Hours Worked and Lifetime Earnings Inequality*

Alexander Bick
Federal Reserve Bank of St. Louis, CEPR

Adam Blandin
Vanderbilt University

Richard Rogerson
Princeton University, NBER

September 16, 2024

Abstract

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

JEL Codes: D15, E21, E24, J22, J31, J24
Keywords: Lifetime Earnings, Hours worked, Human Capital, Inequality

---

*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, Goethe University, I-85 Macro Workshop, LAEF Workshop on Labor Markets and Macroeconomic Outcomes, LAEF and St. Louis Fed Human Capital Conference, LMU Munich, NBER Conference on Economic Activity Across Time and Space, NBER Summer Institute Micro Data and Macro Models, Society for Economic Dynamics Annual Meeting in Barcelona, Stanford Institute for Theoretical Economics Labor Market and Policies Workshop, Universtiy of Connecticut, Carleton University, University of Bonn, University of North Carolina at Charlotte, MIT, Princeton University, and Vanderbilt University for helpful commments and suggestions. Kevin Bloodworth II, Elizabeth Harding, and Siyu Shi provided outstanding research assistance. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

1 Introduction

Lifetime earnings in the US are highly unequal; for example, recent work by Guvenen et al. (2022) finds that lifetime earnings at the 75th percentile are 2.7 times larger than at the 25th percentile. Assessing the factors behind these differences in lifetime earnings is one of the core objectives of the literature on inequality. While the literature has investigated several factors, little attention has been devoted to the role of labor supply. In this paper we argue that labor supply, and in particular its interaction with human capital accumulation, plays an important role in accounting for inequality in lifetime earnings.

Our paper offers three contributions. First, we document the nature and extent of heterogeneity in lifetime hours of work using the National Longitudinal Study of Youth 1979 (hereafter, NLSY79). Second, we quantify the effect of this heterogeneity on lifetime earnings inequality using a Ben-Porath model of human capital accumulation extended to include a labor supply decision. Third, we decompose the effect of lifetime hours inequality on lifetime earnings inequality into a direct channel due to higher hours and a human capital channel due to higher human capital accumulation.

The first step in our analysis is to document the extent and nature of differences in lifetime hours of work. While cross-sectional differences in annual 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 characterize the stochastic process for annual hours at the individual level and compute differences in lifetime hours of work. We highlight three findings. First, differences in lifetime hours are large, even when restricting attention to workers with high labor force attachment. For our sample of highly attached male workers, the interquartile range reflects differences of more than 400 hours per year over their life-cycle, which is almost 20 percent of median lifetime hours. Second, cross-sectional differences in annual hours reflect both a transitory and a permanent component. In particular, individual hours are not well captured by a simple AR(1) process. Third, we document a strong positive correlation between lifetime hours and life-cycle earnings growth.

The second step in our analysis examines the forces that shape human capital accumulation in a simple Ben-Porath model. This framework provides a very intuitive mechanism connecting heterogeneity in lifetime hours of work with heterogeneity in human capital accumulation. We show analytically that higher expected future hours of work increase the incentive for human capital accumulation today. Heterogeneity in expected future hours will thus give rise to heterogeneity in human capital accumulation and generate heterogeneity in lifetime earnings.

The third step in our analysis carries out a quantitative analysis of the sources of lifetime earnings inequality in a rich heterogeneous agent life-cycle economy. As in much of the existing literature, we allow for heterogeneity in initial human capital and learning ability, as well as persistent shocks to income. Our key extension is to allow for heterogeneity in tastes for work that allows our model to replicate the salient facts of heterogeneity in hours of work over the life-cycle.

We calibrate the parameters of this model to match a large set of moments that characterize the distribution of earnings and hours over the life-cycle. Matching these moments requires two sources of preference heterogeneity: a permanent component and a persistent but transitory component. While the transitory component generates a substantial amount of cross-sectional variation in annual hours, the permanent component is the dominant source of variation in lifetime hours.

We use the calibrated model to assess the factors that generate inequality in lifetime earnings and hours. Four key results emerge. First, almost 20 percent of the variance of log lifetime earnings comes from hours heterogeneity. Second, between a third and a half of this effect is attributed to human capital accumulation. Third, roughly 90 percent of the effect of hours heterogeneity on lifetime earnings inequality is due to preference heterogeneity. Fourth, these results are driven by the effect of preference heterogeneity on differences in lifetime hours as opposed to differences in annual hours. When we eliminate preference heterogeneity, a substantial amount of cross-sectional variation in annual hours remains, but there is almost no variation in lifetime hours and almost no effect of hours on lifetime earnings.

Our model of hours and earnings 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 that 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. Motivated by a recent regulation adopted in France that limits weekly hours to be no greater than 48, we use our model calibrated to the US to evaluate the partial equilibrium consequences that would result if the US were to adopt such a regulation. Our model implies that this would reduce both mean earnings and inequality in lifetime earnings. Importantly, the workers who are most negatively affected by this regulation are spread throughout the lifetime earnings distribution.

Our paper lies at the intersection of three large literatures—inequality, human capital accumulation, and labor supply. These literatures are too vast to attempt a summarize here, so we restrict our discussion to a small set of closely related papers that have inequality as their core focus.[1] Huggett, Ventura and Yaron (2006, 2011), and Guvenen, Kuruscu and Ozkan (2014) all use a Ben-Porath model to study inequality in a life-cycle setting. Relative to them, our key contribution is to assess the role of hours inequality for lifetime earnings inequality. The two papers by Huggett,

[1] In particular, we do not review the large literature that studies labor supply in models of human capital accumulation. While not unrelated to inequality, many of these papers have focused on labor supply elasticities and the effects of policy. Heckman (1976*a*,*b*) and Shaw (1989) are early contributions. Subsequent contributions of note include Keane and Wolpin (1997), Imai and Keane (2004), Wallenius (2011, 2013), Keane and Wasi (2016), Blundell et al. (2016)), Blundell et al. (2021), and Fan, Seshadri and Taber (2024). Stantcheva (2017) and Badel, Huggett and Luo (2020) study optimal taxation of labor income in models with human capital accumulation.

Ventura and Yaron (2006, 2011) study lifetime earnings inequality but assume no heterogeneity in total hours and so cannot speak to the effect of heterogeneity in lifetime hours on lifetime earnings. Guvenen, Kuruscu and Ozkan (2014) use their model to study the effect of progressive taxes on inequality. They have an endogenous choice of total hours, but they do not allow for preference heterogeneity, which (as we show) implies essentially no heterogeneity in lifetime hours.

Manuelli, Seshadri and Shin (2012) also study lifetime labor supply and human capital accumulation in a Ben-Porath model. But they focus on inequality across countries due to differences in length of working life, whereas we focus on inequality across individuals within the US due to heterogeneity in hours along the intensive margin.[2]

Kaplan (2012) and Heathcote, Storesletten and Violante (2014) study hours heterogeneity in a life-cycle setting that abstracts from human capital accumulation. Kaplan (2012) specifically focuses on the reduction in cross-sectional variance of hours in the early part of the life-cycle, an issue that we do not explicitly address, as this is not such a prominent feature in our sample of men with high labor market attachment. Heathcote, Storesletten and Violante (2014) study cross-sectional inequality in both wages and hours over the life-cycle. Neither analysis considers the distinction between variation in annual hours versus lifetime hours, and neither studies the connection between hours inequality and lifetime earnings inequality.

An outline of the paper follows. Section 2 documents properties of lifetime hours and earnings using the NLSY79. Section 3 establishes analytically that higher future hours lead to higher human capital accumulation in a Ben-Porath model. Section 4 describes our quantitative model and Section 5 calibrates it. Section 6 presents our main results about the role of preference heterogeneity for inequality. Section 7 studies the effects of the French regulation on hours and Section 8 concludes.

2 Empirical Facts on Hours and Earnings Inequality

This section describes how we use the NLSY79 to create a balanced panel of individual observations on annual earnings and annual hours worked from age 25 to 55. We document several properties of lifetime hours for this balanced panel, in particular how the dispersion in lifetime hours varies as we consider samples with varying degrees of labor force attachment. We then present a variety of statistics for hours and earnings for a subsample of highly attached males that will be used in the quantitative analysis later in the paper.

[2] Glover, Mustre-del Río and Pollard (2023) examine how differences in years of work contribute to lifetime earnings differences between whites and blacks.

2.1 Data and Measurement

Our empirical analysis is based on the NLSY79, a longitudinal study of 12,686 individuals born between 1957 and 1964.[3] Respondents were recruited and initially interviewed in 1979, when they were between 14 and 22 years old. The NLSY79 was conducted annually through 1994, after which it was conducted biennially, occurring in even numbered years. The most recent available data are from the interview year 2020. The last year for which we have earnings data is 2019, at which point the youngest individuals in the sample are 55 and the oldest are 63.

The NLSY79 records starting and end dates of each job a person holds and collects usual weekly hours worked in each job since the last interview in which the respondent participated. Although interviews are biennial after 1994, the employment history thus covers all years. From this, one can construct a weekly history of employment status and usual weekly hours worked across all jobs since 1978. We aggregate this information to produce annual measures of weeks worked, usual weekly hours, and total annual hours.

Respondents also report their annual earnings for the calendar year preceding the interview year. Thus, whereas employment histories are collected for all years, earnings are not collected for even numbered years starting in 1994. Information is collected separately for two categories of earnings: (i) income from wages, salary, commissions, or tips from all jobs before deductions for taxes or anything else last year, and (ii) income received from a farm/business owned last year. Our measure of earnings is the sum of these two components, and following Guvenen et al. (2022) we deflate earnings with the personal consumption expenditures index normalized to one in 2013. We measure hourly wages as annual earnings divided by annual hours.[4]

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

To create this balanced panel we need to address the issue of missing values. In addition to the years in which earnings data are not collected, there are also missing data for individual

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

[4] Following the procedure in Bick, Blandin and Rogerson (2024), we make two adjustments to control for hourly wage outliers. First, wages below half of the Federal minimum wage are set equal to that value, with earnings adjusted accordingly. Second, and motivated by the evidence in Bick, Blandin and Rogerson (2024), we treat wage observations in the top 0.1% of the wage distribution as due to misreported hours, and thus set weeks worked and weekly hours to missing for these observations.

Table 1: Hours Worked in Our Balanced Panel

	All	Men	Women
Mean	1819.3	2103.8	1559.5
Coefficient of Variation			
Annual	0.54	0.44	0.62
Lifetime	0.38	0.29	0.41
Individuals	6335	3029	3306

responses to hours and/or earnings questions in some interviews. We impute missing values using the interpolation procedure described in Bick, Blandin and Rogerson (2024). Given that we rely on interpolation to fill in missing values and our goal is to have a balanced panel through age 55, we remove from our sample any individuals who do not have a complete interview at age 55 or older. In particular, this will remove all individuals who leave the sample prior to age 55. We also remove individuals for whom there are not sufficiently nearby observations to be used in the imputation procedure. Specifically, we impose that for any year in which employment status, hours worked, or earnings are missing for the entire year, there must be at least one observation within the previous or next five years.[5] One potential concern with the balanced panel we create is that the effects of attrition and our selection criterion based on missing values may not be random. In Bick, Blandin and Rogerson (2024) we argue that selection effects are not significant and our imputation method works well. In particular, we show that life-cycle profiles of employment, hours worked, and earnings (both means and standard deviations where applicable) in our balanced sample are comparable to those in the Current Population Survey for the same cohorts. The lifetime earnings distribution in our sample also closely aligns with the distribution for the same cohorts in Social Security Administration data. We refer the reader to that paper for more details.

Implementing our procedure yields a balanced panel data set with annual data on weeks worked, usual weekly hours, total hours, earnings, and average hourly wages for a sample of 6335 individuals between ages 25 and 55. This balanced panel will be of interest for many issues that feature labor supply decisions in a life-cycle setting. A novel issue of particular interest to us in this paper is the extent to which lifetime hours vary across individuals, and we compare differences in lifetime hours with differences in annual hours. For the remainder of this paper we define (annualized) lifetime hours for an individual in our sample to be the sum of annual hours from age 25 to 55 divided by 31, the number of annual observations for each individual. When studying the properties of annual hours we will pool the annual observations from our balanced panel.

Table 1 presents several summary statistics for our balanced panel, both for the overall sample

[5]In Bick, Blandin and Rogerson (2024), we document the effects of varying this threshold and conclude that five years strikes a reasonable balance between maximizing sample size and minimizing measurement error in our imputation procedure.

Table 2: Distribution of Attachment in our Balanced Panel

	1	5	10	15	20	25	30	31
All	98.5%	96.7%	93.9%	89.2%	80.7%	69.6%	47.1%	36.3%
Men	99.1%	98.2%	96.3%	93.5%	87.1%	78.7%	59.0%	46.8%
Women	97.9%	95.4%	91.6%	85.2%	74.8%	61.1%	36.1%	26.6%

as well as separately for men and women.[6] As is well known from previous work, the dispersion in annual hours is large: For the overall sample the coefficient of variation is 0.54. The dispersion in lifetime hours is about one-third smaller than the dispersion in annual hours, indicating that some of the variation at annual frequency averages out over time. Nonetheless, it remains true that the dispersion in lifetime hours is also substantial, with a coefficient of variation equal to 0.38. There is much greater dispersion for women than for men in both measures, but for both groups the coefficient of variation drops by roughly one third when moving from annual hours to lifetime hours.

2.2 Sample Selection

The balanced panel we created in the previous subsection did not impose any explicit criterion regarding labor force attachment. As a result, our sample displays considerable heterogeneity in terms of attachment. Following Guvenen et al. (2022), we consider an individual to be "attached" in a given year if they work at least 520 hours. Using this definition, we can compute the years of attachment for each individual in our balanced panel. Table 2 displays the fraction of individuals in our sample that meet various thresholds for years of attachment, both for the overall sample as well as by gender.

