distribution of Lifetime Earnings and Hours

(a) Earnings

(b) Earnings and Hours

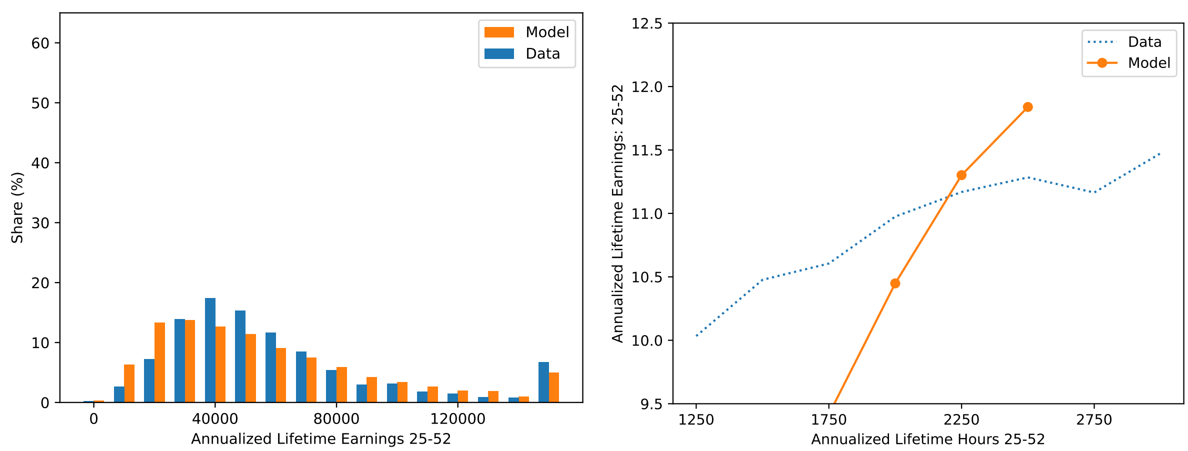

ings over the life cycle. Loosely speaking, from the perspective of generating rising inequality in earnings over the life cycle, heterogeneity in learning ability and heterogeneity in tastes for work act as substitutes. The right panel of Figure 14 shows that relying more heavily on heterogeneity in learning ability generates far too steep a relationship between lifetime hours and lifetime earnings.

7.2 Implications

In this subsection we repeat some of the analysis of Section 6 in the calibrated model without preference heterogeneity. As noted above, this calibrated model attributes a much larger role to heterogeneity in learning ability. Heterogeneity in learning ability will give rise to heterogeneity in

Table 8: Variance of log Lifetime Earnings: Homogeneous Preferences

	Benchmark	$\sigma_z = 0$	$\sigma_x = 0$	$\sigma_\alpha = 0$
Variance	0.387	0.232	0.163	0.362
Percent of Benchmark	100%	59.9%	42.1%	93.5%

human capital accumulation even in the absence of any heterogeneity in time allocations. Consistent with this, we find that heterogeneity in time allocations is much less important in this model. In particular, forcing all individuals to have the average life cycle profile for both production and investment hours only reduces cross-sectional dispersion in log earnings by 16 percent, whereas in the full model this value was 40 percent.

In this model there are only two factors that generate the increasing variance of human capital over the life cycle: shocks to human capital and heterogeneity in learning ability. Repeating our earlier analysis we find that setting $\sigma_z = 0$ reduces the increase in the variance of log human capital by almost 90 percent, while setting $\sigma_\alpha = 0$ reduces it by almost 50 percent. The importance of z shocks is relatively unchanged, but we now find a much larger role for heterogeneity in learning ability.

Lastly, we assess the contribution of various factors to the heterogeneity in lifetime income. Table 8 reports results.

We now find that heterogeneity in initial human capital is much more important than shocks to z. Heterogeneity in learning ability is much less important than the other two factors, but is much more important than in the full model. From this exercise we conclude that the main messages from Huggett, Ventura and Yaron (2011) about the relative importance of different factors is relatively robust to the inclusion of an endogenous labor-leisure choice. But as shown in the previous section, they are not robust to extending the model in a way that allows it to match the observed heterogeneity in hours.

When we use this model to evaluate the French regulation that restricts total hours, the results are dramatically different. We now find that both the mean and variance of log lifetime earnings decrease by only 0.1 percent. These smaller effects occur precisely because the model without preference heterogeneity does not generate sufficient variation in hours and the variation that it does generate exhibits very little persistence. For evaluating the effects of this particular policy, our full model provides very different answers than the model without preference heterogeneity.

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 have argued that heterogeneity in lifetime hours of work generated by heterogeneity in preferences plays a quantitatively important role in shaping lifetime earnings inequality via its impact on human capital accumulation. In particular, we find that heterogeneity in preferences accounts for more than 25 percent of the variation in the log of lifetime earnings as well as more than 25 percent of the increase in the variance of log earnings over the life cycle between the ages of 30 and 50. One key message from our analysis is that it is important to include preference heterogeneity in analyses of inequality.

An important by-product of our analysis is to document properties of lifetime hours of work using the long panel in the NLSY79. While we show that there is indeed an important transitory component to cross-sectional differences in hours of work, we show that there is also a very persistent component. We show that a mix of permanent and transitory heterogeneity provides a good match to the autocorrelation properties of the data. In particular, we show that preference heterogeneity is needed to generate the large differences in lifetime hours found in the data. That is, shocks to human capital and heterogeneity in initial human capital and learning ability are not sufficient to generate the magnitude of differences found in the data.

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.

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.

Blandin, Adam. 2018. "Learning by doing and ben-porath: Life-cycle predictions and policy implications." *Journal of Economic Dynamics and Control*, 90: 220–235.

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.

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.

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.

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, Lance Lochner, and Christopher Taber. 1998. "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents." 1.

Hosseini, Roozbeh, Karen A Kopecky, and Kai Zhao. 2022. "How important is health inequality for lifetime earnings inequality?"

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.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song. 2010. “Earnings inequality and mobility in the United States: Evidence from social security data since 1937.” *The Quarterly Journal of Economics*, 125(1): 91–128.

Kuruscu, Burhanettin. 2006. “Training and Lifetime Income.” *American Economic Review*, 96(3): 832–846.

Manuelli, Rodolfo, Ananth Seshardi, and Yongseok Shin. 2012. “Lifetime Labor Supply and Human Capital Investmen.” National Bureau of Economic Research.

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.

Rubinstein, Yona, and Yoram Weiss. 2006. “Chapter 1 Post Schooling Wage Growth: Investment, Search and Learning.” In . Vol. 1 of *Handbook of the Economics of Education*, , ed. E. Hanushek and F. Welch, 1–67. Elsevier.