Is Earnings Inequality Driven by Talent or Tastes?*

Ian Fillmore[†] Trevor S. Gallen[‡]

March 8, 2023

Abstract

Do tastes for leisure play an important role in determining earnings inequality? Using NLSY79 data, we fit a simple lifecycle model in which workers differ in (i) their human capital at age 30, (ii) their talent for accumulating human capital, and (iii) their tastes for leisure. We find that tastes play an important role in earnings inequality and a model without taste differences cannot match the data. Eliminating talent differences would cut the variance of pre-tax earnings in half, while eliminating taste differences would cut the variance by three-fourths. We conclude by exploring implications of our findings for optimal taxation.

Keywords: Income inequality, human capital, tastes for leisure, optimal taxation

JEL Codes: J22, J24, H21, H24

*The authors would like to thank Marcus Berliant, Eliza Forsythe, Yana Gallen, Jonathan Heathcote, Lutz Hendricks, Mark Huggett, Greg Kaplan, Ben Lockwood, Louis Kaplow, Magne Mostad, Casey Mulligan, Steven Levitt, and seminar participants at the Midwest Macro meetings, Midwest Economics Association meetings, Purdue University, Olin School of Business, DePaul University, the National Tax Association meetings, Indiana University - Purdue University Indianapolis, and the 2019 Society for Economic Dynamics meetings for helpful feedback and comments.

[†]Department of Economics, Washington University in St. Louis, ianfillmore@wustl.edu.

[‡]Department of Economics, Krannert School of Management, Purdue University, tgallen@purdue.edu.

1 Introduction

Do differences in tastes for leisure play an important role in determining income inequality? Casual empiricism suggests that at least some of the differences in people's life outcomes are driven by differences in their life goals and tastes. However, beginning with Mirrlees (1971), the literature on optimal tax design has traditionally assumed that workers possess the same tastes and only differ in their productivity or wages, which are assumed to be exogenous rather than emerging from endogenous human capital investments earlier in life. These simplifying assumptions have proven useful in understanding the short run incentive effects of income taxes, but they cannot provide an adequate explanation for the sources of income inequality. And the sources of income inequality matter greatly when considering the optimal amount of income redistribution. More broadly, whether and how income inequality is a problem, and whether and how public policy should address it, depend crucially on the sources of income inequality.

In this paper, we begin by discussing a simple extension of the classic Ben-Porath model of lifecycle human capital investment (Ben-Porath, 1967) that allows people to differ in their tastes for leisure. The model combines endogenous human capital investment over the lifecycle with a labor-leisure decision. In the model, agents differ in their initial human capital at age 30, their talent for acquiring human capital, and their tastes for leisure.[1] Neal and Rosen (2005) emphasize that workers with low tastes for leisure not only sacrifice little by giving up leisure time to invest in human capital, but the benefits of investing in human capital are higher because they expect to work more later in life. Consequently, individuals with low tastes for leisure will tend to invest more in human capital when young, leading to higher wages when older. And just like talent, small differences in tastes for leisure can compound into large differences in human capital and earnings later in life. Unlike more sophisticated models in the literature, we omit many dimensions of labor supply and human capital investment. For instance, we ignore household decisions, health shocks (Hokayem and Ziliak, 2014), fertility (Rosenzweig and Wolpin, 1980), the presence of children (Blundell et al., 2005; Cherchye et al., 2012), and involuntary unemployment and search frictions (Low et al., 2010). Rather than attempt to explicitly model these factors, we take a similar approach to Kaplan (2012) and mitigate their impact through our choice of

[1] It is important to note that taste differences do not serve as a residual in our model as in, for example, MaCurdy (1981). Rather, we estimate taste differences from the joint distribution of work hours and earnings over the lifecycle.

sample, age span, and data treatment. We focus on ages 30–44 to avoid years of schooling when workers are young and retirement as workers get older. We abstract from transitory health shocks, involuntary unemployment, and search frictions by estimating an auxiliary statistical model that filters out "transitory" shocks and measurement error in earnings and work hours. Lastly, we minimize the importance of fertility shocks by focusing on highly-attached, prime-age males, whose labor supply is less affected by fertility (Angrist and Evans, 1998). Although much more complex models have been developed and estimated in the literature, we opt for a simple and transparent model of lifecycle human capital investment and labor supply, to highlight the key moments in the data behind our results. Although prior work has considered the interaction between optimal taxation and human capital investment (Stantcheva, 2015; Kapička, 2015) and the implications of taste differences for optimal taxation (Fleurbaey and Maniquet, 2006; Lockwood and Weinzierl, 2015; Fleurbaey and Maniquet, 2018; Bergstrom and Dodds, 2021), our model highlights an important interaction between lifecycle dynamics and differences in tastes for leisure. As pointed out by Neal and Rosen (2005), incorporating taste differences into a model of human capital investment introduces a *dynamic* channel by which relatively small differences in preferences can compound into large wage and earnings differences later in life.

Motivated by our model, we examine three pieces of reduced-form evidence that, although difficult to explain without taste differences, arise naturally in a model with different tastes for leisure. Our data come from the National Longitudinal Survey of Youth 1979 (NLSY79) First, we find a high variance in "permanent" work hours—a standard deviation between seven and eight hours per week that does not decline with age—even among prime-age males who are highly attached to the labor force and even after filtering out transitory shocks and measurement error. Second, college graduates work about three more hours per week than non-graduates, even after conditioning on the current wage.[2] Third, past leisure predicts future leisure, even after conditioning on wages and lifetime income.[3] We find that a worker who consumes 10% more leisure between ages 30–34 tends to consume 2.6% more leisure between ages 50–54. Although all three pieces of evidence are inconsistent with a labor supply model with identical tastes, the evidence is perfectly consistent with a model of human capital investment wherein individuals

[2]Without taste differences, after conditioning on the current wage college graduates should consume *more* leisure than non-graduates, due to a pure income effect. With taste differences, it may be that college graduates have a lower taste for leisure which causes them to invest in more human capital *and* work more hours.

[3]In a standard lifecycle model of labor supply, the first order condition for leisure implies that past leisure hours should not predict future leisure hours after accounting for lifetime income and the current wage.

with a lower taste for leisure optimally choose to invest more in human capital.

With these three pieces of reduced-form evidence, we then turn to estimating our structural model using indirect inference. We fit an auxiliary statistical model to the data that estimates the joint distribution of the levels and slopes of earnings and work hours over the lifecycle. The auxiliary model serves two purposes: first, it filters out transitory shocks and measurement error, and second it provides a parsimonious statistical representation of the joint distribution of earnings and work hours over the lifecycle. We then choose the joint distribution of talent, tastes for leisure, and initial human capital to match the joint distribution of earnings and work hours over the lifecycle. We find that our model fits the data well, while a model without differences in tastes for leisure cannot match the joint distribution of earnings and work hours.

We find a large role for taste differences in driving income inequality. We calculate the elasticities of the mean, variance, and skewness of log earnings with respect to the variances of talent, taste, and initial human capital. We find that, while lowering the variance of talent by one percent would reduce earnings inequality by 1.24 percent, lowering the variance of tastes by one percent would reduce earnings inequality by a bit more, 1.38 percent. At the margin, tastes are at least as important as talent in driving income inequality. We also simulate counterfactuals wherein we completely eliminate differences in talent and in tastes for leisure. Completely eliminating talent differences cuts in half the variance of log earnings at age 44, while completely eliminating taste differences has an even larger effect, cutting the variance of log earnings by three-quarters.

Our model is able to distinguish talent from tastes using the joint distribution of earnings and hours over the lifecycle. We show that observing earnings alone is insufficient. Although increasing talent or decreasing tastes for leisure both increase earnings growth by increasing human capital investment, they have very different implications for work hours. Lowering tastes for leisure will cause an agent to take less leisure, thereby raising both earnings and work hours and generating a positive covariance between the two. On the other hand, producing the same earnings profile by instead raising talent and/or initial human capital, will cause work hours to fall due to an income effect, thereby generating a negative covariance. This suggests that the strong positive covariance between earnings and work hours over the lifecycle that we observe in the data is key to our finding that tastes for leisure are an important driver of income inequality. To verify this intuition, we re-estimate the model but constrain it to have no taste differences, and we confirm that this no-taste-differences model misses badly on the covariance of earnings and work hours—rather than match the strong positive covariance between the two, the model generates a negative covariance. We then manually decrease the covariance between log earnings and work hours at age 30 from its estimated value of 2.04 to zero which causes the estimated variance in tastes to fall by half and the no-taste-differences model to fit much better. If, in addition, we manually decrease the correlation between earnings at age 30 and subsequent earnings growth from -0.2 to -0.5, the variance of tastes would drop even further and the fit of the no-taste-differences model would improve even more.[4] Thus, these two moments are key to understanding the role of tastes for leisure in driving income inequality.

Having estimated our model and found significant differences in taste, we examine the consequences for the optimal tax if the same earnings inequality were due to variation in talent or variation in tastes for leisure. We find that the drivers of income inequality alter the optimal tax schedule, although the direction and magnitude of the effect depend on the relative curvature of consumption and leisure in the utility function and, disturbingly, the scaling of utility. For someone earning \$100,000, we find that the marginal tax is much higher if utility is scaled so that taste differences reflect differences in the marginal utility of consumption rather than leisure. What's more, decreasing taste differences would lower the marginal tax rate in the former case while raising it in the latter. While interpersonal utility comparisons are always fraught, this sensitivity of optimal taxes to the (unknowable) scaling of utility functions raises both philosophical and practical concerns.[5] One might argue on philosophical grounds that redistribution on the basis of tastes alone is unjustified (Lockwood and Weinzierl, 2015; Fleurbaey and Maniquet, 2006). But in a dynamic setting where wages, work hours, and earnings all depend on both talent *and* tastes for leisure, it is not obvious how an income tax could or should unravel the two.[6]

Our paper makes three contributions to the literature on income inequality and optimal taxation. First, motivated by the standard labor-leisure first order condition, we report reduced-form evidence that tastes for leisure differ. Second, we fit a simple lifecycle model of human capital

[4]The correlation between earnings at age 30 and subsequent earnings growth governs the degree to which lifecycle earnings paths cross after age 30.

[5]Aguiar and Hurst (2007) point out that the rise in leisure hours among less-educated workers also complicates a welfare analysis of rising income inequality.

[6]Fleurbaey and Maniquet (2006) do consider the case of an optimal income tax when agents differ in both skills and tastes for leisure. But their setting is static, and their results do not apply when future wages depend on both talent and tastes for leisure.

investment and labor supply to NLSY79 data on highly-attached, prime-age male workers. We find tastes for leisure differ substantially with, if anything, more of an effect on income inequality than talent. Two moments are key for this result: the large positive covariance of initial hours and initial income, and the relatively small covariance between earnings at age 30 and subsequent earnings growth. Third, we find that taste differences affect optimal taxes, but the magnitude and direction of the effect is disturbingly sensitive to the scaling of utility.