The table shows that there is substantial heterogeneity in the extent of attachment over the life-cycle. While more than 90 percent of all men have at least 15 years of attachment between ages 25 and 55, and almost 80 percent have at least 25 years of attachment, less than 50 percent have the full set of 31 years of attachment. Not surprisingly, women display considerably less attachment than men: Only around 85 percent have at least 15 years of attachment, and less than 30 percent have the full 31 years.

The extent of dispersion in both annual and lifetime hours changes considerably as one varies the attachment threshold. To see this, Table 3 reports mean hours and the coefficient of variation for both annual and lifetime hours for four different attachment thresholds. The values in this table are for the overall sample. Results by gender are contained in Appendix Tables B.1 and B.2.

[6]The sample of annual observations include some zero values for hours. For this reason we report the coefficient of variation for hours rather than the standard deviation of log hours.

Table 3: Distribution of Hours Worked for Different Attachment Thresholds

	1	10	20	31
Mean	1837.1	1897.4	2041.8	2320.4
Coefficient of Variation				
Annual	0.53	0.49	0.40	0.26
Lifetime	0.36	0.32	0.24	0.16
Individuals	6237	5946	5111	2297

As we vary the threshold from one year of attachment to 31 years of attachment, mean hours increase by almost 500 hours, and the coefficients of variation fall: from 0.53 to 0.26 for annual hours and from 0.36 to 0.16 for lifetime hours. Importantly, even when the attachment threshold is set to 31 years, the dispersion in lifetime hours remains large: Given that mean hours are roughly 2300, a one standard deviation difference amounts to more than 350 annual hours, or roughly seven hours per week. Note that the large gap in our measure of dispersion for annual and lifetime hours prevails independently of the attachment threshold.

Our model and analysis later in this paper build on the framework in Huggett, Ventura and Yaron (2011) and Guvenen, Kuruscu and Ozkan (2014), both of which effectively abstract from the participation decision. With this in mind, we apply two additional criteria in order to select a subsample of our balanced panel that minimizes the potential importance of non-participation spells. First, and following the two previously mentioned papers, we exclude women from our analysis, since women face more frequent and longer spells of non-participation in the labor market due to child births and child rearing and our model will abstract from fertility. Second, we focus on the subsample of males who have 31 years of attachment from age 25 to 55. This leaves us with a sample of 1418 men. We have also carried out our analysis on the sample of males with at least 20 years of attachment. This has relatively little effect on any of the results that we stress later on. Results for this alternative sample are contained in Appendix B.3.

As noted above, the results in Table 3 show that a higher attachment threshold tends to decrease the variation in lifetime hours. In view of this, the results that we present later in this paper for the role of variation in lifetime hours in our highly attached sample might reasonably be viewed as conservative when considering a broader sample.

2.3 Properties of Hours and Earnings for the High Attachment Sample

In this subsection we document properties of hours and earnings for our highly attached sample of males that will be used in our quantitative analysis. Several of our figures will document life-

Figure 1: Hours Worked Over the Life-Cycle, Ages 25 – 55

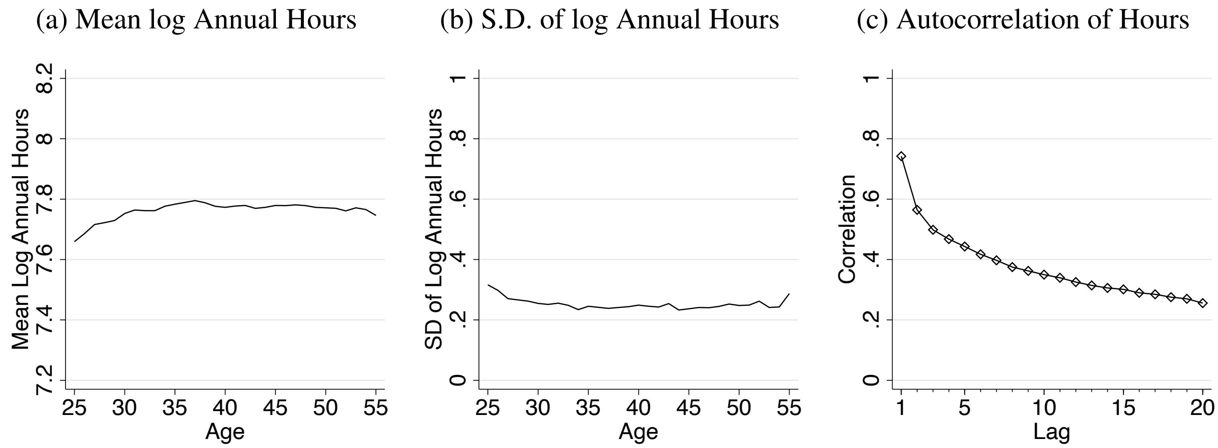

cycle profiles for various outcomes. A common challenge in this context is to distinguish between cohort, time, and age effects. Because of the relatively narrow age range of participants initially surveyed by the NLSY79, we feel it is appropriate to view our data as capturing outcomes for one specific cohort of individuals over their life-cycle. Conditional on focusing on a single cohort, the life-cycle profiles we display in this section should be understood as reflecting both time and age effects. While we do not attempt to distinguish between time and age effects in this section, the model that we use later in this paper will incorporate time effects in the form of time variation in the wage rate per efficiency unit of labor. This assumption in conjunction with the structure of the model will implicitly decompose these profiles into age and time effects.

2.3.1 Properties of Hours

Figure 1 presents the mean and standard deviation of annual hours worked over the life-cycle. Over the first ten years of the life-cycle there is a substantial increase (12.4 log points) in the mean of log hours and a substantial decrease (-7.1 log points) in the standard deviation of log hours.[7] But from the early thirties onward both profiles exhibit little trend.

As noted earlier, the cross-sectional variation in annual hours of work is large: An individual between ages 30 and 50 who works one standard deviation above the mean works roughly 50% more hours than someone working one standard deviation below the mean.

The extent to which differences in annual hours translate into differences in lifetime hours depends on the persistence of annual hours at the individual level. To document the persistence

[7] This decrease in the standard deviation of hours in the early part of the life-cycle was the focus of the analysis by Kaplan (2012) that we cited in the introduction, though we note that the reduction for our highly attached sample is smaller than the reduction he found for a broader sample.

Table 4: The Distribution of Lifetime Hours and Its Components

Percentile	Annualized Lifetime Hours	Weeks per Year Worked	Hours per Week Worked
5	1982.7	49.2	40.4
10	2054.8	50.6	40.7
25	2155.1	50.8	42.4
50	2340.3	51.1	45.8
75	2588.4	51.0	50.8
90	2904.4	51.5	56.4
95	3141.6	50.7	62.0

Notes: Individuals are sorted into percentiles according to their annualized lifetime hours. *Weeks per year worked* and *Hours per week worked* are the average values for all individuals in a given percentile and the two adjacent percentiles. Note that the product of the two variables is slightly different than the average annualized lifetime hours for each percentile.

of hours differences over time, we exploit the long panel dimension of our sample to compute the autocorrelation of hours at lags ranging from one to twenty years. To construct autocorrelations at lag length $t > 0$, we collect all length-t pairs of hours $\{(h_{i,a}, h_{i,a+t})\}_{a=25}^{55-t}$ from our sample of individual hours profiles $\{h_{i,a}\}_{a=25}^{55}$ and compute the pairwise correlation of this collection. Figure 1c plots the resulting correlation coefficients for lag lengths $t = 1, ..., 20$.[8]

The key message from Figure 1c is that although there is some tendency for mean reversion, hours display considerable persistence over long horizons.[9] The 1-year autocorrelation of hours is 0.741. The autocorrelation falls by about half, to 0.431, when moving to a lag length of 5 years, after which it continues to decrease but at a slower rate. At 20 years the autocorrelation is still substantially above zero, at 0.251. 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 0.741 at a lag of one year would have an autocorrelation of effectively zero at a twenty-year lag.

The first two columns of Table 4 show that the high persistence of hours at long lags generates large differences in lifetime hours. The interquartile range is more than 400 hours per year, and the 90-10 gap is more than 800 hours per year. Even for this sample of highly attached males, these are large differences in lifetime hours corresponding to almost one fifth and one third of median hours, respectively.

Variation in lifetime hours for our highly attached sample can come from variation in average

[8]One may worry that the survey design and our imputation procedure creates too much persistence in annual hours worked. We address this concern in the Appendix. In particular, Appendix Figure B.1a shows that the autocorrelation profile excluding imputed values for annual hours is virtually identical to the profile shown in Figure 1c.

[9]It is of interest to know if the properties of the autocorrelation function vary with the attachment threshold used to construct our sample. In Appendix B.3 we display the autocorrelation functions for the subsample requiring 20 years of attachment and show that the quantitative properties of the autocorrelation function are quite robust to this variation.

annual weeks worked or average weekly hours. The last two columns in Table 4 show how these margins vary across the lifetime hours distribution. Virtually all of the difference in lifetime hours of work is accounted for by differences in average weekly hours rather than average weeks worked. This does not imply that individuals in our sample do not experience any variation in weeks worked across years. Rather, it implies that differences in weeks worked across individuals in our highly attached sample tend to average out over time. Table 5 confirms this more formally by presenting results for a variance decomposition.[10]

Table 5: Variance Decomposition of Annual and Lifetime Log Hours

		% Explained by		
	Var(log Hours)	Var(log Weeks)	Var(log Weekly Hours)	2 Cov
Annual	0.067	25.2%	72.0%	2.8%
Lifetime	0.022	4.5%	94.6%	0.9%

2.3.2 Properties of Earnings

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. See, for example, the figures in Huggett, Ventura and Yaron (2006, 2011), or Heathcote, Storesletten and Violante (2014). 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 2 displays the mean and standard deviation of the log of annual earnings over the life-cycle. Mean earnings slightly more than double between ages 25 and 55, with a noticeable flattening after age 40. The standard deviation of log earnings increases by 21.5 log points between ages 25 and 55. These patterns are broadly consistent with the results based on the PSID, as reported in (**?**), though we note that their sample selection criteria differ from ours.[11]

We compute autocorrelations for earnings using the same method that we used for hours. Figure 2c plots the results. The 1-year autocorrelation of earnings is 0.932 and declines to 0.493 at a twenty-year lag. We conclude that earnings are even more persistent than hours. Once again, these two values are inconsistent with a simple AR(1) process: An autocorrelation of 0.932 at a one-year

[10] In Appendix B.3 we examine the extent to which this property remains true if we only require twenty years of attachment. There we find that outside of the bottom ten percent of the lifetime hours distribution, differences in weekly hours worked are the most important driver of differences in lifetime hours. For workers in the bottom ten percent there is also significant variation in the number of years worked and average weeks worked.

[11] In particular, we require individuals to have at least 520 hours in all years, whereas they include any annual observation that includes at least 520 hours even if the individual associated with that observation works less than 520 hours at a different date.

Figure 2: Cross-section of Earnings Over the Life-Cycle, Ages 25 – 55

(a) Mean log Annual Earnings (b) S.D. of log Annual Earnings (c) Autocorrelation of Earnings

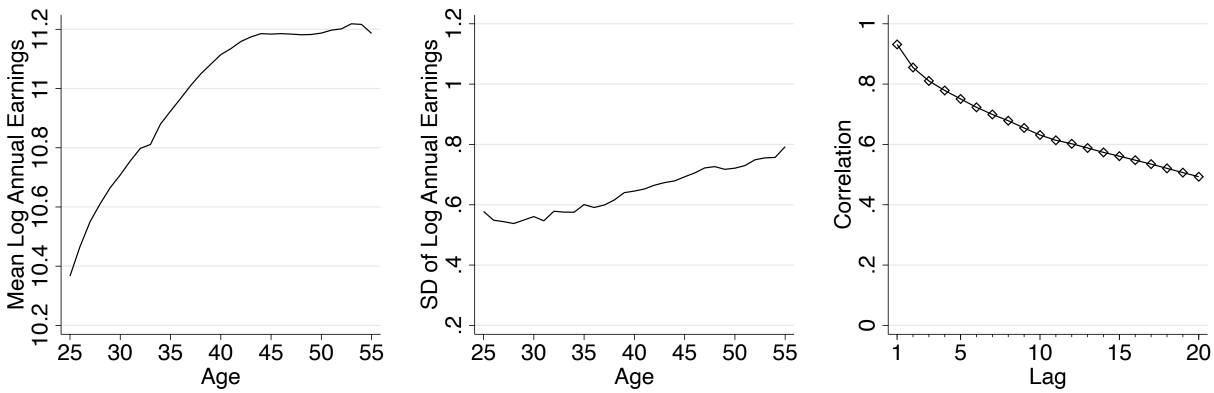

lag would imply a correlation of only around 0.245 at a twenty-year lag.[12]

As with hours, persistent earnings differences imply sizable variation in average lifetime earnings. We follow Guvenen et al. (2022) and compute (annualized) lifetime earnings as the sum of earnings between ages 25 and 55 divided by 31. We find that on average, a man in our sample earns $76624 per year from age 25-55. The standard deviation is $51828 per year, which is 67.6% of the mean. The interquartile range is $40920, and the interquartile ratio is 2.12.

2.3.3 Joint Properties of Lifetime Hours and Earnings

In this section we document some aspects of the relationship between hours worked and earnings. We begin by examining the relationship between lifetime hours and the growth rate of earnings over the life-cycle. To construct a measure of life-cycle earnings growth for individuals in our sample, we average an individual's earnings between ages $25-29$ to get a measure of initial earnings, and we average their earnings between ages $51-55$ to get a measure of final earnings. Earnings growth is then just the percent change between initial and final earnings.

Figure 3a plots average earnings growth by lifetime hours bins. The relationship is essentially flat below 1750 annual hours, though recall that very few individuals in our sample lie in this region. For hours above 1750 the relationship is increasing: Focusing on the two data points that

[12]As was the case for hours, one may worry that our imputation procedure for earnings mechanically increases the persistence in annual hours worked. Appendix Figure B.1 shows the autocorrelation profile excluding imputed values for annual earnings alongside the profile including imputed earnings. When doing this we distinguish between two time periods. Before 1994, earnings were collected annually and during this period the two profiles are virtually identical. From 1994 onwards, earnings are only reported every other year. The two-year auto-correlation in the directly reported data is slightly lower than when including imputed values. From lag 4 onwards, the two profiles again lie almost on top of each other.

Figure 3: Lifetime Hours and Earnings

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

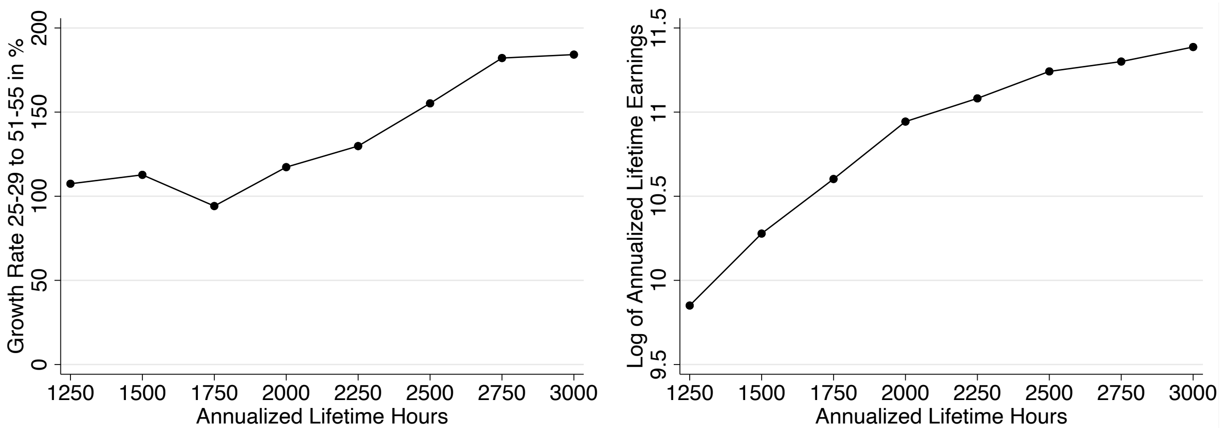

Notes: The 3000 hours bin includes anyone with 3000 or more annualized lifetime hours.

reflect the interquartile hours range, the average earnings growth of individuals who work $2,500-2,749$ hours per year are 38 percentage points higher than the earnings growth of individuals who work $2,000-2,249$ hours.

Column 1 of Table 6a formalizes this, reporting the results of regressing earnings growth on log lifetime hours. The coefficient is positive and statistically different than zero, and the implied effects are economically significant as well. In particular, the interquartile range for lifetime hours of about 20% implies a difference of almost 30 percentage points for life-cycle earnings growth.

The human capital model that we develop later in the paper will allow for heterogeneity in learning ability, which is also a potential source of heterogeneity in life-cycle earnings growth. For this reason it is of interest to assess the extent to which the previous result holds when controlling for a measure of learning ability. To pursue this we run the same regression controlling for AFQT scores, which we view as a plausible proxy for learning ability. Results are shown in Column 2 of Table 6a. While AFQT scores have a positive and statistically significant effect on life-cycle earnings growth, including them barely affects the coefficient on lifetime hours.[13]

The previous result about the relationship between lifetime hours and life-cycle earnings growth is consistent with a mechanism in which differences in lifetime hours lead to differences in human capital accumulation. The Ben-Porath model that we study in this paper will feature such a mechanism. Another perspective on this relationship comes from examining the relationship between lifetime hours and lifetime earnings. Holding wages constant, the elasticity of earnings with respect to hours is equal to unity in a model that does not feature human capital accumulation. Figure 3b shows that lifetime earnings are strongly increasing in lifetime hours. Column 1 of Table 6b

[13]Note that this sample is slightly smaller due to missing AFQT scores for some individuals.

Table 6: Lifetime Hours and Earnings Regressions

(a) Life-Cycle Earnings Growth

	(1)	(2)
log Lifetime Hours	1.58***	1.59***
	(0.35)	(0.34)
AFQT Percentile		0.02***
		(0.00)
Constant	1.37***	1.39***
	(0.05)	(0.05)
N	1418	1361
R^2	0.01	0.07

(b) log of Lifetime Earnings

	(1)	(2)
log Lifetime Hours	1.33***	1.31***
	(0.10)	(0.09)
AFQT Percentile		0.01***
		(0.00)
Constant	11.08***	11.08***
	(0.01)	(0.01)
N	1418	1361
R^2	0.12	0.30

Notes: In both regressions, log annualized lifetime hours and AFQT percentiles are demeaned such that the constants are comparable across specifications (1) and (2). Standard errors are in parentheses. *,** , and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

shows the result of regressing log lifetime earnings on log lifetime hours. The earnings-hours elasticity is 1.3 and exceeds unity at the one percent significance level. This elasticity barely changes when we include a respondent's AFQT percentile as an additional regressor in Column 2.[14]

We noted earlier that focusing on our highly attached sample may offer a conservative estimate of the role of lifetime hours given that broader groups tend to have greater dispersion in lifetime hours. We close this section by noting that the (reduced form) relationship between lifetime hours and life-cycle earnings growth becomes stronger when we apply a less stringent attachment threshold. For example, when requiring only 20 years with at least 520 annual hours, the hours coefficient for the life-cycle earnings growth regression is 1.70 without AFQT scores and 1.55 with AFQT scores. The hours coefficient for the lifetime earnings regression is 1.81 without AFQT scores and 1.61 with AFQT scores. (See Figure B.4 and Table B.4 in the Appendix.) This reinforces the idea that our strict sample selection criteria might be viewed as providing a lower bound on the role of lifetime hours worked for lifetime earnings.

[14] Figure 3b reveals that the slope of the relationship is higher below 1750 hours, a region in which there are very few observations. We have also run these regressions excluding the bins below 1750 and found that the results are only modestly affected. Without controlling for AFQT scores the coefficient on lifetime hours is 1.25, and controlling for AFQT scores it is 1.22. Both are statistically significant at the 1 percent level.

3 Lifetime Hours and Human Capital Accumulation: A First Look

As documented in the previous section, a key feature of inequality is that the cross-sectional dispersion of log earnings (and wages) increases substantially over the life-cycle. The literature has focused on two proximate causes: the accumulation of persistent shocks and heterogeneity in learning ability. In this section we argue that heterogeneity in hours worked, which the literature has largely neglected, constitutes a third relevant factor. We make this argument in the context of a standard Ben-Porath model. The key economic channel captured by this model is that individuals who expect to work longer hours throughout their lifetime have greater incentives to invest in human capital.

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. In the second subsection we highlight the mechanism linking lifetime work hours and incentives to accumulate human capital. To facilitate transparency, here we assume that total hours in each period of working life are exogenously fixed at some level and ask how changes in this level affect incentives to invest in human capital. Our quantitative model in Section 4 will endogenize total hours.

3.1 A Simple Ben-Porath Model

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

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

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

The individual is endowed with initial human capital x_1. At each age t prior to retirement the individual chooses to divide their h units of time between production (n_t) and investing in human capital (s_t). Because the individual does not value leisure, it follows that $n_t + s_t = h$ during this period of the life-cycle. Time devoted to human capital investment at age t augments the individual's human capital at age $t+1$:

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

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

A competitive labor market offers a time invariant wage rate of w per efficiency unit of production labor services, so the individual receives labor earnings equal to $n_t x_t w$ at age t. While the individual faces the same wage per efficiency unit of production labor services at all ages, the payment per unit of time spent in production is equal to $x_t w$. The individual can borrow and lend at the time invariant interest rate r. The only constraint on borrowing is that the individual cannot die with negative assets.

3.2 The Mechanics of Human Capital Investment

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

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

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

This condition can be interpreted as requiring that the marginal value of time allocated to production should equal the marginal value of time allocated to investment. The left-hand side is the effective wage for an individual at age t and thus reflects the value of a marginal increase in production time at age t holding all else constant. Turning to the right-hand side, higher investment today produces a stream of benefits in all future periods until retirement, and the overall benefit is the present value of this stream. To understand the terms in this sum, note that the term $\alpha\phi x_t^{\phi} s_t^{\phi-1}$ reflects the marginal increase in human capital at age $t+1$ as a result of a marginal increase in time devoted to investment. This investment will also increase human capital in period $t' > t+1$ by the amount $\alpha\phi x_t^{\phi} s_t^{\phi-1}(1-\delta)^{t'-(t+1)}$. The value of this additional human capital at age t' is the product of three terms: the increase in human capital at age t', the total number of hours worked at age t', 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. Note that it is total future hours that enter on the right-hand side and not only future production hours. This is because higher human capital increases productivity for both activities and the marginal value of time is equated across the two activities.

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

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

An important point highlighted by Equation (1) is that the quantitative effect of both higher learning ability and higher total hours on time devoted to investment will likely depend 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.

4 A Structural Model of Lifetime Hours and Earnings

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

4.1 Model

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

A key element of our analysis is to extend the basic Ben-Porath model from the previous section to account for the differences in lifetime hours across individuals. As documented in Bick, Blandin and Rogerson (2022), observables explain very little of the variation in hours worked across individuals. This leads us to allow for heterogeneity in preferences. Our specification of preference heterogeneity is in turn motivated by the properties of the autocorrelation of hours documented in Figure 1c. The fact that the autocorrelation declines over time suggests a transitory component to preference heterogeneity, while the fact that the autocorrelation plateaus above zero at longer lags suggests a permanent component. Previous research has allowed for either permanent preference heterogeneity (e.g., Kaplan (2012), Heathcote, Storesletten and Violante (2014), Keane and Wasi (2016) and Bick, Blandin and Rogerson (2022)) or transitory preference heterogeneity (e.g., Imai and Keane (2004) and Chang et al. (2020)), but not both. We show below that allowing for both is quantitatively important.

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

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

where $c_{i,t}$ is consumption at age t, $h_{i,t}$ is total time devoted to work at age t, ψ_i is an individual-specific time-invariant preference shifter, and $\pi_{i,t}$ 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 $\pi_{i,t}$ is assumed to follow:

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

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

where the innovations $\nu_{i,t+1}$ are assumed to be iid over time and across individuals. The parameters β, σ, and γ satisfy $0 < \beta < 1$, $\sigma > 0$, and $\gamma > 0$ and are common to all individuals.

As in Section 3, total hours of work at age t for individual i, $h_{i,t}$, are allocated between producing ($n_{i,t}$) and investing in human capital ($s_{i,t}$); but, following Huggett, Ventura and Yaron (2011), the human capital accumulation process is now specified as:

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

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

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

This shock is iid over time and across individuals. Note that while the shock $z_{i,t}$ 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 (ψ_i and α_i) and two initial conditions ($\pi_{i,1}$ and $x_{i,1}$). In our quantitative analysis we assume that $x_{i,1}$ and ψ_i are joint lognormally distributed:

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

We assume that α_i is lognormally distributed, and to economize on the size of our state space we assume that $x_{i,1}$ is perfectly correlated with α_i.[15] While our analysis does not explicitly model occupations and the potential for human capital accumulation opportunities to vary across occupations, we view the heterogeneity in learning ability α_i as reflecting both the innate ability of an individual to learn as well as the learning abilities inherent in the occupational choice that best suits the individual's innate skills. We impose that the initial distribution for $\pi_{i,1}$ is the ergodic distribution for the π_t process. Because the idiosyncratic shocks for this process are uncorrelated to all other variables, we assume that $\pi_{i,1}$ is uncorrelated to all other variables.

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

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

The parameter $\tau_0 \in [0,1)$ determines the overall level of the income tax, while the parameter $\tau_1 \in [0,1)$ determines the progressivity of the tax. To minimize the state space, we model a simplified Social Security system. In particular, we assume that all individuals in our model receive a transfer from the government equal to I_{ss} in each period beginning in period T_R. Because our quantitative analysis will focus on the highly attached male sample described in Section 2 that is eligible for very few transfers, we abstract from all transfers other than Social Security.

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

[15] Guvenen, Kuruscu and Ozkan (2014) also assumes that learning ability and initial human capital are perfectly correlated. 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.

constant productivity growth that is labor augmenting.

We let w_t denote the wage per efficiency unit that the cohort faces at age t. An individual of age t with human capital x that supplies n units of time to production earns pre-tax labor income nxw_t. There are no state contingent securities, but individuals are able to borrow and save at the interest rate r subject to the natural borrowing constraint. An age t individual with human capital x_t and assets k_t who is not retired has the following period budget equation:

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

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

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

When connecting our model with the data we assume that log hours and log earnings are subject to classical measurement error with standard deviations of σ_{mh} and σ_{me}, respectively.

4.2 The Individual's Problem

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

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

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

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

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

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

5 Calibration

In this section we describe how we calibrate our model to match salient features of earnings and hours over the life-cycle and report on the ability of the calibrated model to match targeted and untargeted moments.

5.1 Calibration Procedure

Our model features 24 parameters: three price parameters (r, w_1, and g_w), three common preference parameters (β, σ, and γ), three technology parameters for the human capital accumulation process (δ, ϕ, and σ_z), two 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), four parameters characterizing the tax and transfer system (τ_c, τ_0, τ_1, and I_{ss}), and five parameters characterizing the joint distribution of x_0 and ψ (μ_x, μ_ψ, σ_x, σ_ψ, and $\rho_{x,\psi}$).