2 Lifecycle Model of Human Capital Investment

In this section, we present a simple model of human capital investment. The model closely follows the extension of Ben-Porath (1967) by Neal and Rosen (2005). The Ben-Porath model of human capital investment is a canonical model for understanding earnings over the lifecycle. In the model, agents maximize the net present value of earnings over their lifetime, trading off between paid work time and human capital investment, which does not earn a wage but does increase future wages. Agents are characterized by the triple $(A, \phi, \overline{k})$. A captures talent, or the agent's ability to acquire new human capital, ϕ captures his relative taste for leisure, and $\overline{k}$ is his initial level of human capital. In each period, the agent allocates his time between labor n_t, leisure ℓ_t, and human capital investment s_t. Human capital grows according to the law of motion

$$k_{t+1} = (1-\delta)k_t + A(s_t k_t)^\gamma \tag{1}$$

where $\gamma \in (0, 1)$. Talent (A) reflects the agent's efficiency at acquiring new human capital. The agent's wage in period t is $w_t = Rk_t$ and depends solely on his human capital. Note that talent does not directly raise wages. Rather, talented individuals find it easier to accumulate human capital over time.

The standard Ben-Porath model does not include a leisure choice and abstracts from consumption by assuming that the agent has access to complete markets. The result is that the agent trades off labor and human capital investment over the lifecycle so as to maximize the net present value of lifetime income. As Neal and Rosen (2005) show, one can incorporate a leisure decision into the Ben-Porath model. Period utility depends on consumption c_t, leisure ℓ_t, and the taste parameter ϕ which determines the agent's relative taste for consumption vs leisure. The agent maximizes lifetime utility subject to a period time constraint and a lifetime budget constraint. The agent solves

$$\begin{aligned}
\max_{\{c_t s_t \ell_t n_t\}} \quad & \sum_{t=1}^{T} \beta^{t-1} U(c_t, \ell_t; \phi) \\
s.t. \quad & s_t + \ell_t + n_t = 1 && (2) \\
& s_t,\ \ell_t,\ n_t \geq 0 \\
& \sum_{t=1}^{T} \left(\frac{1}{1+r}\right)^{t-1} c_t = \sum_{t=1}^{T} \left(\frac{1}{1+r}\right)^{t-1} w_t n_t && (3) \\
& w_t = R k_t \\
& k_{t+1} = (1-\delta) k_t + A (s_t k_t)^{\gamma} \\
& k_1 = \overline{k}
\end{aligned}$$

where our lifecycle budget constraint in equation (3) implicitly assumes perfect capital markets, in which agents are able to borrow and save without limit, so long as they begin and end life with zero assets. The Kuhn-Tucker conditions for this problem are

$$\begin{aligned}
[c_t]: \quad & \beta^{t-1} U_c(c_t, \ell_t; \phi) = \left(\frac{1}{1+r}\right)^{t-1} \lambda \\
[\ell_t]: \quad & U_\ell(c_t, \ell_t; \phi) = \mu_t \\
[n_t]: \quad & \left(\frac{1}{1+r}\right)^{t-1} \lambda w_t = \beta^{t-1}(\mu_t + \mu_t^n) \\
[n_t]: \quad & \mu_t^n n_t = 0;\ \mu_t^n \geq 0;\ n_t \geq 0 \\
[s_t]: \quad & A\gamma (s_t k_t)^{\gamma-1} k_t \lambda \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_\tau = \mu_t
\end{aligned}$$

where λ is the Lagrange multiplier on the lifetime budget constraint (3), μ_t is the Lagrange multiplier on the period time constraint (2), and μ_t^n is the multiplier on the non-negativity constraint for labor.[7] Note that human capital investment s_t will always be positive (except in the last period) because the marginal product of human capital investment is infinite at zero. In the case where ℓ_t, $n_t > 0$ so that the agent is at an interior solution, $\mu_t^n = 0$ and the first

[7]Because the agent is maximizing utility, rather than the net present value of earnings as in the classic Ben-Porath model, there is no closed-form solution.

order conditions for ℓ_t and s_t become

$$[\ell_t]: \quad \beta^{t-1} U_\ell(c_t, \ell_t; \phi) = \left(\frac{1}{1+r}\right)^{t-1} \lambda w_t \tag{4}$$

$$[s_t]: \quad \left(\frac{1}{\beta(1+r)}\right)^{t-1} \frac{R(s_t k_t)^{1-\gamma}}{A\gamma} = \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_\tau. \tag{5}$$

To interpret these conditions, consider the case where $\beta = \frac{1}{1+r}$ and the utility function is separable in consumption and leisure. In this case, consumption will be constant over the lifecycle, and the agent will set the marginal utility of leisure proportional to the wage in each period. The wage will rise over the lifecycle as the agent accumulates human capital, causing leisure to decline. The right hand side of (5) captures the benefits of an additional unit of human capital today, in terms of increased future earnings. The left hand side captures the costs of investing in one more unit of human capital today, in terms of earnings in the current period. Equation (5) illustrates how the agent's human capital investment today depends on his future labor supply—the more he anticipates working in the future, the stronger his incentive to invest in human capital today. Due to this mechanism, individuals with a low taste for leisure choose to invest more in human capital when young, causing static differences in tastes to generate differences in labor earnings growth.

3 Data and Reduced Form Evidence

Our data come from the 1979 National Longitudinal Survey of Youth (NLSY79). We restrict our sample to strongly attached, prime-age, males, which we define to be those who we observe in the NLSY79 working in at least fourteen of the fifteen years between ages 30 and 44, and who never report working zero hours in a year.[8] For each respondent, we observe total labor income across all jobs as well as total hours worked across all jobs in the previous year.[9] We expect taste differences for our sample to provide a lower bound for taste differences in the population.

In Figure 1, we present some simple reduced form evidence that tastes for leisure vary across people and are correlated with human capital investment. To begin, the left graph of

[8]We exclude people in years when they are self-employed or working without pay. After 1996, the NLSY79 was only conducted every other year, so we count an appearance in the data after 1996 as "two years" for the purpose of inclusion in the sample.

[9]Within the context of our model, reported work hours include both paid work time as well as human capital investment. We drop observations of respondents in years when they were self-employed.

Figure 1 illustrates the amount of variation across people in the number of hours worked per week. The solid line plots mean weekly hours worked from ages 30 to 44 for highly-attached, prime-age males in the NLSY79. The dashed lines plot one standard deviation bands in weekly hours across individuals. In calculating these bands, we use a random slope model to filter out "transitory" variation in hours (coming from measurement error or other transitory shocks) and focus on the variance of "permanent" hours worked (see Section 4.4 for details). Even among highly-attached, prime-age males, and even after removing measurement error and transitory shocks, work hours vary enormously across individuals, and this variance does not decline with age.

As a second piece of evidence, consider equation (4) with no taste differences and contrast the labor supply decisions of college graduates and non-graduates. By assumption there are no differences in taste, so work hours for college graduates may be higher because their wages are higher (substitution effect) or lower because their lifetime income is higher (income effect). However, if we condition on the current wage then only the income effect remains, and we should expect college graduates to work fewer hours than non-graduates who have the same wage at age t, particularly at later ages when human capital investment is minimal.[10] The right graph of Figure 1 shows the opposite is true. We regress weekly hours worked at age t on a dummy for whether the respondent was a college graduate, controlling for the log wage at age t.[11] We find that college graduates actually work *more* hours than non-graduates with the same wage, and that this difference rises slightly over the lifecycle. Although this evidence is inconsistent with a labor supply model with identical tastes, the evidence is perfectly consistent with a model of human capital investment wherein individuals with a lower taste for leisure optimally choose to invest more in human capital (i.e. graduating from college).[12]

Finally, Table 1 provides a third piece of evidence that is consistent with variation in tastes for leisure. As seen in equation (4), the first order condition for leisure implies that, if utility is separable in consumption and leisure and there are no taste differences, the marginal utility of leisure is proportional to (i) the current wage and (ii) the marginal utility of lifetime income

[10] By conditioning on the current wage, we are comparing college graduates who received unfavorable transitory wage shocks to non-graduates who received favorable transitory shocks.

[11] The regression specification also includes a cubic polynomial in age interacted with the dummy for being a college grad.

[12] While taste differences are not the focus of their study, Imai and Keane (2004) estimate a lifecycle model of labor supply and find significant differences in the disutility of labor parameter by education level (see Table 4 of their paper).

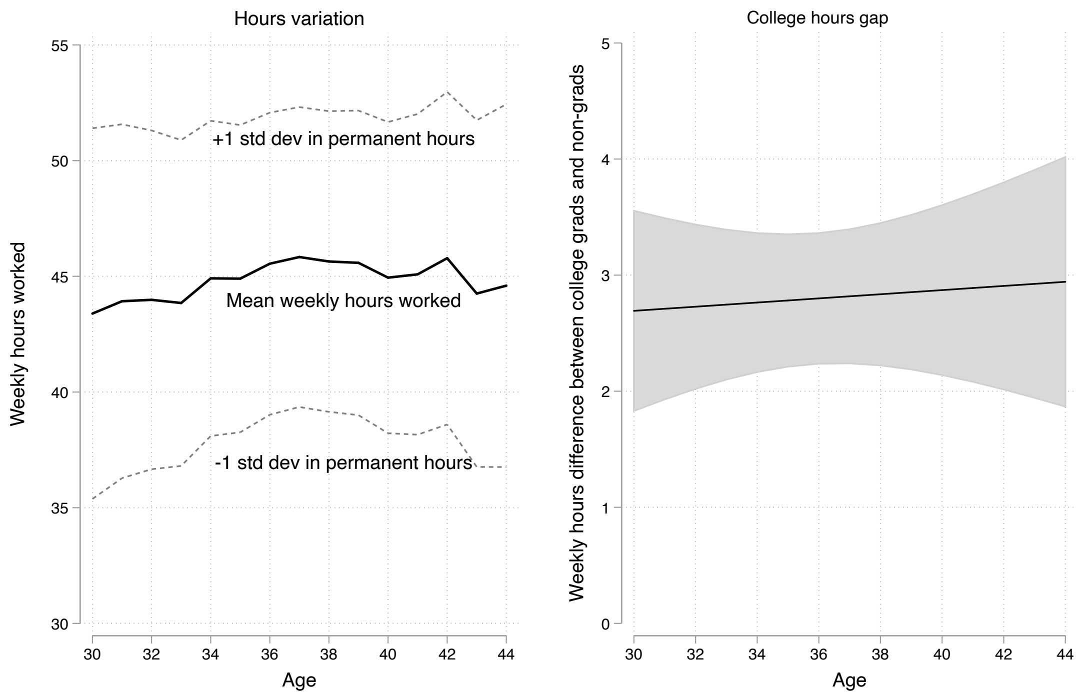

Figure 1: The left graph plots mean weekly work hours for ages 30 through 44. The dotted lines plot one standard deviation in "permanent" hours where the variance due to measurement error and other transitory shocks has been removed (see Section 4.4 for details). The right graph plots the estimated difference in weekly hours worked between college graduates and non-graduates with 95 percent confidence intervals. The graph is based on a regression of weekly hours at age t on a dummy for being a college graduate interacted with age, controlling for log wage at age t.

(λ). Therefore, after controlling for wages and lifetime income, past leisure should not predict future leisure—that is, there should be no third "omitted variable" that connects past and future leisure hours. On the other hand, if individuals differ in their tastes for leisure, then these tastes will function as an omitted variable that will introduce a positive correlation between past and future leisure. In Table 1, we report estimates of the following regression

$$\log(\ell_{i,54}) = \beta \log(\ell_{i,34}) + \delta_1 \log(w_{i,54}) + \delta_2 \log(\text{LifeInc}_i) + u_{i,54} \tag{6}$$

where $\log(\ell_{i,34})$ and $\log(\ell_{i,54})$ are (log) annual hours of leisure during ages 30–34 and 50–54, $\log(w_{i,54})$ is the (log) wage during ages 50–54, and $\log(\text{LifeInc}_i)$ is an estimate of (log) lifetime income.[13] In Table 1, β, the coefficient of interest, is positive and highly significant.[14] The estimates imply that, if you compared two workers with the same wage at age 54 and the same lifetime income but the first worker took 10 percent more leisure at age 34, then that worker will take 2.6 percent more leisure at age 54.

Table 1: Regressing Leisure at 50–54 on Leisure at 30–34

	Log leisure hours at 50–54
Log leisure hours at 30–34 (β)	0.26**
	(0.076)
Log wage at 50–54 (δ_1)	0.13***
	(0.031)
Log lifetime income (δ_2)	-0.22***
	(0.040)
Observations	778

* $p < .05$, ** $p < .01$, *** $p < .001$

Table 1: This table reports estimates from regressing log leisure hours at ages 50–54 on log leisure hours at ages 30–34, controlling for wages at ages 50–54 and lifetime income. Standard errors are reported in parentheses. Data are from the NLSY79. Sample weights were used.

In summary, using the NLSY, we find three patterns consistent with variation in tastes for leisure. First, the standard deviation of work hours among highly-attached, prime-age males is

[13] Lifetime income was calculated by taking the present value of all labor income reported between ages 18–54 using a discount rate of 5 percent. Wages at 50–54 were calculated by dividing total labor income reported between ages 50–54 by total hours worked between ages 50–54. We choose the age range 30–34 to mitigate concerns about formal schooling when young and the age range 50–54 to mitigate concerns about on-the-job training contaminating our measure of wages when older.

[14] The estimated coefficient on the (log) wage is positive due to a well-known data issue called division bias. This bias can be corrected with an instrument, but we do not worry about that here since we are not directly interested in this coefficient.

7–8 hours per week, even after removing variation due to transitory shocks and measurement error. Second, despite college graduates having higher lifetime income, they work more hours than non-graduates with the same wage, not fewer. Third, although past leisure should not predict future leisure after controlling for current wage and lifetime income, we find that past and future leisure are positively correlated. Although at odds with a model of identical tastes for leisure, these results fit naturally within a model where tastes for leisure vary.

4 Empirical Strategy

4.1 Model Specification for Estimation

To quantify the importance of taste differences, we adapt the model from Section 2 for empirical estimation. First, we specify a functional form for the agent's period utility function, using the common preferences of MaCurdy (1981). Second, we assume that agents are subject to an "HSV" tax function (Heathcote et al., 2017) parameterized by scale parameter λ and progressivity parameter τ.[15] Third, we assume agents face a borrowing constraint, so that assets B_t for every time period t cannot fall below $\overline{B}$. Agents maximize the net present value of period utility:

$$U_i(c,\ell;\phi_i) = \frac{c_t^{1-\sigma}}{1-\sigma} - \phi_i \frac{\epsilon}{1+\epsilon}(1-\ell_t)^{\frac{1+\epsilon}{\epsilon}} \tag{7}$$

Subject to period budget constraint:

$$c_t + B_t = B_{t-1} + \lambda(r_t B_{t-1} + w_t n_t)^{1-\tau} \tag{8}$$

And borrowing constraint:

$$B_t \geq \overline{B}$$

Agents face no uncertainty about the future.[16] They begin life at age 30 endowed with a triple of talent (A), taste for leisure (ϕ), and initial human capital ($\overline{k}$) and end life at age 65.

[15] The functional form for the HSV tax function is $inc_{post} = \lambda(inc_{pre})^{1-\tau}$ where inc_{pre} and inc_{post} are pre- and post-tax income. When the progressivity parameter τ equals zero, the function produces a flat tax rate of $1-\lambda$.

[16] We discuss the consequences of heterogeneous beliefs in Section 6.3.

4.2 Calibration

Although we estimate the joint distribution of talent A, taste ϕ, and initial human capital $\overline{k}$ by indirect inference, we calibrate several other parameters directly. In our baseline calibration, we choose $\gamma = 0.62$ and $\delta = 0.05$, consistent with Hendricks (2013).[17] We set σ so that the

Table 2: Calibrated Parameters

Parameter	Symbol	Value	Notes
Ben-Porath diminishing returns	γ	0.62	(Hendricks, 2013)
Human Capital depreciation	δ	0.05	(Hendricks, 2013)
Elas. of intertemporal subst.	σ	2	EIS = 0.5 (Basu and Kimball, 2002)
Discount factor	β	0.95	(Gomme and Rupert, 2007)
Human capital wage	R	20	Normalization
Borrowing Constraint	$\overline{B}$	-12	Equivalent to -$62,400

Table 2: This table reports our directly-calibrated parameter values.

elasticity of intertemporal substitution is 0.5, consistent with both long-run labor supply (Basu and Kimball, 2002) and micro-studies on the parameter (Havranek et al., 2015). The discount rate β is chosen to be 0.95, consistent with Gomme and Rupert (2007), while the net-of-tax interest rate r is $1/\beta - 1$, so that absent any financial frictions, households would choose equal consumption in every period. We also set R, a normalization constant in our model, to be 20. Thus, individual i's potential earnings in year t is given by $Rk_{it} = 20k_{it}$.[18] Finally, our borrowing constraint is set to the equivalent of $62,400. This is chosen so that the maximum debt service ratio an average-income household could accrue immediately is 8%, which is near the 95th percentile of the ratio of revolving debt payments to income for 30 year-olds in the Survey of Consumer Finances (SCF).[19]

[17] While some of the literature suggests a Ben-Porath technology that's nearly linear with near-zero depreciation to match changes in the wage distribution under skill-biased technological change (see, for instance, Guvenen and Kuruscu (2012)), Hendricks (2013) models schooling choice closely and finds that, because near-linear models with zero depreciation see no human capital accumulation after age 45, they (incorrectly) predict near-perfect comovements of the wage profiles of older cohorts.

[18] For convenience when interpreting results, we convert fraction of annual hours worked into annual hours worked. Doing so changes the scaling of human capital so that Rk_{it} can now by interpreted as an hourly wage. Consequently, in our results we will report the hourly wage Rk_{it}, rather than raw human capital.

[19] Based on SCF variable PIRREV for 30 year-old households in 1989 (when members of the NLSY79 would have been approximately 30 years old).

4.3 Identification

In this section we illustrate how, within our model, the lifecycle paths of earnings *and* work hours allow us to disentangle tastes and talent. In Figure 2 we plot the paths of earnings and work hours for an agent with triple $\{A = 0.25, \phi = 0.73, \overline{k} = 33.7\}$.[20] We also plot paths for two alternative triples. In the first, we lower the taste parameter ϕ by 10 percent which raises the lifecycle earnings path; in the second, we increase talent A and initial human capital $\overline{k}$ to match this alternative earnings path. Now, imagine two agents: one which follows the baseline earnings path in Figure 2 and another which follows the alternative earnings path. From observations on lifecycle earnings alone we cannot tell whether these agents differ in tastes or in talent, but, as illustrated in the right panel of Figure 2, we can distinguish taste from talent if we also observe lifecycle work hours. If a person were earning more over the lifecycle because he had a lower taste for leisure, then he would work more hours, while if he were earning more because he is more talented, then he would work fewer hours at early ages.[21] Thus, despite producing the same lifecycle *earnings* paths, the lower-taste-for-leisure and higher-talent triples generate different paths for *work hours*. Lowering the taste for leisure shifts work hours up. In contrast, increasing talent actually lowers work hours early in life due to both an income effect and an intertemporal substitution of work towards later periods of higher wages. Thus, we can distinguish taste from talent by looking at the lifecycle paths of both earnings and work hours together. Indeed, the strong positive correlation in the data between earnings and work hours at age 30 is a key reason why we find that taste differences play a large role in income inequality.

4.4 Estimation With Indirect Inference

We estimate the joint distribution of A, ϕ, and $\overline{k}$ by indirect inference. Our auxiliary statistical model is a random slope model, augmented with year controls. The model estimates the means, variances, and covariances of the levels and slopes of both log earnings and work hours. Then we solve for a joint distribution over A_i, ϕ_i, and $\overline{k}_i$ that matches those moments as closely as possible when we simulate the model.

We fit the following auxiliary statistical model to lifecycle earnings and work hours between

[20] The kinks in the earnings and work hours paths occur because the agent hits his borrowing constraint.

[21] Although it is not shown in Figure 2, higher talent may raise work hours at later ages as wages continue to rise.

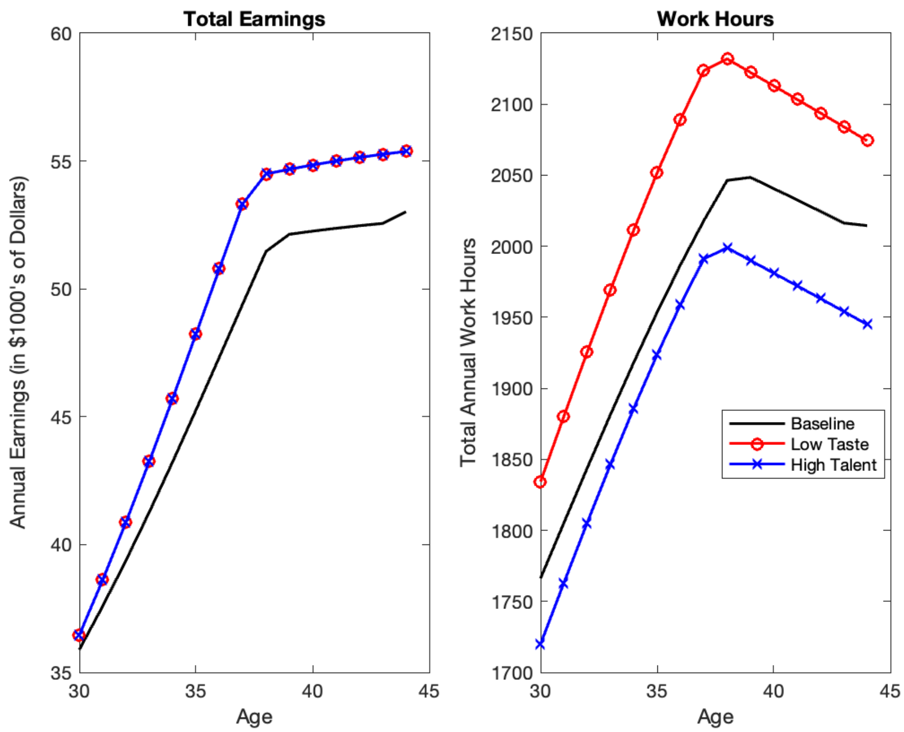

Figure 2: This figure solves the model for three separate triples $(A, \phi, \overline{k})$: a baseline triple (black lines), a low-taste-for-leisure triple (red lines) wherein ϕ is lower than in baseline, and a high-talent triple (blue lines). The high-talent triple was set by choosing talent A and initial human capital $\overline{k}$ to match the earnings path of the low-taste triple. Although the low-taste and high-talent triples are observationally equivalent when looking at earnings alone, they can be distinguished by including data on work hours.

ages 30 and 44.