Ten parameters are set externally and summarized in Table 7. We normalize the parameter $\mu_x = 0$ since it is not identified separately from w_1. Setting the period length equal to a year, we set $r = 0.02$ and $\beta = 1/(1+r) = .9804$. We adopt commonly used values for the two curvature parameters σ and γ: $\sigma = 1$, and $\gamma = 0.3$. Tax function parameters are set according to the estimates in Heathcote, Storesletten and Violante (2017): $\tau_0 = 0.81$ and $\tau_1 = 0.181$, and the consumption tax τ_c is set to 0.07. Following Huggett, Ventura and Yaron (2011), we set $\delta = 0.02$. We set $g_w = 0.005$ based on average growth of wages for males aged 25-29 between 1981 and 2018 in CPS data. Details of this calculation are provided in Appendix B4.

For the remaining fourteen parameters we choose values using a simulated method of moments procedure in which we choose parameters such that when we solve the model it matches the value of fourteen moments.[16] We first list the fourteen moments that we target and later provide some intuition about the connection between model parameters and these moments.

Because the distributions of hours and earnings are the core outcomes of interest in our analysis, we target moments that describe these distributions. Specifically, we target five moments related to hours, six moments related to earnings, and two moments related to joint properties of earnings and hours. The five hours moments that we target are the overall mean and standard deviation of log hours for the pooled sample of all individuals between ages 25 and 55, and values for the autocorrelation of individual hours at lags of 1, 10, and 20 years. The six earnings moments that

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

Table 7: Externally Calibrated Parameter Values

Parameter	Interpretation	Value	Source
β	Patience	0.9804	Huggett, Ventura and Yaron (2011)
r	Interest rate	0.02	$1/\beta$
σ	CRRA	1.0	–
γ	Frisch elasticity	0.3	–
δ	Human capital depreciation	0.02	Huggett, Ventura and Yaron (2011)
τ_0	Tax Rate	0.81	Heathcote, Storesletten and Violante (2014)
τ_1	Tax Progressivity	0.181	Heathcote, Storesletten and Violante (2014)
τ_c	Consumption Tax	0.07	McDaniel (2007)
μ_x	Mean of $\log x_0$	0.0	Normalization
g_w	Exogenous wage growth rate	0.005	CPS

we target are mean log earnings at ages 30 and 50, the standard deviation of log earnings at 30 and 50, and the autocorrelation of log earnings at lags of 1 and 20 years. The two joint hours and earnings moments that we target are the correlation between hours and earnings at age 30, and the slope of the relationship between lifetime hours and life-cycle earnings growth, as depicted in Figure 3a. The final moment that we target is the ratio of the Social Security transfer to mean annual earnings.

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

Given values for all of the other parameters, the five moments of the hours distribution can be used to pin down values for μ_ψ, σ_ψ, ρ_π, σ_π, and σ_{mh}. Intuitively, the value of μ_ψ is tightly linked to the cross-sectional mean of annual hours. Each of the other four parameters influences both the variance of log annual hours as well as the shape of the hours autocorrelation profile. In particular, measurement error, transitory preference heterogeneity, and permanent preference heterogeneity will each generate very different profiles for the hours autocorrelation function. Loosely speaking, the autocorrelation function in the full model will be a weighted average of the these profiles, with the weights influenced by the relative variances of the three components. The overall variance of annual hours is affected by the scale of each of the three variances. Similar logic implies that holding all else constant, the values for σ_z and σ_{me} can be chosen to target values of the autocorrelation function for log earnings.

Holding other parameters fixed, w_1, σ_x, μ_α, and σ_α will impact properties of the earnings distribution. The values of w_1 and σ_x will impact the level and variance of earnings for young

Table 8: Internally Calibrated Parameter Values

Parameter	Interpretation	Value	Moment
w_1	Wage in first period	25.5	Mean log earnings, age 30
σ_x	SD of $\log x_0$	0.325	SD log earnings, age 30
μ_α	Mean of $\log \alpha$	-2.3801	Mean log earnings, age 55
σ_α	SD of $\log \alpha$	0.2	SD log earnings, age 55
μ_ψ	Mean of $\log \psi$	3.8234	Mean log annual hours, age 25-55
σ_ψ	SD of $\log \psi$	0.525	SD log annual hours, age 25-55
$\rho_{\alpha,\psi}$	Corr. of $(\log \alpha, \log \psi)$	-0.15	Correlation of hours and earnings, age 30
σ_π	SD of $\log \pi$	0.4	Hours autocorrelation profile
ρ_π	Autocorrelation of $\log \pi$	0.88	Hours autocorrelation profile
σ_{mh}	SD measurement error	0.1	Hours autocorrelation profile
σ_{me}	SD measurement error	0.17	Earnings autocorrelation profile
σ_z	SD human capital shock	0.1	Earnings autocorrelation profile
ϕ	HC elasticity wrt investment	0.57	Lifetime hours, earnings growth
I_{ss}	Social Security benefit	0.188	$0.4\bar{e}$

workers, μ_α will influence mean life-cycle earnings growth, and σ_α will influence the extent to which earnings become more dispersed with age. Because initial human capital impacts earnings of young individuals and the value of ψ impacts hours, the correlation parameter $\rho_{x,\psi}$ will influence the correlation of earnings and hours for young workers. We note that because learning ability and initial human capital are perfectly correlated, the value of $\rho_{x,\psi}$ also controls the correlation between learning ability and tastes for work.

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

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

Table 8 displays the calibrated parameter values. Here we note four features of the calibrated values. First, although we use a novel moment to help identify the parameter ϕ, the calibrated value of 0.57 is within the broad range of estimates found in the literature, though toward the lower end.[18]

Second, we find a modest negative correlation between permanent tastes for work and initial human capital, with $\rho_{x,\psi}$ equal to -0.15. Because we impose that learning ability α and initial human capital are perfectly correlated, our calibration also implies a modest negative correlation between permanent disutility for work and learning ability. We previously argued that high learning ability and high hours are both sources of higher life-cycle earnings growth. The relatively weak correlation between ψ and α implies that these two channels operate somewhat independently of each other.

Third, the transitory component of preference heterogeneity is very persistent and contributes substantially to the cross-sectional variation in tastes for work. Specifically, the variance of $\log \pi$ in the ergodic distribution of the transitory process is equal to 0.709, which is more than twice as large as the variance in the log of the permanent component ψ, which is equal to 0.276. Our calibration procedure also implies a substantial amount of measurement error in hours and earnings, with $\sigma_{mh} = 0.10$ and $\sigma_{me} = 0.17$, respectively.

Fourth, our procedure yields a value of 0.10 for σ_z. This is somewhat smaller than the value of 0.11 used by Huggett, Ventura and Yaron (2011).[19] Our estimate comes from requiring our model to match the autocorrelation properties of log earnings at long lags, whereas their estimate was

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

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

[19] It is relevant to note that our time frame and sample selection differ from theirs. They present estimates of the shock process variance for several different samples and find that imposing a higher minimum earnings threshold does lower the estimated variance. They also find that the estimated variance is slightly lower for the later part of their sample period, which is the relevant period for our sample.

based on the stochastic properties of wage changes at short lags for older individuals. A robust finding from the literature on job loss is that older workers are less likely to experience job loss but that the consequences of job loss are much larger for them than younger workers.[20] For this reason we prefer not to impose that measured shocks for older workers are a good estimate of the average shocks faced by workers of all ages. If we impose $\sigma_z = 0.11$ and do not require our model to match the autocorrelation properties of log earnings at longer lags, we find that the model exhibits too little persistence in log earnings at lags longer than five years.

5.2 Fit of the Calibrated Model

In this subsection we report on the ability of the model to match both targeted and untargeted moments. We begin by examining the model's ability to fit the targeted moments, beginning with moments related to earnings. Results are shown in the three panels of Figure 4. Overall, the model closely tracks the evolution of the mean and standard deviation of log earnings over the life-cycle, as well as the persistence of earnings. As a reminder, for the mean and standard deviation, we targeted the moments at ages 30 and 50, and for the autocorrelation we targeted lags 1 and 20. One small discrepancy to note is that the model implies a slightly concave profile for the standard deviation, whereas the profile in the data is slightly convex. This same pattern is present in the calibrated model of Huggett, Ventura and Yaron (2011); see their Figure 2(b).[21] We also note that the mean of log earnings in the model for ages 25-30 is higher than in the data, which – as we will discuss momentarily – stems from a similar discrepancy between hours in the model and the data.

Next we consider the moments for hours. Results are shown in the three panels of Figure 6. Panel (a) shows the age profile for log mean hours, and Panel (b) shows the age profile for the standard deviation of log hours. Recall that our calibration procedure targeted the overall cross-sectional mean and variance of log hours, but not the values for any particular age. In both the model and the data these profiles are relatively flat between the ages of 30 and 50, so matching the overall sample value leads to a reasonable fit to the life-cycle profiles. In Section 6 we will show that preference heterogeneity plays a key role in allowing the model to generate the amount of dispersion in hours found in the data.

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

[20] See, for example, the results and discussion in Couch and Placzek (2010). White (2010) documents this fact using the NLSY.

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

Figure 4: Model Fit for Earnings

(a) Mean of log Earnings (b) S.D. of Log Earnings (c) Persistence of Earnings

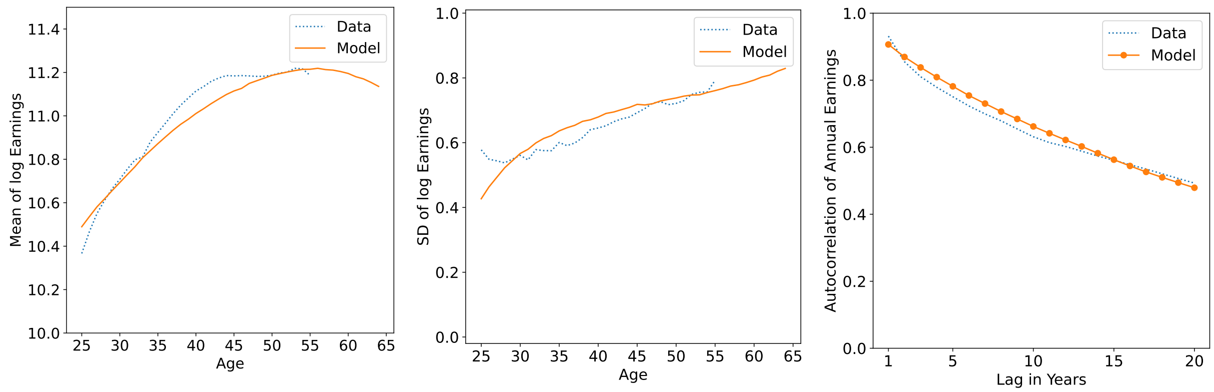

Figure 5: Model Fit for Hours

(a) Mean of log Hours (b) S.D. of log Hours (c) Persistence of Hours

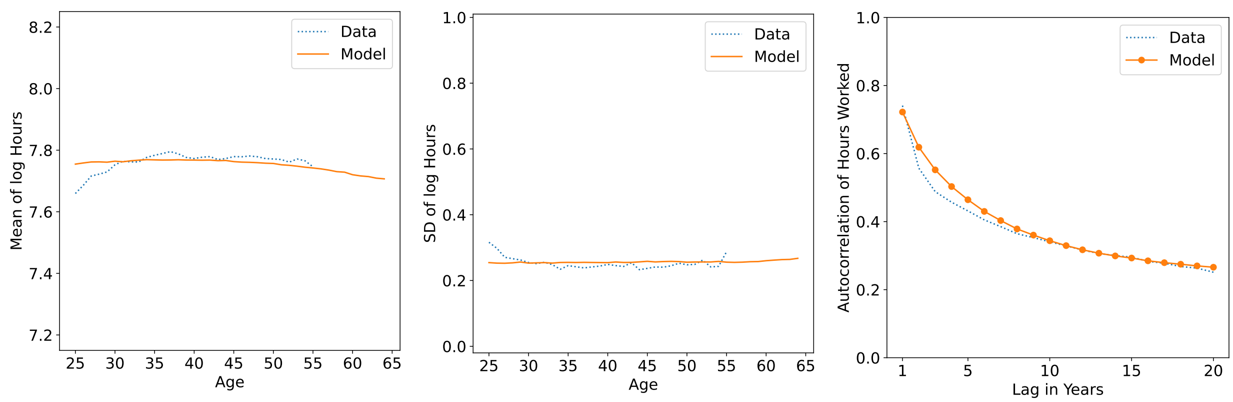

the data over this age range. These two properties are intimately related: the increased prevalence of individuals with relatively low annual hours of work in this age range tends to both decrease the mean and increase the variance. We offer two potential rationalizations for this discrepancy. The first is that individuals in this age range are more likely to have spells of unemployment as they search for a good match, a feature that our model abstracts from. (See Kaplan (2012) for an analysis that incorporates this feature.) The second is our assumption that the data on hours reflect both production and investment time. This age range is the period of highest investment in human capital, so if some investment time is not included in reported work hours in the data, our model would be expected to overestimate work hours at young ages.[22]

Panel (c) shows the autocorrelation function for hours. We targeted values at lags of one, ten, and twenty, and Panel (c) shows that the model does a good job of capturing the entire profile.