$$\log(earn_{ity}) = \alpha_{i,earn} + \beta_{i,earn}t + \zeta_y + u_{it} \tag{9}$$

$$work_{ity} = \alpha_{i,work} + \beta_{i,work}t + \xi_y + v_{it} \tag{10}$$

$earn_{ity}$ is total labor earnings[22] of person i at age t in year y while $work_{ity}$ is total weekly work hours (including human capital investment). $t = 0$ corresponds to age 30, so the intercept $\alpha_{i,earn}$ represents person i's expected log earnings at age 30. Each person i has his own lifecycle earnings profile governed by $\alpha_{i,earn}$ and $\beta_{i,earn}$, but each year his earnings are also subject to a shock u_{it} which captures both transitory shocks and measurement error. Finally,

[22]Labor earnings are converted to 2012 dollars using the CPI-U.

we also allow for a year-specific effect ζ_y to capture aggregate economic fluctuations.[23] We estimate an analogous specification for work hours. For each person i, the coefficient vector $(\alpha_{earn}, \beta_{earn}, \alpha_{work}, \beta_{work}) \sim N(\mu, \Sigma)$ while the residual vector $(u_{it}, v_{it}) \sim N(0, \Omega)$. In other words, the intercepts and slopes are all jointly normal with some mean μ and covariance matrix Σ. The residuals for earnings and hours, u_{it} and v_{it}, are also jointly normal and are allowed to be correlated, but residuals are assumed to be independent over time and across individuals. We use maximum likelihood to estimate μ, Σ, and Ω, as well as the year fixed effects ζ and ξ. For indirect inference, we focus on targeting μ and Σ which gives us fourteen moments to match (four means, four variances, and six covariances). Let $\widehat{\theta}$ denote a vector containing all fourteen of these estimated moments and $\mathbf{V}$ a diagonal matrix containing the variances of each estimated moment. The column labeled "Data" in Table 4 reports the estimates $\widehat{\theta} \equiv (\widehat{\mu}, \widehat{\Sigma})$ from the model in equations (9) and (10).

With an estimate $\widehat{\theta}$ of the parameters of the auxiliary statistical model in equations (9) and (10), we estimate the economic model as follows. Let $\Gamma = \{\{(A_i, \phi_i, \overline{k}_i)\}_{i=1}^{10}, \epsilon, \lambda, \tau\}$ represent a vector of ten triples, the Frisch elasticity of labor supply, and the two tax parameters. Given the calibrated parameter values in Table 2, we choose Γ to minimize the following objective function subject to perfectly matching average and marginal tax rates:

$$\min_{\Gamma} \quad \left(\theta(\Gamma) - \widehat{\theta}\right)' \mathbf{V}^{-1} \left(\theta(\Gamma) - \widehat{\theta}\right) \quad \text{s.t. } \mathcal{T}(\Gamma) = \widehat{\mathcal{T}}$$

where $\theta(\Gamma) \equiv (\mu(\Gamma), \Sigma(\Gamma))$ denotes the moments obtained from the auxiliary statistical model in (9) and (10) after simulating data from the model for each of the ten triples in Γ. $\mathcal{T}(\Gamma)$ denotes the median marginal and median average tax rates in our model simulation while $\widehat{\mathcal{T}}$ denotes the median marginal and median average tax rates for all households, reported in Table 4 of Guner et al. (2014). The ten triples in the vector Γ serve as a discrete approximation of the continuous joint distribution of $(A, \phi, \overline{k})$.[24] This allows us to avoid imposing functional form restrictions on the joint distribution of talent, tastes, and initial human capital.

[23] The NLSY79 allows us to distinguish age effects from year effects because not all respondents were born in the same year. In appendix A, we include occupation fixed effects to account for the possibility that, for some occupations, hours and earnings may be bundled. For example, the high earnings of heart surgeons and CEO's may come paired with high hours. Our findings are not sensitive to including or excluding these occupation controls.

[24] Kennan (2006) shows that the optimal discrete approximation of a continuous distribution is one with equal weights.

Table 3: Estimated Means, Variances, and Covariances

	Talent (A)	Taste (ϕ)	Initial Human Capital ($R\overline{k}$)
Means	0.251	0.725	33.652
		Covariances	
Talent (A)	0.011	-0.007	0.995
Taste (ϕ)		0.215	-4.548
Initial Human Capital ($R\overline{k}$)			563.409

Table 3: This table depicts means, variances, and covariances for our ten estimated triples depicted graphically in Figure 3

5 Estimation Results

We find substantial variation in tastes for leisure. In Table 3, we report the means, variances, and covariances of talent A, taste ϕ, and initial human capital $\overline{k}$. The coefficient of variation for taste is estimated to be 0.64. Figure 3 plots three bivariate scatter plots of the ten triples in $\widehat{\Gamma}$. Perhaps unsurprisingly, human capital at age 30 is positively correlated with talent, but both are negatively correlated with taste for leisure. This makes sense if human capital at age 30 were the result of earlier investments, which were themselves driven by talent and tastes. It is helpful to look at some triples in isolation to gain some intuition. The triple with the highest talent also has the second highest initial human capital but the second lowest taste for leisure. An individual with this triple would work many hours and invest heavily in human capital early on, thereby causing his wages and earnings to rise quickly, peaking around age 55. Over time, he would reduce his human capital investment while increasing both consumption and leisure. In contrast, the triple with the lowest talent also has low initial human capital and a relatively high taste for leisure. An individual with this triple would invest little in human capital, causing his wage and earnings to slowly fall as his human capital depreciates. Despite his high taste for leisure, he works and saves aggressively early in life, but over time his labor supply falls and he draws down his savings so that his consumption is flat over the lifecycle. The majority of the triples are clustered around the same values for talent and initial human capital but differ widely in taste for leisure. Although these individuals have quite similar earnings paths over the lifecycle, their work hours differ substantially.

Figure 3 also reports our estimates of ϵ and our two tax parameters λ and τ. Note that in the presence of a binding borrowing constraint, the ϵ can no longer be interpreted as a traditional Frisch elasticity of labor supply. We estimate a value of 1.57 for this parameter which would traditionally mean that work hours were very responsive to transitory wage shocks. However, the borrowing constraint tempers that response by forcing agents to work in early years, even though their wages are relatively low, in order to finance their consumption. The combined effect of a large ϵ and binding borrowing constraint is that we are able to match the relatively flat average growth in work hours over the lifecycle, despite agents experiencing significant growth in their earnings. The tax parameters λ and τ are pinned down by matching the mean marginal and average tax rates in the data. Under our estimated tax parameters, an agent with \$100,000 of pre-tax earnings would face a tax bill of \$13,849, giving him an average tax rate of 13.8 percent and a marginal tax rate of 21.7 percent.[25]

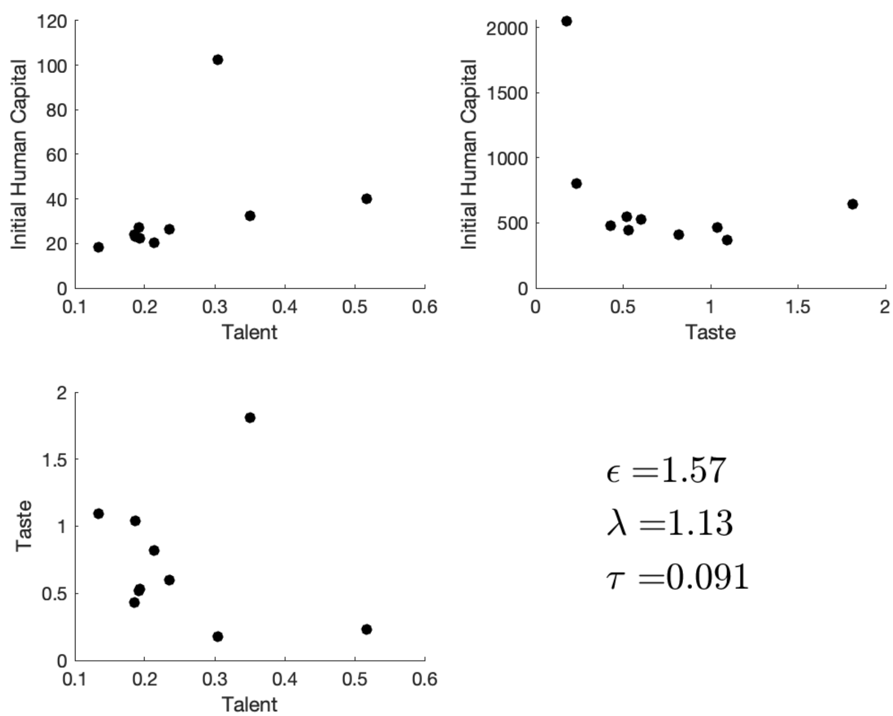

Figure 3: This figure plots our ten estimated triples of talent A, taste for leisure ϕ, and initial human capital at age 30 as an hourly wage $R\bar{k}$. It also reports our estimates of the Frisch elasticity ϵ and the two tax parameters λ and τ.

Table 4 reports our model's goodness of fit. Column 1 of Table 4 reports each of the estimated moments in $\widehat{\theta}$ from the auxiliary statistical model. Column 2 reports $\theta(\widehat{\Gamma})$, which are the

[25]Post-tax earnings are given by $earn_{post} = 5200\lambda(earn_{pre}/5200)^{1-\tau}$ (the scaling by 5200 is a convenient normalization). The average tax rate is thus $1 - \frac{earn_{post}}{earn_{pre}}$ and the marginal rate is $1 - \lambda(1-\tau)(earn_{pre}/5200)^{-\tau}$.

corresponding moments that come from simulating data from the model under $\widehat{\Gamma}$. It also reports in square brackets the difference between each simulated and estimated moment, divided by the standard error of the estimated moment. Our model fits the estimated moments quite well, resulting in an objective value of only 0.02. In contrast, column 3 reports the simulated moments when we constrain the model to have no taste differences by forcing each triple to have the same value for ϕ. Without taste differences the model cannot replicate several of the estimated moments, especially the positive covariance between the intercepts α_{earn} and α_{work}, and the objective function is much larger at 186.7. Lastly, Column 4 of Table 4 comes from the same no-taste-differences model as in column 3, but forced to match the earnings moments: Mean(α_{earn}), Mean(β_{earn}), Var(α_{earn}), Var(β_{earn}), and Cov($\alpha_{earn}, \beta_{earn}$)[26]. As a result, the estimator misses badly on moments involving work hours. Column 4 shows that, while the no-taste-differences model can closely match both the mean and variance of earnings over the lifecycle, it cannot simultaneously match the data on work hours.