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

Figure 6: Model Fit for the Relationship between Earnings and Hours

(a) log Earnings-log Hours Correlation (b) Lifetime Hours and Mean Earnings Growth (c) Lifetime Hours and Median Earnings Growth

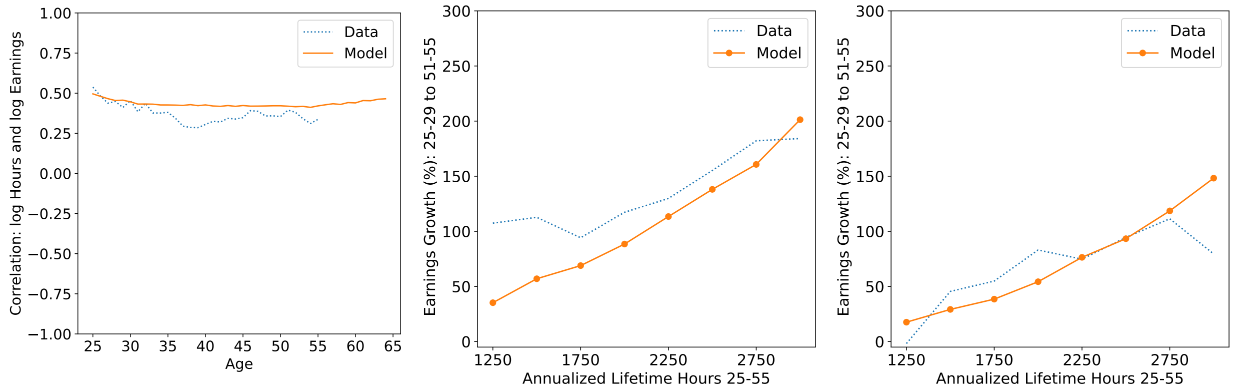

Next we consider the two moments that involve relationships between earnings and hours. Figure 6a shows the age profile for the correlation between earnings and hours in the model and in the data. The only targeted value was the correlation at age 30. Both in the data and in the model the age profile for this correlation displays a very modest downward drift, though in the data there is an additional dip between ages 35 and 45.

Figure 6b shows the relationship between lifetime hours and life cycle earnings growth in the model and in the data. Our calibration targeted the slope of this profile for annualized lifetime hours over the range of 1750-2750. While the profile in the model lies slightly below the profile in the data, the slope of the model profile does a good job of tracking the slope of the empirical profile over this range. We note that sample sizes outside of the range 1750-2750 tend to be small and that mean growth rates can be heavily influenced by outliers. To further assess this, Figure 6c shows the same relationship when using median earnings growth rather than mean earnings growth.

We now turn to moments that were not explicitly targeted by our calibration procedure. We did not explicitly target any moments based on wages, but given that we capture the patterns for earnings and hours individually and do a reasonable job of matching the profile for the correlation between earnings and hours, it is not surprising that our model does a reasonable job of matching the properties of wages. The three panels of Figure 7 display the results. Panel (a) shows that the model closely tracks the evolution of mean log wages over the life-cycle. Panel (b) shows the model slightly understates the standard deviation of wages between the ages 25 and 30 but fits the profile quite well between ages 30 and 55. The understatement of the standard deviation at young ages results from the fact that the model understates the standard deviation of hours over this age range. Panel (c) displays the correlation between hours and wages over the life-cycle. Consistent with our results for earnings, the model captures the fact that this correlation is near zero, but the model generates a slightly larger positive trend over the life-cycle than is present in the data.

Figure 7: Model Fit for Wages

(a) Mean of log Wages (b) S.D. of log Wages (c) log Wage-log Hours Corr.

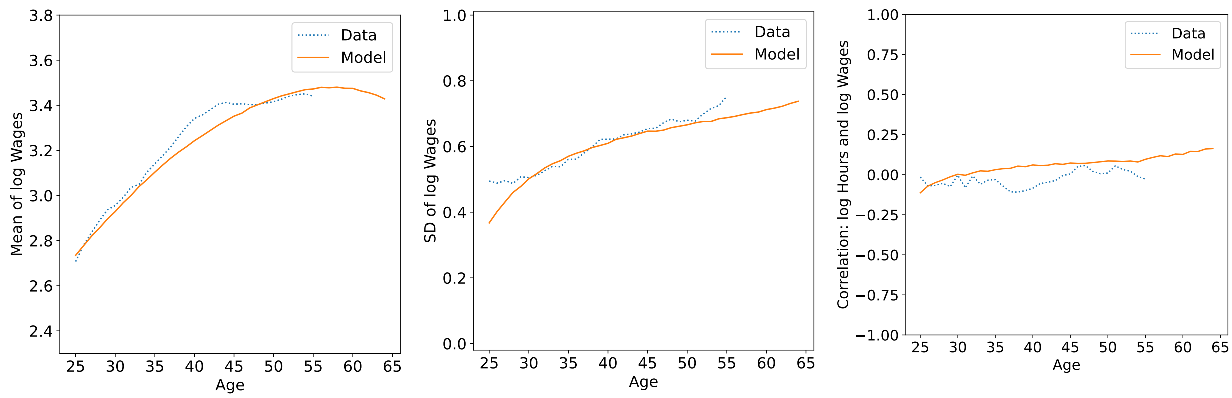

It is noteworthy that our model closely captures the properties of both earnings and wages, as these measures are of independent interest. When studying consumption and wealth inequality, inequality of earnings is of primary importance; but, when studying inequality of opportunity, it is inequality in wage rates that is of particular interest. Notably, Huggett, Ventura and Yaron (2011) studied only properties of earnings, and Guvenen, Kuruscu and Ozkan (2014) studied only properties of wages.

Figure 8 displays the distribution of annualized lifetime hours and annualized lifetime earnings in both the model and the data. Neither of these lifetime distributions were explicitly targeted. For earnings, we targeted the cross-sectional variance at ages 30 and 50 and two values of the autocorrelation function. For hours, we targeted the mean and variance for log annual hours in the overall cross-section and three values of the autocorrelation function.

The main observation from Panel (a) is that our model does not generate sufficient concentration in the middle of the distribution. This reflects the well-known issue that models with log normally distributed heterogeneity cannot generate the high concentration found in the distribution of annual hours worked. Simply put, the large spike of individuals who report 2000 annual hours (40 hours a week for 50 weeks) cannot be well approximated with a normal distribution. Although the concentration in annualized lifetime hours is less severe, the issue is still present. While we could generate additional concentration in the hours distribution by introducing non-linearities into the mapping from hours to earnings, as in Bick, Blandin and Rogerson (2022), we have opted not to do so in order to better focus on other features.

Turning to lifetime earnings in Panel (b), we highlight two discrepancies between model and data. First, it is again true that the model does not generate the level of concentration found in the data, though the extent of the discrepancy is more modest compared to the case of hours. This is intuitive; because there is substantial heterogeneity in wages for workers in the hours bin

Figure 8: Model Fit of Lifetime Earnings and Hours

(a) Annualized Lifetime Hours (b) Annualized Lifetime Earnings (c) Annual. Lifet. Earn. & Hours

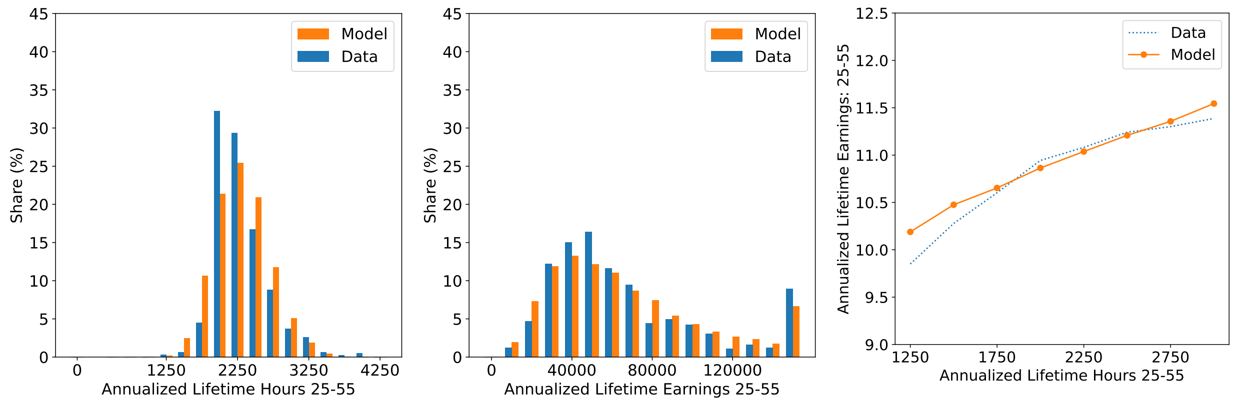

Table 9: Model Fit of Lifetime Inequality

	Var of log Liftime	
	Earnings	Hours
Data	0.314	0.021
Model	0.322	0.025

containing 2000, the concentration in hours yields much less concentration in earnings. Second, our model does not generate sufficient mass in the right tail of the distribution. This is a well-known issue in the literature for models like ours with log normal shocks. Panel (b) confirms that we do not have sufficient mass in the final bin, which corresponds to individuals with annualized lifetime earnings greater than $150,000.

Table 9 complements these figures by showing the variance of log lifetime earnings and hours in the model and the data. Consistent with the previous discussion, it shows that the lack of concentration leads our model to generate too much variance for both earnings and hours, though the difference for earnings is smaller. Although our model generates a variance of log lifetime hours that is too high relative to the data, it generates an interquartile gap of 1.24, which is very close to the gap of 1.20 in the data, suggesting that the lack of concentration in the middle of the distribution is the main source of the discrepancy in variances.

Panel (c) of Figure 8 shows how average annualized lifetime earnings vary across bins of the annualized lifetime hours distribution. Although not targeted, the model and data profiles track each other quite closely, except at the lowest and highest hours bins. The discrepancy arises because the model-based relationship is fairly linear, whereas the relationship in the data is slightly concave. This suggests that it might be reasonable to introduce some additional non-linearities to the investment production function at both low and high hours.

Lastly, we examine whether the model generates a reasonable joint relationship between lifetime earnings, lifetime hours, and ability. Recall that in Section 2 we ran a regression of log lifetime earnings on log lifetime hours and AFQT scores, which we suggested was a reasonable proxy for learning ability in our model. (Recall that in our model, learning ability is perfectly correlated with initial human capital.) When we run the analogous regression in the model, replacing AFQT scores with learning ability, we get coefficients of 1.20 on lifetime hours and 0.01 on learning ability, compared with 1.31 and 0.01, respectively, in the data regression reported in Table 6b.

6 Sources of Lifetime Earnings Inequality

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

6.1 Heterogeneity in Hours of Work and Earnings Inequality

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

Results for this counterfactual are reported in Table 10. The first row repeats statistics for our benchmark model, and the results in the second row show that exogenously imposing equal hours leads to a reduction of 19 percent in the variance of log lifetime earnings. This counterfactual also reduces the growth in earnings inequality over the life-cycle: In our benchmark calibration, the cross-sectional variance of log earnings increases by 0.267 from age 30 to 55, whereas in this

Table 10: The Role of Hours Heterogeneity

	log Lifetime Earnings		Δ_{55-30} log Earnings	
	Var.	% of BM	Var.	% of BM
Benchmark (BM)	0.322	100%	0.258	100%
Counterfactual: $n_{i,t} + s_{i,t} = \bar{h}$	0.261	81.0%	0.191	74.2%

Notes: The first row reports statistics from our benchmark model over the age range 25-55. The second row reports the same statistics for the counterfactual economy where total hours worked are restricted to be the same across all individuals and ages. Columns labeled % of BM report the percent of the benchmark values that remain in the counterfactual.

counterfactual the increase falls to 0.191, a drop of more than 25 percent.

Notably, the results in Table 10 reveal that the drop in the variance of log lifetime earnings, 0.061, is much larger than the drop in the variance of log lifetime hours, 0.025. In a static textbook model of labor supply, hours heterogeneity for otherwise identical workers directly leads to heterogeneity in earnings. As highlighted in Section 3, a key feature of the Ben-Porath model is that heterogeneity in hours will also lead to heterogeneity in human capital accumulation, which will in turn amplify the direct effect of hours heterogeneity on earnings heterogeneity. This raises the question, how important human capital accumulation is for the lifetime earnings effects reported in Table 10?

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

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

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

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

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

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

Table 11: The Role of Human Capital

	Var.	% of BM
Benchmark (BM)	0.322	100%
Counterfactual: $n_{i,t} + s_{i,t} = \bar{h}$	0.261	81.0%
Direct Channel	0.279	86.7%
Human Capital Channel	0.290	90.1%

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

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

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

Results of this decomposition exercise are reported in Table 11. The direct channel delivers a 13.3 percent decrease in the variance of log lifetime earnings, while the human capital channel delivers a 9.9 percent decrease. The sum of these two channels is greater than the total effect. This is perhaps intuitive since time devoted to production and investment are complements in terms of generating lifetime earnings.

Assessing the contribution of the human capital channel requires that one take a stand on how to assign the interaction effects. For our headline number we assign the interaction effects proportionately. If we do this, the share of the overall decline in earnings inequality from removing hours heterogeneity that is due to the human capital channel is 43 percent $\left(= \frac{100-90.1}{(100-86.7)+(100-90.1)}\right)$. We can also generate an interval of values by considering the two extremes in which we assign all of the interaction effects to either the direct or human capital channels. If we assign all of the interaction effects to the human capital channel, then the human capital channel is assigned everything not accounted for by the direct channel, which implies a 30 percent $\left(= \frac{86.7-81.0}{100-81.0}\right)$ share for the human capital channel. If we assign all of the interaction effects to the direct channel, this figure becomes 52 percent $\left(= \frac{100-90.1}{100-81.0}\right)$. Focusing on either our headline number of 43 percent or the range of 30 to 52 percent, we conclude that the human capital channel is quantitatively important.

Table 12: The Role of Preference Heterogeneity

	Annual				Lifetime			
	log Earnings		log Hours		log Earnings		log Hours	
	Var.	% of BM	Var.	% of BM	Var.	% of BM	Var.	% of BM
Benchmark (BM)	0.481	100%	0.065	100%	0.322	100%	0.025	100%
$\sigma_\psi = \sigma_\pi = 0$	0.389	81.0%	0.010	15.9%	0.267	82.8%	0.001	2.4%
$\sigma_\psi = 0$	0.437	90.9%	0.048	74.0%	0.278	86.3%	0.008	31.7%
$\sigma_\pi = 0$	0.432	89.8%	0.027	41.8%	0.309	96.1%	0.018	70.3%

Notes: The first row reports statistics from our benchmark model over the age range 25-55. The subsequent rows report the same statistics under various counterfactuals that eliminate one or more sources of preference heterogeneity. Columns labeled % of BM report the percent of the benchmark values that remain in each counterfactual.

6.2 The Role of Preference Heterogeneity

The previous subsection showed that heterogeneity in life-cycle hours profiles across individuals is an important contributor to lifetime earnings inequality. A central feature of our model is the introduction of preference heterogeneity in order to capture the salient features of hours heterogeneity. But preference heterogeneity is not the only source of hours heterogeneity; heterogeneity in initial human capital, heterogeneity in learning ability, and heterogeneity in the realization of human capital shocks might also generate important differences in time allocations across individuals. In this subsection we assess the contribution of preference heterogeneity. In Section 6.3 we consider the contribution of other sources.

6.2.1 Preference Heterogeneity and Lifetime Earnings Inequality

To assess the contribution of preference heterogeneity to lifetime earnings inequality, we consider a counterfactual in which we eliminate preference heterogeneity in our calibrated model by setting $\psi_i = \mu_\psi$ and $\pi_{i,t} = 1$ for all i and t. This counterfactual corresponds to the specification in Guvenen, Kuruscu and Ozkan (2014), as they had an endogenous choice of total hours but homogeneous, time invariant preferences. We then re-solve the model, assuming that each individual experiences the same sequence of shocks to human capital and measurement error as in the benchmark economy. The second row of Table 12 shows the results from this exercise. To facilitate comparison, the first row repeats the results for our benchmark model.

Not surprisingly, this counterfactual dramatically reduces the overall cross-sectional variation in log annual hours: Pooling all observations for individuals between ages 25 and 55, the variance of log annual hours drops from 0.065 in our benchmark model to 0.010 in the counterfactual economy without preference heterogeneity, a drop of almost 85 percent. Importantly, although the model without preference heterogeneity still generates a non-trivial amount of heterogeneity in hours across individuals, it features virtually no dispersion in lifetime hours. In the benchmark model the variance of log lifetime hours is 0.025, and in the counterfactual exercise it decreases to 0.001.[23] The amount of heterogeneity in annual hours that remains in our model when we eliminate preference heterogeneity is in line with the amount of hours dispersion generated in the calibrated model of Guvenen, Kuruscu and Ozkan (2014), which as noted earlier assumed no preference heterogeneity. Because they did not study the properties of lifetime hours, our result about the almost total elimination of variance in lifetime hours is novel.

Eliminating preference heterogeneity leads to a 17.2 percent decrease in the variance of log lifetime earnings (from 0.322 to 0.267). This is 90 percent of the effect reported in Table 10. Importantly, it remains true that the drop in the variance of log lifetime earnings is much larger than the drop in the variance of log lifetime hours. Specifically, the drop in the variance of log lifetime earnings (0.051) is roughly twice the drop in the variance of log lifetime hours (0.024).

The previous counterfactual eliminated both permanent and transitory preference heterogeneity. It is also of interest to examine the relative importance of these two components of preference heterogeneity. To do this, we repeat the previous exercise but eliminate each component separately. Results are reported in the third and fourth rows of Table 12.

When we eliminate only permanent preference heterogeneity, the variance of log lifetime earnings decreases by 12.9 percent, whereas when we eliminate only transitory preference heterogeneity, the variance of log lifetime earnings decreases by 3.9 percent. It follows that permanent preference heterogeneity is more than three times as important as a source of lifetime earnings inequality, despite the fact that the variance of $log\,\pi$ in its ergodic distribution is more than twice as large as the variance of $log\,\psi$.

Permanent preference heterogeneity is also more important as a source of lifetime hours inequality. Shutting down permanent preference heterogeneity reduces the variance of log lifetime hours by almost 70 percent, whereas shutting down transitory preference heterogeneity reduces it by roughly 30 percent. These effects on lifetime hours inequality are quite distinct from the effects on dispersion in annual hours for the pooled sample of all annual observations between ages 25 and 55. Eliminating permanent preference heterogeneity reduces the variance of pooled, cross-sectional log annual hours by roughly 25 percent, whereas eliminating transitory preference heterogeneity reduces it by almost 60 percent. The key point is that some of the differences in hours due to transitory preference shocks tend to average out over the life-cycle, so transitory preference shocks are very important for the variation in annual hours but much less important for the

[23] We have also carried out an exercise in which we study a model without preference heterogeneity. We calibrate it following the same procedure as in Section 5, but dropping all of the moments that include hours, and setting the values of ϕ and $\sigma_{m,h}$ equal to their values in our benchmark calibration. Consistent with the results just reported, this model misses dramatically on the overall cross-sectional variance in log hours and generates virtually none of the variance in log lifetime hours found in the data.

Table 13: The Role of Human Capital

	Var.	% of BM
Benchmark (BM)	0.322	100%
$\sigma_\psi = \sigma_\pi = 0$	0.267	82.8%
Direct Channel	0.279	86.5%
Human Capital Channel	0.296	91.7%
$\sigma_\psi = 0$	0.278	86.3%
Direct Channel	0.291	90.2%
Human Capital Channel	0.299	92.8%

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

variation in lifetime hours.

The amplification effect on lifetime earnings is also larger for the permanent component. The ratio of the change in the variance of log lifetime earnings to the change in the variance of log lifetime hours is 2.59 $\left(=\frac{0.322-0.278}{0.025-0.008}\right)$ for permanent preference heterogeneity and 1.86 $\left(=\frac{0.322-0.309}{0.025-0.018}\right)$ for transitory heterogeneity.

These results highlight the importance of capturing the extent to which differences in lifetime hours are due to permanent versus transitory preference heterogeneity, or, equivalently, the importance of matching the properties of the autocorrelation function for hours. We will return to this point in subsection 6.2.3.

6.2.2 The Role of the Human Capital Channel

It is again of interest to examine the relative importance of the direct and human capital channels introduced in the previous section. To do this, we repeat the decomposition exercise from Section 6.1. Results are presented in Table 13 for two counterfactuals: one where we eliminate all preference heterogeneity, and another where we eliminate only permanent preference heterogeneity.

The sum of the direct and human capital channels exceeds the total effects in both panels so that assessing the role of human capital again requires taking a stand on how to assign the interaction terms. Carrying out the same calculations as in the previous section, we arrive at a headline number of 38 percent when shutting down both permanent and transitory components.[24] While somewhat

[24]Considering the two extreme cases for assigning the interaction effect yields an interval of values ranging from 22 to 48 percent.

smaller than the headline number in Section 6.1, it remains true that the human capital channel is quantitatively significant.

If we repeat this exercise for the case in which we eliminate only the permanent component of preference heterogeneity, we arrive at a headline number of 42 percent.[25] This value is larger than the one computed when both components of preference heterogeneity were eliminated. This finding is intuitive: Human capital accumulation is most important for young individuals, and, as highlighted in Section 3, it is heterogeneity in expected future hours of work that create heterogeneity in the incentives to accumulate human capital. And heterogeneity in expected future hours of work is most heavily influenced by the permanent component of preference heterogeneity.

Another perspective on these results considers how eliminating permanent preference heterogeneity affects the two different components of time allocation: production and investment. Specifically, we can produce separate annualized measures of lifetime hours devoted to production and investment, analogous to how we defined lifetime (total) work hours previously. When we do this we find that, relative to the benchmark model, eliminating permanent heterogeneity reduces the variance of log lifetime production hours by 31.7 percent and the variance of log lifetime investment hours by 52.2 percent.

6.2.3 The Importance of Modeling Permanent and Transitory Preference Heterogeneity

The results presented in the previous subsections establish two features of our calibrated model. First, whereas the transitory component of preference heterogeneity is the dominant source of the variance in annual hours, the permanent component of preference heterogeneity is the dominant source of variance in lifetime hours. Second, for understanding the sources of variance in lifetime earnings, the variance in lifetime hours is more important than variance in annual hours. Our analysis in Section 3 provides intuition for this second finding.

These results suggest that it is important to model the different sources of hours variation. To highlight the quantitative significance of this message, we consider how our results would be affected if we had instead started with a model that only allows for permanent preference heterogeneity. That is, we consider a model that does not include the shocks to π and carry out the same calibration procedure as described earlier, with two modifications. First, we set σ_{mh} equal to its value in our benchmark model. Second, we no longer target any of the values for the hours autocorrelation function. These modifications reduce both the number of moments being targeted and the number of parameters that need to be assigned by three.

Implementing this change results in effects that are are dramatically larger than in our bench-

[25]The interval of values from considering the extreme cases for assigning the interaction effect yield an interval from 28 to 53 percent.

mark model. If we compute the effect of removing all hours heterogeneity, as in Section 6.1, the reduction in the variance of log lifetime earnings is now 0.111, versus 0.061 in the benchmark model. And when we eliminate all preference heterogeneity, as in Section 6.2.2, the reduction in the variance of log lifetime earnings is 0.110, versus 0.055 in the benchmark model. That is, the effect of preference heterogeneity on lifetime earnings inequality is doubled, and preference heterogeneity accounts for more than 99 percent of the effect of all hours variation.

The reason the effects become so much larger is that this specification results in too much variation in lifetime hours relative to the data. In our benchmark calibration, the variance of the log of lifetime hours was equal to 0.025, but in this alternative exercise it more than doubles to 0.058. Both models match the variance in log annual hours, but with only permanent heterogeneity the persistence of hours is higher, so a given variance in log annual hours translates into a much higher variance of log lifetime hours.

One reason this exercise is of interest is that several papers only incorporate permanent preference heterogeneity yet target the cross-sectional variance of log annual hours found in the data. The results of this exercise suggest that using permanent preference heterogeneity to match the variance in annual hours will generate too much variation in lifetime hours and overstate the importance of preference heterogeneity for lifetime earnings inequality.

6.3 Other Sources of Inequality

The previous subsections have focused on the role of heterogeneous time allocations and preference heterogeneity as sources of lifetime earnings inequality, as these are the dimensions that are novel in our analysis relative to the literature. In this subsection we examine the contribution of other sources of heterogeneity to lifetime earnings inequality and compare our results with those in Huggett, Ventura and Yaron (2011).[26]

Results for lifetime measures are reported in Table 14. (Table C.5 in the Appendix provides results for lifetime and annual measures.) The entries in this table represent the share of overall variance remaining after the indicated channels are eliminated. For comparison, the first three rows repeat the earlier results of eliminating preference heterogeneity, as shown in Table 12. When we eliminate human capital shocks in our model, the variance of log lifetime earnings decreases by 40.4 percent (see row 4, column 1). This is roughly the same as in Huggett, Ventura and Yaron (2011). This is perhaps not too surprising given that our shock process has a variance that is quite close to theirs, but the implicit implication is that an endogenous hours margin does little to affect

[26]While we think it is of interest to compare our findings with those in Huggett, Ventura and Yaron (2011), we remind the reader that our analyses are not strictly comparable due to differences in sample selection and other details. For example, they start their model at age 23 and follow individuals until 65, whereas we start ours at age 25 and follow them until 55.

Table 14: Sources of Variation in Lifetime Earnings and Hours

	Lifetime log Earnings	Lifetime log Hours
$\sigma_\psi = \sigma_\pi = 0$	82.8%	2.4%
$\sigma_\psi = 0$	86.3%	31.7%
$\sigma_\pi = 0$	96.1%	70.3%
$\sigma_z = 0$	59.6%	99.2%
$\sigma_x = \sigma_\alpha = 0$	51.5%	96.8%
$\sigma_x = 0$	57.9%	100.8%
$\sigma_\alpha = 0$	76.5%	96.4%
$\sigma_x = \sigma_\alpha = \sigma_\psi = 0$	45.0%	30.1%
$\sigma_z = \sigma_\pi = 0$	54.9%	69.5%

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

the propagation of human capital shocks to lifetime earnings.

Huggett, Ventura and Yaron (2011) offered a binary decomposition of the sources of lifetime earnings inequality between initial conditions (initial human capital and learning ability) and shocks to human capital. When we eliminate heterogeneity in initial human capital and learning ability in our model, the variance of log lifetime earnings decreases by 48.5 percent (see row 5, column 1), substantially smaller than the 61.5 percent value found by them. In terms of matching patterns in the data, preference heterogeneity partially substitutes for both initial human capital differences as well as differences in learning ability.

The previous calculation considered the joint effect of initial human capital and learning ability. The fact that we assume perfect correlation between initial human capital and learning ability perhaps makes it natural to consider them jointly, but one can also assess the relative importance of initial human capital and learning ability. When we do this we find that the effect of initial human capital on lifetime earnings inequality is almost twice as large as the effect of learning ability.

We can also consider a broader version of the binary decomposition between initial conditions and shocks. In our model, permanent preference heterogeneity represents heterogeneity in initial conditions, while transitory preference heterogeneity represents heterogeneity in shocks. The bottom row of Table 14 reports these results. Our model implies that roughly 55 percent of the variance of log lifetime earnings is due to initial conditions.

It is also of interest to examine the sources of inequality in lifetime hours. Eliminating human capital shocks has an almost negligible effect on the variance of log lifetime hours, reducing it by less than 1.0 percent, while jointly eliminating heterogeneity in initial human capital and learning ability has a larger but still modest effect, reducing the variance of log lifetime hours by 3.2 percent. Combined with our earlier results, it follows that preference heterogeneity, specifically permanent preference heterogeneity, is the overwhelming source of heterogeneity in lifetime hours.

7 A Policy Application

Two messages emerge from the results in the previous section. First, variation in hours worked has important effects on the variance of earnings, both through a direct channel that affects production time and a dynamic channel that affects human capital accumulation. Second, variation in hours of work are driven by preference heterogeneity. The first message implies that policies that affect the distribution of hours of work may have important effects on both the mean and dispersion of lifetime earnings. The second message implies that such policies may be a source of misallocation if they change relative work hours across individuals with differing levels of disutility of work.