The auxiliary model in equations (9) and (10) implies the following linear (quadratic) relationship for the means (variances and covariance) of log earnings and work hours over the lifecycle.

$$\begin{aligned}
\mathbb{E}[\log(earn_t)] &= \mathbb{E}[\alpha_{earn}] + \mathbb{E}[\beta_{earn}]t && (11)\\
\mathbb{E}[work_t] &= \mathbb{E}[\alpha_{work}] + \mathbb{E}[\beta_{work}]t && (12)\\
\mathrm{Var}[\log(earn_t)] &= \mathrm{Var}[\alpha_{earn}] + 2\,\mathrm{Cov}[\alpha_{earn}, \beta_{earn}]t + \mathrm{Var}[\beta_{earn}]t^2 && (13)\\
\mathrm{Var}[work_t] &= \mathrm{Var}[\alpha_{work}] + 2\,\mathrm{Cov}[\alpha_{work}, \beta_{work}]t + \mathrm{Var}[\beta_{work}]t^2 && (14)\\
\mathrm{Cov}[\log(earn_t), work_t] &= \mathrm{Cov}[\alpha_{earn}, \alpha_{work}]\\
&\quad + (\mathrm{Cov}[\alpha_{earn}, \beta_{work}] + \mathrm{Cov}[\alpha_{work}, \beta_{earn}])\,t\\
&\quad + \mathrm{Cov}[\beta_{earn}, \beta_{work}]t^2 && (15)
\end{aligned}$$

By substituting the estimated means, variances, and covariances from column one of Table 4 into equations (11)–(15), Figure 4 plots in red the implied lifecycle paths of the means, variances, and covariance of log earnings and work hours. It plots in black the corresponding paths from simulating our model (column 2 of Table 4). As we might expect from how closely the simulated

[26]This is done by downweighting all other moments by a factor of 10,000, so the model still attempts, after matching earnings moments, to match labor moments.

Table 4: Estimated and Simulated Moments

$\log(earn_{ity}) = \alpha_{i,earn} + \beta_{i,earn}t + \zeta_y + u_{it}$ $work_{ity} = \alpha_{i,work} + \beta_{i,work}t + \xi_y + v_{it}$				
Moment	Estimated	Simulated		
		Main Model	No Taste Diff	
			Target All	Target Earn
Mean(α_{earn})	10.63	10.63	10.64	10.63
	(0.02)	[0.00]	[0.23]	[−0.00]
Mean(β_{earn})	0.024	0.024	0.015	0.024
	(0.007)	[−0.03]	[−1.37]	[−0.01]
Mean(α_{work})	43.83	43.83	43.88	58.15
	(0.26)	[0.00]	[0.21]	[55.43]
Mean(β_{work})	-0.017	-0.008	0.060	0.111
	(0.093)	[0.10]	[0.83]	[1.38]
Var(α_{earn})	0.29	0.29	0.22	0.29
	(0.01)	[0.00]	[−5.81]	[0.00]
Var(β_{earn})	0.00119	0.00119	0.00122	0.00119
	(0.00007)	[0.04]	[0.43]	[0.00]
Var(α_{work})	70.23	70.24	65.94	149.90
	(3.62)	[0.00]	[−1.18]	[22.02]
Var(β_{work})	0.439	0.437	0.447	0.478
	(0.034)	[−0.04]	[0.25]	[1.15]
Cov(α_{earn},β_{earn})	-0.0039	-0.0039	-0.0112	-0.0039
	(0.0013)	[0.03]	[−5.61]	[−0.00]
Cov(α_{earn},α_{work})	2.04	2.04	-0.76	-4.78
	(0.28)	[−0.00]	[−10.02]	[−24.45]
Cov(α_{earn},β_{work})	-0.11	-0.11	-0.17	-0.18
	(0.03)	[−0.01]	[−2.13]	[−2.49]
Cov(β_{earn},α_{work})	-0.05	-0.05	-0.12	-0.13
	(0.02)	[−0.01]	[−2.95]	[−3.50]
Cov(β_{earn},β_{work})	0.010	0.010	0.014	-0.001
	(0.002)	[−0.01]	[1.84]	[−5.28]
Cov(α_{work},β_{work})	-3.21	-3.21	-3.21	0.15
	(0.02)	[−0.01]	[0.03]	[155.75]
Median Average Tax Rate	0.07	0.07	0.07	0.07
Median Marginal Tax Rate	0.15	0.15	0.15	0.15
Objective Value		0.02	186.65	28462

Table 4: This table reports estimated moments from the data along with their simulated counterparts. Standard errors are reported in parentheses in column 1. Columns 2–4 report simulated model moments. In square brackets we we report the difference between the moment estimated from the data and the model simulated moment, divided by the standard error of the estimate. Data are from the NLSY79. Sample weights were used.

and estimated moments match, the estimated and simulated paths are nearly identical. For comparison, Figure 4 also plots in blue the implied paths from our no-taste-differences model (column 3 of Table 4). Without taste differences, the simulated paths deviate significantly from the estimated paths in several respects. First, the no-taste-differences model cannot reproduce the slight downward trend in mean work hours between ages 30 and 44. Second, it understates the variances of both log earnings and work hours. And third, it predicts a negative covariance between log earnings and work hours as opposed to the positive covariance estimated from the data.

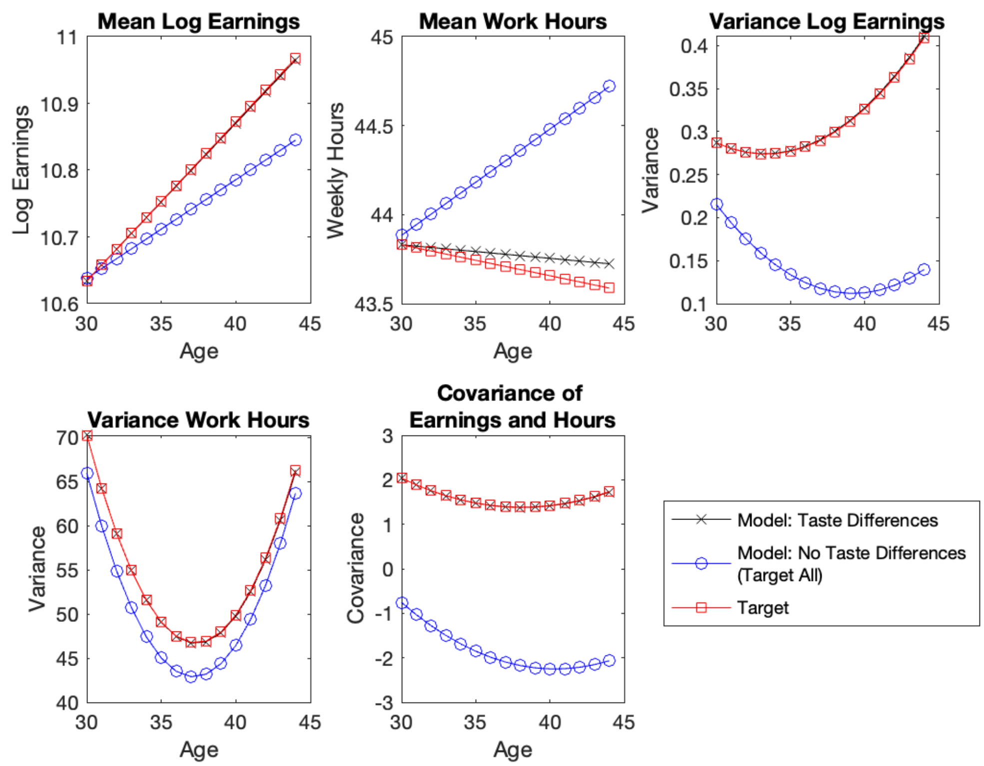

Figure 4: This figure plots in red the implied paths of the means, variances, and covariance of log earnings and work hours based on the estimated moments in Table 4. The black lines plot the fit of our model with taste differences (column 2 of Table 4), while the blue lines with circles plot the fit of a model with no taste differences that targets all moments (column 3 of Table 4).

We also evaluate the model's fit by estimating and simulating four untargeted moments, reported in Table 5. The first three untargeted moments come from the regression reported in Table 1 of log leisure at age 50–54 on log leisure at ages 30–34, log wage at ages 50–54, and lifetime earnings. Our model replicates the coefficient on log leisure reasonably well and misses on the coefficients of log wage and lifetime income while the no-taste-differences models miss badly on all three coefficients. That said, the estimated values of δ_1 and δ_2 may be impacted by division bias, so we might not expect any of the models to match those moments. For our fourth untargeted moment, we regress a person's rank of earnings at ages 50–54 on the rank of his earnings at ages 30–34. The idea behind this moment is that it should tell us something about the degree to which earnings paths cross between ages 30–54. If there are many crossings, then the rank correlation should be small or even negative. We estimate a coefficient of 0.64, which indicates relatively few crossings in this age range. Our model produces a rank correlation of 0.42, which, although statistically different, is qualitatively similar. The no-taste-differences model that is fitted to all the moments in Table 4 produces a rank correlation of only 0.20, implying much more crossing than in the data. The no-taste-differences model matches this moment much better when it is forced to match earnings moments, although as seen above, doing so causes it to miss other moments badly. Because the two models that match the earnings moments in Table 4 produce very similar rank correlations, we interpret this as evidence that our model, estimated only using data for ages 30–44, still provides a reasonable fit for earnings in later years. The bottom row of Table 5 reports the sum of squared differences between the estimated and simulated moments, inversely weighted by the variance of each estimated moment. As in Table 4, the model with taste differences significantly outperforms the no-taste-differences models.

6 Discussion

6.1 How much does taste variation contribute to earnings variation?

To understand the effects of talent, taste, and initial human capital on the earnings distribution, we perform the following decomposition. Denoting the mean, variance, and skewness of log earnings at age t by $\mu_{E,t}$, $V_{E,t}$, and $\mathcal{S}_{E,t}$, we totally differentiate each with respect to the

Table 5: Untargeted Moments

$\log(\ell_{i,54}) = \beta \log(\ell_{i,34}) + \delta_1 \log(w_{i,54}) + \delta_2 \log(\text{LifeInc}_i) + u_{i,54}$ $IncRank_{i,54} = \gamma IncRank_{i,34} + u_{i,54}$				
Moment	Estimated	Simulated		
		Main Model	No Taste Het	
			Target All	Target Earn
β	0.26	0.21	-0.82	-0.40
	(0.08)	$[-0.67]$	$[-14.22]$	$[-8.68]$
δ_1	0.13	-0.02	-0.01	-0.39
	(0.03)	$[-4.82]$	$[-4.62]$	$[-16.91]$
δ_2	-0.22	-0.03	0.58	1.30
	(0.04)	[4.73]	[19.99]	[37.98]
γ	0.64	0.42	0.20	0.49
	(0.03)	$[-7.80]$	$[-15.33]$	$[-5.29]$
Objective Value		106.87	858.05	1832.28

Table 5: This table reports untargeted moments from the data along with their simulated counterparts. Standard errors are reported in parentheses in column 1. Columns 2–4 report the difference between each simulated and estimated moment, divided by the standard error, in square brackets. Data are from the NLSY79. Sample weights were used.

variances of taste, talent, and initial human capital (V_A, V_ϕ, $V_{\overline{k}}$):

$$
\begin{aligned}
d\mu_{E,t} &= \frac{\partial \mu_{E,t}}{\partial V_A} dV_A + \frac{\partial \mu_{E,t}}{\partial V_\phi} dV_\phi + \frac{\partial \mu_{E,t}}{\partial V_{\overline{k}}} dV_{\overline{k}} \\
dV_{E,t} &= \frac{\partial V_{E,t}}{\partial V_A} dV_A + \frac{\partial V_{E,t}}{\partial V_\phi} dV_\phi + \frac{\partial V_{E,t}}{\partial V_{\overline{k}}} dV_{\overline{k}} \\
d\mathcal{S}_{E,t} &= \frac{\partial \mathcal{S}_{E,t}}{\partial V_A} dV_A + \frac{\partial \mathcal{S}_{E,t}}{\partial V_\phi} dV_\phi + \frac{\partial \mathcal{S}_{E,t}}{\partial V_{\overline{k}}} dV_{\overline{k}}
\end{aligned}
$$