There are many policies that directly affect the distribution of hours of work. As one example, several European countries have enacted legislation to reduce the standard workweek. Perhaps the most publicized of these was the decision by France to reduce the standard workweek to 35 hours. Another class of policies aims to compress the distribution of hours worked from above. For example, all advanced economies have legislation that stipulates a threshold for overtime hours and the level of the overtime premium. France offers a more direct example of this class of policies. Specifically, in a pair of laws adopted between 1998-2000 and rolled out between 2000-2002 ("Aubry I and Aubry II"), France adopted a regulation imposing that most workers could work no more than 48 hours per week.[27]

Many researchers have studied the impact of reductions in the standard workweek.[28] We believe our model has important implications regarding the effects of policies that seek to compress the hours distribution from above, and so in this section we use our calibrated model to assess the direct effects of such a policy. We can carry out this exercise for any specific maximum constraint on hours, but motivated by the French legislation we choose a constraint of 48 hours per week to illustrate the effects. Specifically, we take our benchmark-calibrated model and re-solve for the equilibrium assuming that all workers face a constraint that total hours of work cannot exceed 2496

[27] In the opposite direction, France also recently enacted a regulation imposing a minimum workweek for part-time workers. Because we focus on highly attached males our model is not well suited to study this issue. Carry (2024) develops a model to study the effect of this regulation. Relatedly, in 2024 Greece enacted legislation to make it easier for firms to have a six day workweek that would imply a 48 hour workweek.

[28] Examples include Hunt (1999), Marimon and Zilibotti (2000), Crépon and Kramarz (2002), Rocheteau (2002), Chemin and Wasmer (2009), Raposo and van Ours (2010), Goux, Maurin and Petrongolo (2014), Lopes and Tondini (2022), and Batut, Garnero and Tondini (2023).

($= 48 \times 52$) in any period.[29]

In our partial equilibrium model with exogenous prices, an additional constraint on hours choices cannot have positive effects for any individuals. Policy makers presumably imagined some positive effects associated with this policy, perhaps via general equilibrium effects on prices. For example, they might have thought that limiting supply for some workers would increase demand for the labor supplied by other workers. Our partial equilibrium model does not have anything to say about the potential source of these positive effects and their magnitude, but it can offer guidance about the potential magnitude and distribution of the direct negative effects associated with the additional constraint on hours. Here we focus on these direct effects and leave a full general equilibrium analysis to future work.

Before turning to the results, it is useful to first document some properties of the distribution of hours in our benchmark calibrated model, as this will provide information about how many workers are affected and what their characteristics are. First, almost half—49.7 percent—of annual observations at the individual level violate the constraint. Second, these violations are heavily concentrated: 50 percent of the total violations are accounted for by 24.9 percent of individuals, though only 2.0 percent of individuals are never affected by the constraint. It follows that although almost everyone will experience some negative effect of this regulation, the effects will be highly concentrated.

Third, the violations are concentrated among individuals in the left tail of the ψ distribution: Individuals with ψ values in the lowest quintile account for 35.1 percent of the violations. Intuitively, violations are more likely for individuals that work long hours.

Fourth, the violations occur fairly evenly throughout the life-cycle: 25.9 percent of violations occur between ages 25 and 35, and 22.6 percent occur between ages 55 and 64. This is consistent with the fact that violations are concentrated among those with low ψ values, and which is a permanent characteristic.

Fifth, a substantial fraction of violations occur in the lower part of the lifetime earnings distribution, which, consistent with our earlier calculations, are defined over the age range 25-55. Separating individuals into quintiles of the lifetime earnings distribution and as we move from the lowest to the highest lifetime earnings quintile, the share of total violations are 13 percent, 17 percent, 20 percent, 23 percent, and 26 percent. The fact that higher hours will lead to higher earnings, while holding all else constant, does create some concentration of violations in the highest lifetime earnings quintile. But since violations are concentrated among individuals with low values of ψ,

[29] Because our benchmark economy is calibrated to the US, this exercise should be understood as assessing how outcomes in the US economy would be affected by such a constraint. In particular, our results should not be interpreted as the marginal effect of this constraint on outcomes in the French economy. This would require a model calibrated to French data and incorporating other relevant regulations.

and ψ is only modestly (negatively) correlated with initial human capital, there are many violations among those with relatively low lifetime earnings. Specifically, more than 70 percent of violations are accounted for by individuals with lifetime earnings outside of the top quintile.

Turning to the results of the policy counterfactual, our main finding is that the restriction of holding prices constant has large effects on both the level and dispersion in lifetime earnings. We begin by examining the effects on mean lifetime earnings. Introducing a cap on hours in our model causes mean lifetime earnings to fall by 12.5 percent. This drop in lifetime earnings reflects both a decrease in hours and a decrease in human capital accumulation. Specifically, mean lifetime hours drop by 8.8 percent, while mean human capital in the pooled cross-section sample of individuals aged 25-55 falls by 3.6 percent. While the constraint that we impose applies to total time, the percentage drop in average production time (8.0 percent) is actually smaller than the percentage drop in average investment time (12.3 percent).

While the effects on total hours are concentrated on individuals with low values of ψ, the human capital effects are even more concentrated among these individuals. Human capital investment effects will be largest for young individuals who previously expected to work long hours when old, which are precisely low ψ individuals. Notably, because there is only a modest negative correlation between ψ and initial human capital in our model, the individuals whose human capital accumulation is most affected are spread throughout the lifetime earnings distribution. These high human capital accumulators are the individuals who experience the most upward movement within the earnings distribution over the life-cycle.

Next we turn to the effects of the policy on earnings inequality. As a result of the hours constraint, the variance of lifetime earnings falls by 13.0 percent. The fall in the variance of lifetime hours is even larger, at 66.6 percent. The variance of average human capital falls by 5.0 percent. The decrease in the variance of log lifetime earnings (0.0419) is roughly two and a half times larger than the decrease in the variance of log lifetime hours (0.0166). These results are reminiscent of the results we found in Section 6.2 where we examined the effect of eliminating all preference heterogeneity.

To provide additional perspective on the distribution of changes in lifetime earnings, Figure 9a provides a plot of changes versus two dimensions of heterogeneity: initial human capital (which is perfectly correlated with learning ability) and permanent tastes for work. Consistent with the previous discussion, four features are prominent. First, some groups experience very large losses in lifetime earnings, with losses as large as twenty-five percent. Second, there is a large amount of dispersion in the magnitude of these losses, as some groups in this figure experience losses of less than 5 percent. Third, lower values of ψ are associated with much larger losses. And fourth, higher learning ability is associated with modestly higher losses. This last effect is due to the effect of higher learning ability. In our calibrated model, people with higher learning ability will engage in higher investment in human capital when young and devote more time to production when old.

Figure 9: Changes in Lifetime Earnings and CEV with Hours Cap

(a) Lifetime Earnings (b) Consumption Equivalent Variation

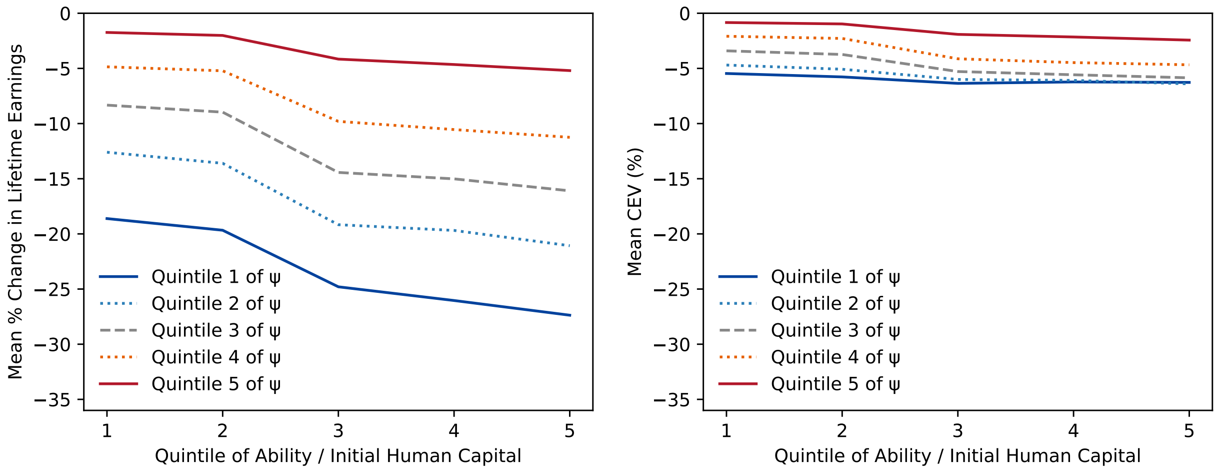

The regulation will dampen these effects for affected individuals.

We can also assess the distribution of welfare losses associated with this regulation. To do this, we solve for the lifetime utility loss for each individual and compute the proportional change in the consumption life-cycle profile in the benchmark model that would produce an equivalent lifetime utility loss. Figure 9b shows these losses as a function of initial human capital and the permanent component of preference heterogeneity. The patterns in this figure mirror those in Figure 9a. In particular, lower values of ψ are associated with larger welfare losses, and holding ψ constant, higher initial human capital is associated with modestly higher welfare losses. The key takeaway from this figure is that the welfare effects are substantial for the most affected groups. All individuals with ψ values in the lowest quintile have consumption equivalent welfare losses that exceed 5 percent. And several other groups also experience welfare losses close to 5 percent.

The overall losses associated with a policy that constrains hours from above are heavily dependent on the distribution of hours in the initial equilibrium. Put somewhat differently, assessing the costs associated with such a policy requires that one have a model that adequately captures the distribution of hours found in reality. We argued earlier that generating sufficient variation in cross-sectional hours and especially lifetime hours requires that one introduce preference heterogeneity and, specifically, a permanent component to preference heterogeneity. In fact, if we examine the consequences of the same hours restriction in our model when we eliminate preference heterogeneity, we find that the regulation on hours has effectively no effect, as the mass of affected individuals drops from 49.7 percent to less than 0.1 percent.

8 Conclusion

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

We close by noting two important areas for future research. The first is to extend our analysis to groups beyond the highly attached male sample that we focus on in this paper. Differences in lifetime hours of work are even larger if we consider broader groups, raising the possibility that one will find even larger effects than we found here. This will require a richer model that explicitly accounts for spells of non-participation.

A second area is to consider a wider range of specific policies that affect the distribution of working time and to incorporate general equilibrium effects. We used our model to carry out a partial equilibrium analysis of a stylized policy that constrains maximum weekly hours of work. We found that such a policy has large effects on both mean lifetime earnings and inequality. Individuals with the largest losses are distributed throughout the lifetime earnings distribution. The recent paper by Carry (2024) carries out a general equilibrium analysis of a policy that places a minimum constraint on workweeks and finds that general equilibrium effects are large.

References

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

Batut, Cyprien, Andrea Garnero, and Alessandro Tondini. 2023. "The employment effects of working time reductions: Sector-level evidence from European reforms." *Industrial Relations: A Journal of Economy and Society*, 62(3): 217–232.

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

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

Bick, Alexander, Adam Blandin, and Richard Rogerson. 2024. “After 40 Years, How Representative Are Labor Market Outcomes in the NLSY79?” Federal Reserve Bank of St. Louis Working Paper 2024-010B.

Blundell, Richard, Monica Costa Dias, Costas Meghir, and Jonathan Shaw. 2016. “FEMALE LABOR SUPPLY, HUMAN CAPITAL, AND WELFARE REFORM.” *Econometrica*, 84(5): 1705–1753.

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

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

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

Carry, Pauline. 2024. “The Effects of the Legal Minimum Working Time on Workers, Firms and the Labor Market.” HAL SciencePo Working papers Main.

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

Chemin, Matthieu, and Etienne Wasmer. 2009. “Using Alsace-Moselle Local Laws to Build a Difference-in-Differences Estimation Strategy of the Employment Effects of the 35-Hour Workweek Regulation in France.” *Journal of Labor Economics*, 27(4): 487–524.

Couch, Kenneth A., and Dana W. Placzek. 2010. “Earnings Losses of Displaced Workers Revisited.” *The American Economic Review*, 100(1): 572–589.

Crépon, Bruno, and Francis Kramarz. 2002. “Employed 40 Hours or Not Employed 39: Lessons from the 1982 Mandatory Reduction of the Workweek.” *Journal of Political Economy*, 110(6): 1355–1389.

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

Glover, Andrew, José Mustre-del Río, and Emily Pollard. 2023. “Lifetime Earnings Differences across Black and White Individuals: Years Worked Matter.” *Economic Review (01612387)*, 108(1).

Goux, Dominique, Eric Maurin, and Barbara Petrongolo. 2014. "Worktime Regulations and Spousal Labor Supply." *American Economic Review*, 104(1): 252–76.

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

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

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

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.

Hunt, Jennifer. 1999. "Has Work-Sharing Worked in Germany?" *The Quarterly Journal of Economics*, 114(1): 117–148.

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

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

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

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

Lopes, Marta C., and Alessandro Tondini. 2022. "Firm-Level Effects of Reductions in Working Hours." Research Institue for the Evaluation of Public Policies (IRVAPP) FBK-IRVAPP Working Papers 2022-05, Bruno Kessler Foundation.

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

Marimon, Ramon, and Fabrizio Zilibotti. 2000. "Employment and distributional effects of restricting working time." *European Economic Review*, 44(7): 1291–1326.

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

Raposo, Pedro S., and Jan C. van Ours. 2010. "How working time reduction affects jobs and wages." *Economics Letters*, 106(1): 61–63.

Rocheteau, Guillaume. 2002. "Working time regulation in a search economy with worker moral hazard." *Journal of Public Economics*, 84(3): 387–425.

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

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

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

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

White, Roger. 2010. "Long-run wage and earnings losses of displaced workers." *Applied Economics*, 42(14).

ONLINE APPENDIX

A Proof for Optimal Human Capital Investment Condition (1)

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

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

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

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

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

This problem can be written recursively as:

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

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

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

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

$$V_t(x,k) = \max_{h,Q,k'} \quad u((1+r)k + whx - C(Q) - k', h) + \beta V_{t+1}(x', k')$$

$$s.t. \quad x' = (1-\delta)x + Q$$

$$h, Q \geq 0, \text{ with equality if } t \geq T_R$$

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

$$FOC(Q_t): \quad C'(Q)u_1(c,h) = \beta V_{t+1,1}(x',k') \tag{A.8}$$
$$Env(x_t): \quad V_{t,1}(x,k) = u_1(c,h)wh + \beta V_{t+1,1}(x',k')(1-\delta) \tag{A.9}$$

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

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

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

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

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

and substituting this into the previous equation yields

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

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

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

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

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

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

B Data

B.1 Hours Worked by Gender and Minimum Years with at Least 520 Annual Hours

Table B.1: Men

	1	10	20	31
Mean	2117.9	2150.5	2241.0	2421.9
Coefficient of Variation				
Annual	0.43	0.41	0.35	0.25
Lifetime	0.28	0.25	0.20	0.15
Individuals	3001	2918	2639	1418

Table B.2: Women

	1	10	20	31
Mean	1579.2	1655.4	1827.4	2139.6
Coefficient of Variation				
Annual	0.60	0.55	0.44	0.25
Lifetime	0.39	0.33	0.23	0.14
Individuals	3236	3028	2472	879

B.2 Comparison of Autocorrelation Profiles

One may worry that the survey design and our imputation procedure creates too much persistence in annual hours worked and annual earnings. Figure B.1a shows that the autocorrelation profile excluding imputed values for annual hours is virtually identical to the profile shown in Figure 1c. Figures B.1b and B.1c show the autocorrelation profile excluding imputed values for annual earnings alongside the profile including imputed earnings. When doing this we distinguish between two time periods. Before 1994, earnings were collected annually, and during this period the two profiles are virtually identical. From 1994 onward, earnings are only reported every other year. The two-year auto-correlation in the directly reported data is slightly lower than when including imputed values. From lag 4 onward, the two profiles again lie almost on top of each other.

FIGURE B.1: Autocorrelation of Annual Hours and Earnings, Ages 25-55

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

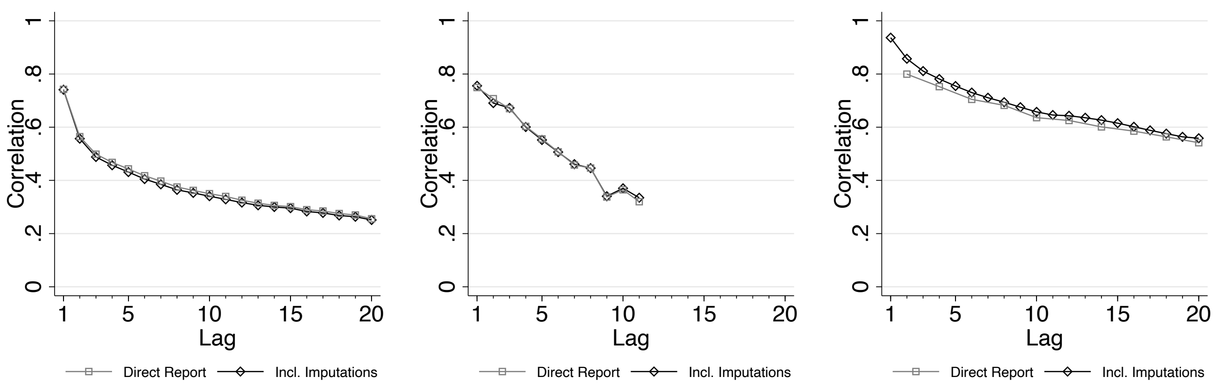

B.3 Requiring Only 20 Years of Employment

The first column of Table B.3 summarizes the distribution over annualized lifetime hours when we require only at least 20 years of employment with at least 520 hours worked. On average, a man in this less-restrictive sample works 2241 hours per year from age 25-55 (compared with 2422 hours when we require at least 520 annual hours in all sample years). The standard deviation in the less-restrictive sample is 451 hours, compared with 374 hours in the more-restrictive sample, which is 20.1% of the mean (compared with 15.5% in the more-restrictive sample). The interquartile range is 485 hours, and the interquartile ratio is 1.24 (compared with 433 hours and 1.20, respectively).

Table B.3: The Distribution of Lifetime Hours and Components

Percentile	Annualized Lifetime Hours	Years Worked	Weeks per Year Worked	Hours per Week Worked
5	1442.6	26.0	44.4	39.3
10	1681.9	27.6	45.9	41.4
25	1992.6	30.0	48.5	42.7
50	2211.0	30.7	50.2	44.5
75	2477.7	31.0	50.9	48.7
90	2782.1	30.9	51.2	54.5
95	3001.7	30.8	51.1	59.1

Notes: We sort individuals into percentiles by their lifetime hours and divide them by 31 years, the number of years an individual is in our sample. *Years worked*, *Weeks per year worked*, and *Hours per week worked* are the average values for all individuals in a given percentile and the two adjacent ones. Note that the product of these three variables is slightly different than the average annualized lifetime hours for each percentile.

In the less-restrictive sample, years worked matter for lifetime hours worked. Note that we define here a year worked as any year with positive hours such that on the individual level the product of columns 2 to 4 is the same as column 1. Columns 2 and 3 of Table B.3 show substantial variation in the number of years worked and weeks worked, particularly at the lower end of lifetime hours distribution. Naturally, in the top half of the distribution most variation comes from weekly hours worked. Weekly hours are however also the most important driver across the entire hours distribution. The 90/10 earnings ratio of years worked and weeks worked is about 1.12, while for weekly hours worked it is 1.31.

Figures B.2 and B.3 show life-cycle and autocorrelation profiles for hours and earnings for the less-restrictive sample. Figure B.4 and Table B.4 show the relationship between lifetime hours and earnings.

FIGURE B.2: Cross-section of Hours Worked Over the Life-Cycle, Ages $25-55$

(a) Mean log Annual Hours (b) S.D. of log Annual Hours (c) Autocorrelation of Hours

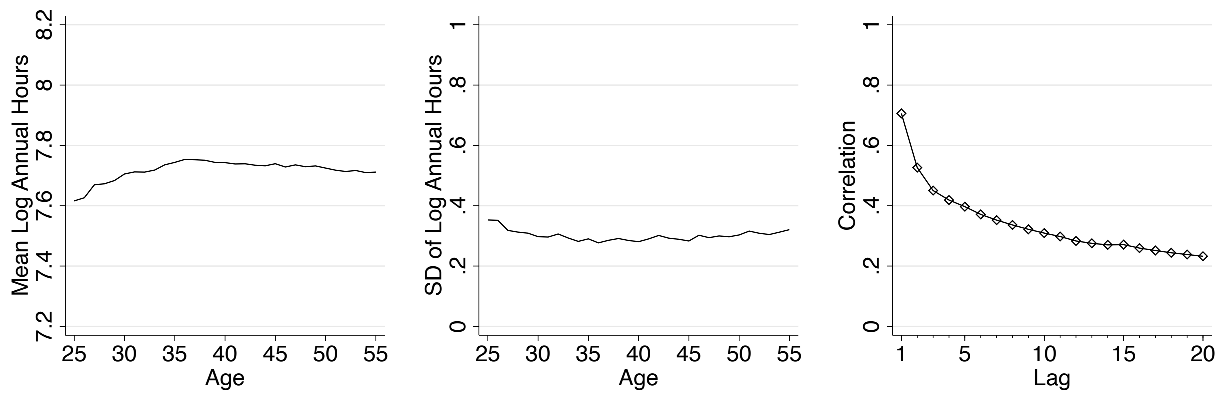

Notes: Moments are conditional on working at least 520 hours per year.

FIGURE B.3: Cross-section of Annual Earnings Over the Life-Cycle, Ages $25-55$

(a) Mean log Annual Earnings (b) S.D. of log Annual Earnings (c) Autocorrelation of Earnings

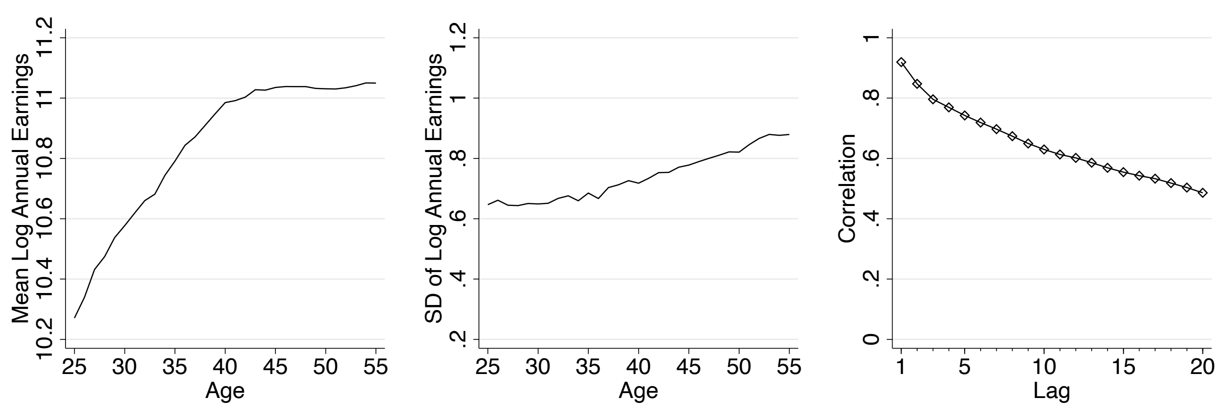

Notes: Moments are conditional on annual hours being at least 520 hours.

FIGURE B.4: The Correlation Between Lifetime Hours and Earnings, Ages $25-55$

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

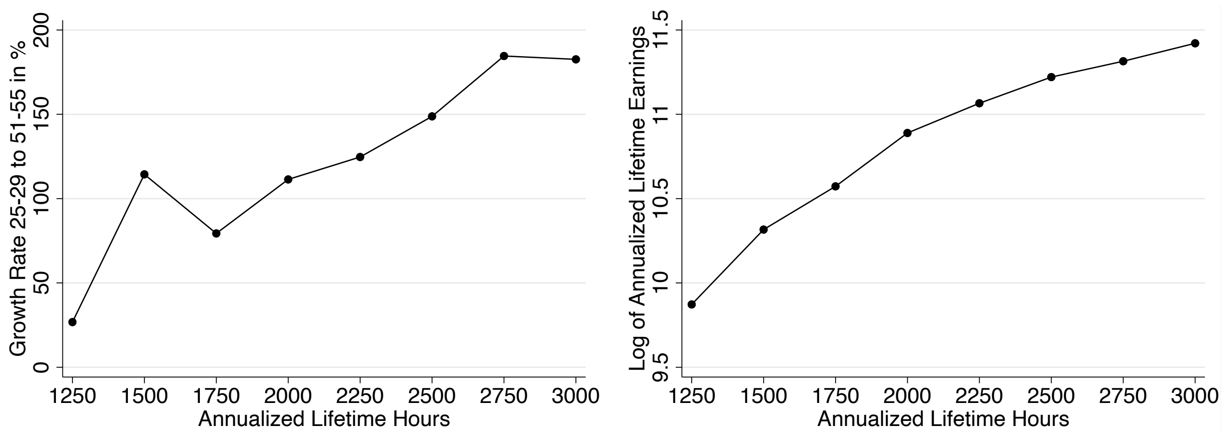

Notes: In (b), we restrict the sample to individuals who work at least 520 hours at each age 25-29 and 51-55.

Table B.4: Lifetime Hours and Earnings Regressions

(a) Life-Cycle Earnings Growth

	(1)	(2)
log Lifetime Hours	1.70***	1.55***
	(0.26)	(0.26)
AFQT Percentile		0.01***
		(0.00)
Constant	1.27***	1.29***
	(0.04)	(0.04)
N	1767	1695
R^2	0.02	0.07

(b) log of Lifetime Earnings

	(1)	(2)
log Lifetime Hours	1.81***	1.61***
	(0.05)	(0.05)
AFQT Percentile		0.01***
		(0.00)
Constant	10.89***	10.89***
	(0.01)	(0.01)
N	2639	2534
R^2	0.33	0.49

Notes: In (a), we restrict the sample to individuals who work at least 520 hours at each age 25-29 and 51-55. In both regressions, log annualized lifetime hours and AFQT percentile are demeaned such that the constants are comparable across specifications (1) and (2). Standard errors are in parentheses. *,**, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

B.4 Average Wage Growth in the CPS

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

C Quantitative Results

Table C.5: Sources of Variation in Earnings and Hours

	Annual		Lifetime	
	log Earnings	log Hours	log Earnings	log Hours
$\sigma_\psi = \sigma_\pi = 0$	81.0%	15.9%	82.8%	2.4%
$\sigma_\psi = 0$	90.9%	74.0%	86.3%	31.7%
$\sigma_\pi = 0$	89.8%	41.8%	96.1%	70.3%
$\sigma_z = 0$	60.5%	99.5%	59.6%	99.2%
$\sigma_x = \sigma_\alpha = 0$	65.3%	99.2%	51.5%	96.8%
$\sigma_x = 0$	79.9%	100.3%	57.9%	100.8%
$\sigma_\alpha = 0$	88.0%	98.8%	76.5%	96.4%
$\sigma_x = \sigma_\alpha = \sigma_\psi = 0$	60.8%	72.0%	45.0%	30.1%
$\sigma_z = \sigma_\pi = 0$	49.9%	41.3%	54.9%	69.5%

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