Or, denoting the percentage change in each variable by Δ, this becomes

$$\Delta \mu_{E,t} = \epsilon^t_{\mu,A} \Delta V_A + \epsilon^t_{\mu,\phi} \Delta V_\phi + \epsilon^t_{\mu,\overline{k}} \Delta V_{\overline{k}} \tag{16}$$

$$\Delta V_{E,t} = \epsilon^t_{V,A} \Delta V_A + \epsilon^t_{V,\phi} \Delta V_\phi + \epsilon^t_{V,\overline{k}} \Delta V_{\overline{k}} \tag{17}$$

$$\Delta \mathcal{S}_{E,t} = \epsilon^t_{\mathcal{S},A} \Delta V_A + \epsilon^t_{\mathcal{S},\phi} \Delta V_\phi + \epsilon^t_{\mathcal{S},\overline{k}} \Delta V_{\overline{k}} \tag{18}$$

where $\epsilon^t_{\mu,A}$, $\epsilon^t_{V,A}$, and $\epsilon^t_{\mathcal{S},A}$ denote the elasticities of the mean, variance, and skewness of log earnings at age t with respect to the variance of talent, and similarly for the variance of taste

V_ϕ and initial human capital $V_{\overline{k}}$.[27] Table 6 reports the elasticities of the mean, variance, and skewness of log earnings with respect to the variances of talent, taste, and initial human capital. It breaks these elasticities out into ages 30–44, 45–65, and 30–65. Increasing the variances of talent, taste, and leisure individually have negligible effects on the mean of log earnings. But when it comes to the variance of log earnings, a one percent increase in the variance of talent raises the variance of log earnings by 0.26 percent at early ages, by 1.48 percent at later ages, and by 1.24 for ages 30–65. At early ages, increasing the variance of tastes has a much larger effect on the variance of earnings (1.69 percent) than talent, but a smaller effect (1.30 percent) at older ages, resulting in an overall effect (1.38 percent) that is only slightly larger than that of talent. For small perturbations, taste variation is slightly more important than talent variation in explaining lifecycle income variance. The effects of talent and tastes on the skewness of earnings differ for early and later ages. Overall, increasing the variance of talent raises the skewness of log earnings by 3.82 percent, while increasing the variance of tastes lowers the skewness of income by 3.86 percent.

Table 6: Moment Elasticities

Age	30-44	45-65	30-65
Panel A: Mean Income			
Talent: A	0.001	-0.036	-0.021
Taste: ϕ	0.024	0.029	0.027
Initial Human Capital: $\overline{k}$	-0.014	0.011	0.001
Panel B: Variance of Income			
Talent: A	0.26	1.48	1.24
Taste: ϕ	1.69	1.30	1.38
Initial Human Capital: $\overline{k}$	0.68	-0.44	-0.21
Panel C: Skewness of Income			
Talent: A	-0.32	1.50	3.82
Taste: ϕ	0.39	-2.51	-3.86
Initial Human Capital: $\overline{k}$	0.21	-0.87	-2.79

Table 6: Panels A, B, and C report the estimated elasticities from equations (16), (17), and (18). Column 1 reports the elasticities when pooling together ages 30–44, column 2 reports them for ages 45–65, and column 3 reports them for all ages 30–65.

[27] There are many ways to calculate these elasticities. For instance, to calculate $\epsilon_{A,t}$, we could hold V_ϕ and $V_{\overline{k}}$ constant at the baseline level, change V_A by 1% and calculate the change in income variance. Or we could do the same thing but holding V_ϕ and $V_{\overline{k}}$ constant at 1% above baseline for the entire exercise. Numerically, it makes little difference.

6.2 Why are tastes so important?

What moments in the data drive our results? That is, what would have to change about the data in order for our model to find a small role for taste differences in generating income inequality? In this section, we show that adjusting just two key empirical moments would dramatically reduce the estimated variance of tastes and would dramatically improve the fit of the no-taste-differences model, which implies two things. First, any data-oriented explanation for our findings (i.e. mismeasurement of various kinds) should focus on these key moments. Second, any theory-oriented explanation for our findings (i.e. that our model is missing something important) needs to be able to account for these two moments.

The first key moment is the covariance between log earnings and work hours. The intuition from Figure 2 tells us that, higher earnings due to talent or initial human capital should correlate with fewer work hours at age 30 while higher earnings due to tastes for leisure should correlate with more work hours. And as we can see in Figure 4, the no-taste-differences model has a very difficult time matching the positive covariance between log earnings and work hours that we observe in the data. From equation (15) we see that the covariance between log earnings and work hours at age 30 is given by $\text{Cov}[\alpha_{earn}, \alpha_{work}]$, and from Table 4 we can see that this moment is precisely where the no-taste-differences model misses the most. But suppose that $\text{Cov}[\alpha_{earn}, \alpha_{work}]$, rather than being strongly positive, had instead been zero. How would changing this single moment have changed our findings about the role of tastes?

We refit the model, but this time targeting $\text{Cov}[\alpha_{earn}, \alpha_{work}] = 0$ instead of its estimated value of 2.04. Table 7 reports the means, variances and covariances of talent A, tastes ϕ, and initial human capital $\overline{k}$ once we set this single moment to zero. The variance of tastes falls by half while the variances of talent and initial human capital both rise. Moreover, Table 8 shows that the fit of the no-taste-differences model improves considerably after setting this covariance to zero. This improvement is partly mechanical since we are moving the moment closer to the value that the no-taste-differences model was already producing. But some of this improvement comes because the no-taste-differences model is no longer driven to miss on other moments in an attempt to try to match the large positive covariance between log earnings and work hours.

Figure 4 shows that the no-taste-differences model also misses on both the level and slope

Table 7: Estimated Means, Variances, and Covariances With Altered Moments

	Talent (A)	Taste (ϕ)	Initial Human Capital ($R\bar{k}$)
Panel A: Baseline			
Means	0.251	0.725	33.652
		Covariances	
Talent (A)	0.011	-0.007	0.995
Taste (ϕ)		0.215	-4.548
Initial Human Capital ($R\bar{k}$)			563.409
Panel B: $Cov[\alpha_{earn}, \alpha_{lab}] = 0$			
Means	0.258	0.676	35.912
		Covariances	
Talent (A)	0.015	-0.010	1.893
Taste (ϕ)		0.104	-3.943
Initial Human Capital ($R\bar{k}$)			905.779
Panel C: $Cov[\alpha_{earn}, \alpha_{lab}] = 0$, $Cov[\alpha_{earn}, \beta_{earn}] = -0.01$			
Means	0.254	0.646	35.912
		Covariances	
Talent (A)	0.011	-0.007	1.515
Taste (ϕ)		0.065	-3.169
Initial Human Capital ($R\bar{k}$)			808.703

Table 7: This table reports means, variances, and covariances for our ten estimated triples, when fitting to two sets of altered moments. Panel A replicates our baseline results reported in Table 3. Panel B reports the results after setting $\text{Cov}(\alpha_{earn}, \alpha_{work}) = 0$, rather than its estimated value of 2.04. Panel C reports the results after also setting $\text{Cov}(\alpha_{earn}, \beta_{earn}) = -0.01$ rather than its estimated value of -0.0039.

Table 8: Estimated and Simulated Moments With Cov($\alpha_{earn}, \alpha_{work}$) = 0

$\log(earn_{ity}) = \alpha_{i,earn} + \beta_{i,earn}t + \zeta_y + u_{it}$

$work_{ity} = \alpha_{i,work} + \beta_{i,work}t + \xi_y + v_{it}$

Moment	Estimated	Simulated		
		Main Model	No Taste Diff	
			Target All	Target Earn
Mean(α_{earn})	10.63	10.63	10.63	10.63
	(0.02)	[0.00]	[−0.05]	[−0.01]
Mean(β_{earn})	0.024	0.024	0.017	0.024
	(0.007)	[−0.00]	[−1.06]	[−0.00]
Mean(α_{work})	43.83	43.83	43.78	58.14
	(0.26)	[−0.00]	[−0.19]	[55.38]
Mean(β_{work})	-0.017	-0.017	0.090	0.113
	(0.093)	[0.00]	[1.15]	[1.40]
Var(α_{earn})	0.29	0.29	0.24	0.29
	(0.01)	[0.00]	[−4.09]	[−0.01]
Var(β_{earn})	0.00119	0.00119	0.00119	0.00118
	(0.00007)	[0.00]	[0.01]	[−0.01]
Var(α_{work})	70.23	70.23	72.53	148.21
	(3.62)	[0.00]	[0.64]	[21.56]
Var(β_{work})	0.439	0.439	0.439	0.475
	(0.034)	[−0.00]	[0.02]	[1.06]
Cov(α_{earn},β_{earn})	-0.0039	-0.0039	-0.0110	-0.0039
	(0.0013)	[0.00]	[−5.44]	[−0.00]
Cov(α_{earn},α_{work})	0.00	0.00	-1.33	-4.78
	(0.28)	[0.00]	[−4.77]	[−17.15]
Cov(α_{earn},β_{work})	-0.11	-0.11	-0.15	-0.18
	(0.03)	[−0.00]	[−1.37]	[−2.49]
Cov(β_{earn},α_{work})	-0.05	-0.05	-0.10	-0.13
	(0.02)	[0.00]	[−2.19]	[−3.37]
Cov(β_{earn},β_{work})	0.010	0.010	0.012	-0.002
	(0.002)	[−0.00]	[0.96]	[−5.39]
Cov(α_{work},β_{work})	-3.21	-3.21	-3.21	0.24
	(0.02)	[0.00]	[0.20]	[160.13]
Median Average Tax Rate	0.07	0.07	0.07	0.07
Median Marginal Tax Rate	0.15	0.15	0.15	0.15
Objective Value		0.00	79.58	29519

Table 8: This table reports estimated moments from the data along with their simulated counterparts. Cov($\alpha_{earn}, \alpha_{work}$) is set to zero rather than its estimated value of 2.04. All other moments are the same as in Table 4. Standard errors are reported in parentheses in column 1. Columns 2–4 report the difference between each simulated and estimated moment, divided by the standard error, in square brackets. Data are from the NLSY79. Sample weights were used.

of the variance of log earnings over the lifecycle. Unlike the positive covariance of log earnings and work hours, the no-taste-differences model *is* able to match these moments (see column 4 of Table 4), but it misses them because matching them closely would cause it to miss several other moments. Thus, on top of zeroing out Cov$[\alpha_{earn}, \alpha_{work}]$, in Table 9 we also alter one additional moment—we change the covariance of the intercepts and slopes of log earnings from -0.0039 to -0.01, which corresponds to changing their correlation from -0.21 to -0.54. Table 9 shows that setting both Cov$[\alpha_{earn}, \alpha_{work}] = 0$ and $Cov(\alpha_{earn}, \beta_{earn}) = -0.01$ further reduces the variance of tastes, and Table 9 shows that changing both moments further improves the fit of the no-taste-differences model.

6.3 Discussion and Robustness

We anticipate several critiques of our findings, though we cannot hope to anticipate them all. In this section, we address some of the most common.

First, perhaps work hours are poorly measured due to outliers that are inflating the variance of earnings. Note that, if this were true, it would not necessarily increase the covariance of hours and earnings nor decrease the covariance between the level and slope of log earnings (the key moments driving our results). That said, in Appendix B, we perform a series of robustness checks to deal with extreme values for work hours. We try various combinations of dropping observations with extremely low hours, censoring low hours, and censoring very high hours. None of these robustness checks meaningfully alters our finding that tastes play a large role in driving earnings inequality.

Second, one plausible explanation for the large positive covariance between log earnings and work hours is that earnings and hours come in occupation-specific bundles.[28] For example, the high earnings of heart surgeons and CEO's may come paired with high hours just due to the nature of the job. If earnings and hours come in bundles, where high hours come paired with high earnings, then this could explain why we find such a large positive correlation of earnings and

[28] It's not immediately clear that occupations themselves pose a problem. As demonstrated by Rosen (1972), a Ben-Porath model of lifecycle human capital investment is equivalent to a model of learning-by-doing with a large number of occupations, each offering different learning opportunities. Occupations with more learning opportunities tend to attract younger workers. In equilibrium, these occupations pay lower wages whereas occupations with fewer learning opportunities pay higher wages and tend to be filled by older workers. Thus, the Ben-Porath framework can be interpreted as a model of endogenous occupation choice over the lifecycle wherein occupations are the means by which workers make human capital investments.

Table 9: Est. and Sim. Moments With Cov($\alpha_{earn}, \alpha_{work}$) = 0 and Cov($\alpha_{earn}, \beta_{earn}$) = −0.01

$$\log(earn_{ity}) = \alpha_{i,earn} + \beta_{i,earn}t + \zeta_y + u_{it}$$
$$work_{ity} = \alpha_{i,work} + \beta_{i,work}t + \xi_y + v_{it}$$

Moment	Estimated	Simulated		
		Main Model	No Taste Diff	
			Target All	Target Earn
Mean(α_{earn})	10.63	10.64	10.63	10.63
	(0.02)	[0.06]	[0.03]	[0.00]
Mean(β_{earn})	0.024	0.023	0.016	0.024
	(0.007)	[−0.07]	[−1.15]	[−0.01]
Mean(α_{work})	43.83	43.86	43.82	53.83
	(0.26)	[0.11]	[−0.04]	[38.71]
Mean(β_{work})	-0.017	0.004	0.083	-0.044
	(0.093)	[0.23]	[1.08]	[−0.29]
Var(α_{earn})	0.29	0.29	0.25	0.29
	(0.01)	[0.27]	[−2.85]	[−0.00]
Var(β_{earn})	0.00119	0.00125	0.00125	0.00119
	(0.00007)	[0.91]	[0.84]	[0.01]
Var(α_{work})	70.23	70.59	69.80	99.90
	(3.62)	[0.10]	[−0.12]	[8.20]
Var(β_{work})	0.439	0.410	0.429	0.191
	(0.034)	[−0.86]	[−0.28]	[−7.34]
Cov(α_{earn},β_{earn})	-0.0100	-0.0081	-0.0128	-0.0100
	(0.0013)	[1.44]	[−2.12]	[−0.00]
Cov(α_{earn},α_{work})	0.00	-0.04	-1.32	-3.79
	(0.28)	[−0.14]	[−4.72]	[−13.58]
Cov(α_{earn},β_{work})	-0.11	-0.13	-0.15	-0.11
	(0.03)	[−0.84]	[−1.49]	[−0.11]
Cov(β_{earn},α_{work})	-0.05	-0.06	-0.08	-0.04
	(0.02)	[−0.33]	[−1.37]	[0.51]
Cov(β_{earn},β_{work})	0.010	0.009	0.013	0.003
	(0.002)	[−0.28]	[1.25]	[−3.06]
Cov(α_{work},β_{work})	-3.21	-3.21	-3.21	0.17
	(0.02)	[−0.03]	[0.02]	[156.84]
Median Average Tax Rate	0.07	0.07	0.07	0.07
Median Marginal Tax Rate	0.15	0.15	0.15	0.15
Objective Value		4.69	43.88	26414

Table 9: This table reports estimated moments from the data along with their simulated counterparts. Cov($\alpha_{earn}, \alpha_{work}$) is set to zero rather than its estimated value of 2.04. In addition, Cov[$\alpha_{earn}, \beta_{earn}$] is set to -0.01 rather than its estimated value of -0.0039. All other moments are the same as in Table 4. Standard errors are reported in parentheses in column 1. Columns 2–4 report the difference between each simulated and estimated moment, divided by the standard error, in square brackets. Data are from the NLSY79. Sample weights were used.

work hours in the data. To address this second concern, we run a robustness check in Appendix section A wherein we include occupation dummies as controls in our auxiliary statistical model. Doing so does lower Cov[$\alpha_{earn}, \alpha_{work}$] somewhat from 2.04 to 1.78, but raises the variance of tastes slightly, from 0.215 to 0.221.

Third, the taste parameter in our model can have multiple interpretations. Throughout the paper, we interpret the taste parameter as reflecting agents' underlying preferences. But it could be interpreted as capturing the opportunity cost of time more broadly, the shadow value of household time. For example, rather than having different tastes for leisure, agents might have different technologies for home production. This alternative interpretation would alter the interpretation of our findings and would have implications for the social welfare function. It would also complicate the optimal tax problem significantly by introducing untaxed home production. Gayle and Shephard (2019) solve a model of optimal taxation with a marriage market, home production, and joint labor supply. The optimal tax policy depends crucially on a number of details, especially which observables the taxing authority is allowed to condition on. Gayle and Shephard (2019) find that accounting for home production flattens the tax schedule because the option to engage in (untaxed) home production effectively raises labor supply elasticities, thereby increasing the efficiency cost of taxing labor. However, like much of the optimal taxation literature, their model is static and therefore cannot allow for the dynamic channel that we highlight in this paper.

Fourth, our results could be explained by differences in beliefs. It is possible to generate the moments in our data by assuming that, rather than different tastes (or talent), agents have different beliefs about the returns to human capital investment (talent in our model) which lead them to make different human capital investment choices. However, we believe it would be hard for such a model to reliably generate the dynamic moments our model finds are important. If some households correctly believe human capital investment has high returns, they will tend to work more hours and earn less in the beginning of the lifecycle compared to other households who incorrectly believe human capital investment has low returns, as they spend more work time accumulating human capital. This tends to cause the covariance between hours and earnings to be negative, counter to what we see in the data. Moreover, such households will tend to decrease work hours later in life, causing the covariance between earnings and hours to decline. While this is one example, it makes clear the binding constraints our empirical estimates tend to put on the class of models that look to find no taste differences. Although an interesting direction for future research, we follow the literature by adopting a standard rational expectations framework where all agents know the human capital production function as well as future returns on human capital.

Finally, we estimate our model using only a select subsample—men who are strongly attached to the labor market—over a limited age range of 30–44. These data decisions were deliberate on our part. Although they each have downsides, will feel that the benefits outweigh the costs. By focusing on strongly attached, prime age males rather than a broader sample, we bias ourselves against finding taste differences. Including women or teenagers or college students or older workers in the sample would almost certainly inflate the variance of tastes for leisure. Moreover, including these other types of workers would make it much harder to justify our simple model of lifecycle labor supply.

6.4 Optimal tax rates

We now explore the consequences of our estimated taste differences for optimal taxation. In doing so, we again stress that our model is very simple and is probably not a reliable guide to the optimal *level* of taxes. Rather, our goal is to explore how optimal taxes *change* in response to a change in the variance of tastes.

The government chooses the tax parameters λ and τ to maximize utilitarian welfare with equal Pareto weights.[29] The population of taxpayers consists of 36 identical cohorts ranging from 30 to 65 years old, with each cohort consisting of the ten triples from our estimate $\widehat{\Gamma}$. Denoting the utility of individual i at age t as $U_{i,t}$ the government's problem simplifies to:

$$\max_{\lambda,\tau} \sum_{i=1}^{N} \sum_{t=1}^{T} U_{i,t}$$

subject to the constraint that tax revenue net of transfers equals revenue in our estimated model. The optimal tax will depend on the joint distribution of talent A, taste ϕ, and initial human capital $\overline{k}$. The first row of Table 10 reports the variance of log pre-tax earnings. The second and

[29] One might argue on philosophical grounds that redistribution on the basis of tastes alone is unjustified (Lockwood and Weinzierl, 2015; Fleurbaey and Maniquet, 2006). But when wages depend on both talent and tastes, it is not obvious how an income tax should unravel the two. Although, utility comparisons are always fraught, especially when tastes differ, we proceed anyway to illustrate additional complications that arise in the face of taste differences.

third rows report the optimal values for the two tax parameters λ and $1-\tau$. Recall that post-tax earnings are related to pre-tax earnings by the formula $earn_{post} = \lambda(earn_{pre})^{1-\tau}$, and $1-\tau$ is the elasticity of post-tax earnings with respect to pre-tax earnings. τ governs the progressivity of the tax system with $\tau = 0$ corresponding to a flat tax of $1-\lambda$. In matching the current U.S. tax code, we estimate values of $\lambda = 1.1288$ and $\tau = 0.0914$. For someone with \$100,000 of pre-tax earnings, this implies a marginal tax rate of 21.7 percent and an average tax rate of 13.8 percent. In contrast, the optimal income tax in our model has $\lambda = 3.246$ and $\tau = 0.577$, which, at \$100,000 of pre-tax earnings, implies a marginal tax rate of 75.1 percent and an average tax rate of 41.1 percent.[30] The fourth row reports the variance of post-tax earnings, which is substantially smaller than pre-tax earnings due to the strong progressivity of the optimal tax. Column 1 reports results under our baseline estimation. In column 2, we eliminate variation in talent by replacing all the values of A in our ten triples with their mean value. Doing so lowers the progressivity of the optimal tax—at \$100,000 of income, marginal and average tax rates fall to 65.2 and 36.4 percent—although some of this effect comes from a halving of the variance of pre-tax earnings. In column 3, we eliminate variation in taste. Despite the central role of talent in generating income inequality in the model, eliminating variation in tastes reduces the variance of log pre-tax earnings by three-quarters, which is *more* than eliminating variation in talent. Comparing pre-tax incomes in Column 1, 2 and 3 is similar in concept to the small perturbation method we report in Table 6, but with larger changes. Tax progressivity falls—at \$100,000 of income, marginal and average tax rates fall to 43.2 and 27.7 percent—although again, much of this decline is due to a dramatic reduction in the variance of pre-tax earnings.

To properly understand how taste differences affect optimal taxes, we need to adjust importance of tastes while holding the level of inequality fixed. We do this by increasing the variance of talent by enough to raise the variance of log pre-tax earnings by ten percent and simultaneously decreasing the variance of tastes to hold the variance of log pre-tax earnings unchanged. After "replacing" ten percent of the taste variation with talent variation, we recalculate the optimal

[30] The optimal progressivity in our model is very high, in large part due to the strong curvature in the utility of consumption, a property pointed out by Diamond and Saez (2011) who note, "Because the government values redistribution, the social marginal value of consumption for top-bracket tax filers is small relative to that of the average person in the economy, and [the value of a dollar to a wealthy person] is small and as a first approximation can be ignored. A utilitarian social welfare criterion with marginal utility of consumption declining to zero, the most commonly used specification in optimal tax models, has this implication." However, we stress that our model is too simple to take the optimal tax schedule seriously as a policy proposal. Rather, our goal is to understand how taste differences alter the optimal tax schedule.

Table 10: Optimal Tax Rates

	Baseline	Var(A)=0	Var(ψ)=0
Variance pre-tax earnings	0.724	0.359	0.164
Tax scale parameter λ	3.246	2.428	1.361
Tax progressivity parameter (τ)	0.577	0.453	0.214
Variance post-tax earnings	0.147	0.107	0.103

Table 10: This table reports the optimal tax parameters when $\alpha = 1$ along with the variance of pre-tax and post-tax earnings. The first column reports the baseline optimal tax. The second column recalculates the optimal tax rate with no talent differences, while the third recalculates the optimal tax with no taste differences.

tax. By replacing some of the variance of earnings due to taste with variation due to talent, we are answering the question "how much does the optimal tax rate vary when tastes play a smaller role in driving earnings inequality?" One final complication relates to the scaling of utility. Thus far, we have assumed that taste differences reflect the disutility of work. By rescaling person i's utility function by $1/\phi_i$, we could reinterpret taste differences as reflecting the marginal utility of consumption. We therefore introduce a monotonic transformation, $\zeta(\phi;\alpha) \equiv \alpha + (1-\alpha)\frac{1}{\phi_i}$ indexed by $\alpha \in [0,1]$ that allows us to control the degree to which taste differences load on consumption rather than leisure. We now rewrite utility as

$$U_i(c,\ell,\phi_i) = \zeta(\phi_i;\alpha)\frac{c_t^{1-\sigma}}{1-\sigma} - \zeta(\phi_i;\alpha)\phi_i\frac{(1-\ell_t)^{1+\eta}}{1+\eta}. \tag{19}$$

Note that $\zeta(\phi_i;0) = \frac{1}{\phi_i}$, which loads taste differences entirely on consumption, while $\zeta(\phi_i;1) = 1$ loads taste differences entirely on leisure. Although different values for α will not change an individual's behavior, α does turn out to matter for the optimal tax. In part, this is because the marginal utility of consumption declines more quickly than the marginal utility of leisure, and in part it is because money can be directly redistributed while time cannot.

Table 11 reports the optimal marginal and average tax rates at \$100,000 of income before and after replacing ten percent of the variation in earnings. Each row reports the optimal tax rates for different values of α. When $\alpha = 1$ and taste differences reflect the marginal utility of leisure, taxes are higher. But they fall considerably when $\alpha = 0$ and taste differences reflect the marginal utility of consumption. What's more, replacing 10 percent of the variation of log earnings due to taste with variation due to talent lowers taxes when $\alpha = 1$, but *raises* taxes when $\alpha = 0$ (see Figure 5). As mentioned above, different values for α amount to a rescaling of the utility function, and are observationally equivalent. Yet, different values for α can have significant effects on the optimal level of taxation and on whether a smaller role for tastes would raise or lower the optimal tax.

Table 11: Optimal Marginal and Average Tax Rates at $100,000 of Income

	Baseline		Counterfactual $Var(A)\uparrow$, $Var(\phi)\downarrow$	
α	Marginal Tax Rate	Average Tax Rate	Marginal Tax Rate	Average Tax Rate
0.00	56.3	27.9	57.2	29.9
0.20	59.3	29.7	59.8	31.5
0.40	62.7	31.8	62.7	33.3
0.60	66.3	34.2	65.9	35.5
0.80	70.4	37.2	69.3	38.0
1.00	75.1	41.1	73.3	41.1

Table 11: This table reports the optimal marginal and average tax rates at $100,000 of income for various levels of α. $\alpha = 0$ attributes all taste differences to the marginal utility of consumption while $\alpha = 1$ attributes all taste differences to the marginal utility of leisure. The second and third columns report the optimal marginal and average tax rates for someone earning $100,000 in our baseline model. The fourth and fifth columns report the optimal tax rates after increasing the variance of talent by enough to raise the variance of log pre-tax earnings by ten percent and simultaneously decreasing the variance of tastes to hold the variance of log pre-tax earnings unchanged.

The large taste differences we find raise normative concerns about redistributing income on the basis of tastes (Fleurbaey and Maniquet, 2006; Lockwood and Weinzierl, 2015). More disturbingly, we also find that the optimal tax depends on the scaling of utility. Since differently-scaled utility functions lead to observationally equivalent behavior, we cannot know how to apportion variation in taste differences between the marginal utility of consumption and the marginal utility of leisure. Therefore, even if we are able to estimate relative tastes, optimal tax policy still depends crucially on untestable assumptions about whether taste differences reflect differences in the marginal utility of consumption or of leisure. Though none of this is a concern if tastes for leisure do not vary across individuals, we find that they do.

7 Conclusion

Our paper makes three related points. First, tastes for leisure will affect human capital investment early in life, and therefore wages later in life. Thus, even if wages depend only on human capital, they will still reflect both talent *and* tastes. Second, motivated by a canonical model

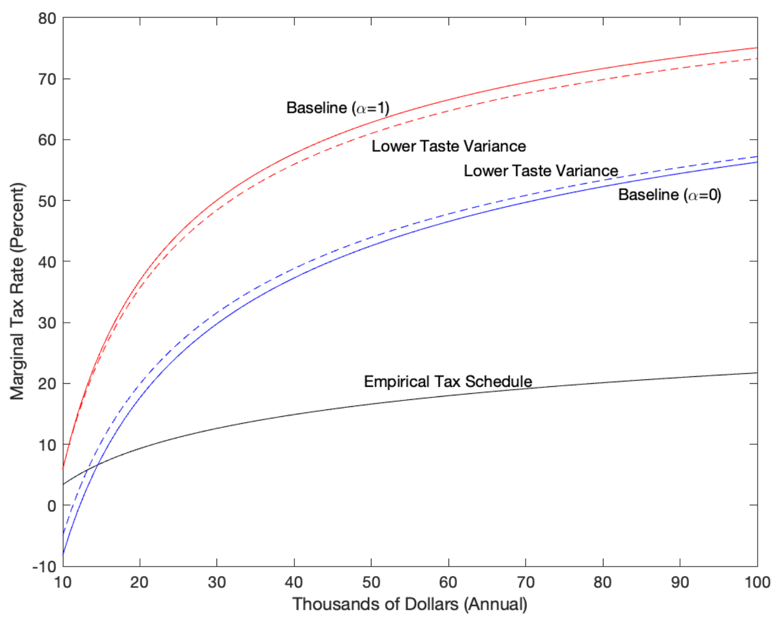

Figure 5: This figure plots several marginal tax rate schedules. The lowest line plots the tax schedule estimated from the data. Next come a pair of lines that plot the optimal schedule when $\alpha = 0$ and taste differences reflect the marginal utility of consumption. The solid line plots the optimal schedule in baseline while the dashed line plots the optimal schedule after increasing the variance of talent by enough to raise the variance of log pre-tax earnings by ten percent and simultaneously decreasing the variance of tastes to hold the variance of log pre-tax earnings unchanged. Finally, the highest pair of lines plot the optimal tax schedule when $\alpha = 1$ and taste differences reflect the marginal utility of leisure.

of labor supply, we find reduced form evidence that tastes for leisure vary. Third, we show that differentiating talent from tastes requires lifecycle data on both earnings and work hours, a simple model of lifecycle labor supply and human capital investment fits the data well and finds substantial variation in tastes, and models without taste differences struggle to match this data. Fourth, we find the optimal level of redistribution depends on the role of taste differences in driving earnings inequality and, to a disturbing degree, on the scaling of utility.

Our results are largely driven by two empirical moments. First, earnings and work hours at age 30 are strongly positively correlated. Second, while it is negative, the covariance between earnings at age 30 and the growth in earnings is not very large in magnitude, which implies relatively few crossings of earnings paths after age 30. When these two facts about lifecycle dynamics are viewed through the lens of our model, we find that tastes are an important driver of income inequality.

We restrict ourselves to a simple model and mitigate unmodeled factors by our choices of sample and age range. We acknowledge the possibility that alternative models could find a smaller role for tastes in determining earnings inequality. But we emphasize that any such model must be able to generate a positive correlation between hours and earnings at early ages along with relatively few crossings of earnings paths after age 30. One possibility might be to incorporate capital imperfections which force otherwise identical individuals into occupational paths with different Mincerian tradeoffs. However, these types of models would themselves have important implications for optimal tax policies, reinforcing our view that the correlation of earnings and work hours when young, along with the correlation of early earnings and earnings growth, are crucial moments to target for any model of optimal taxation.

Our results highlight a thorny issue in optimal taxation. Although structural optimal tax models without taste difference can fit the earnings distribution alone, they will miss badly on the distribution of work hours. However, allowing for taste differences raises other concerns. Should taxes redistribute on the basis of tastes? If not, how could or should an income tax separate talent from tastes when both are embedded in wages? Perhaps most disturbing of all, what are we to make of the sensitivity of optimal taxes to the scaling of utility?

References

Mark Aguiar and Erik Hurst. Measuring trends in leisure: The allocation of time over five decades. *The Quarterly Journal of Economics*, 122(3):969–1006, 2007.

Joshua D. Angrist and William N. Evans. Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size. *American Economic Review*, 88(3):450–477, 1998.

Susanto Basu and Miles S. Kimball. Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption. 2002.

Yoram Ben-Porath. The Production of Human Capital and the Life Cycle of Earnings. *Journal of Political Economy*, 75(4):352–365, 1967.

Katy Bergstrom and William Dodds. Using labor supply elasticities to learn about income inequality: The role of productivities versus preferences. *American Economic Journal: Economic Policy*, 13(3):28–62, August 2021.

Richard Blundell, Pierre-André Chiappori, and Costas Meghir. Collective Labor Supply with Children. *Journal of Political Economy*, 113(6):1277–1306, 2005.

Laurens Cherchye, Bram De Rock, and Frederic Vermeulen. Married with Children: A Collective Labor Supply Model with Detailed Time Use and Intrahousehold Expenditure Information. *American Economic Review*, 102(7):3377–3405, 2012.

Peter Diamond and Emmanuel Saez. The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives*, 25(4):165–190, 2011.

Marc Fleurbaey and François Maniquet. Fair income tax. *Review of Economic Studies*, 73(1): 55–83, 2006.

Marc Fleurbaey and François Maniquet. Optimal Taxation Theory and Principles of Fairness. *Journal of Economic Literature*, 56(3):1029–1079, 2018.

George-Levi Gayle and Andrew Shephard. Optimal Taxation, Marriage, Home Production, and Family Labor Supply. *Econometrica*, 87(1):291–326, 2019.

Paul Gomme and Peter Rupert. Theory, Measurement and Calibration of Macroeconomic Models. *Journal of Monetary Economics*, 54:460–497, 2007.

Nezih Guner, Remzi Kaygusuz, and Gustavo Ventura. Income taxation of U.S. households: Facts and parametric estimates. *Review of Economic Dynamics*, 17(4):559–581, 2014.

Fatih Guvenen and Burhanettin Kuruscu. Understanding the Evolution of the US Wage Distribution: A Theoretical Analysis. *Journal of the European Economic Association*, 10(3): 482–517, 2012.

Tomas Havranek, Roman Horvath, Zuzana Irsova, and Marek Rusnak. Cross-country heterogeneity in intertemporal substitution. *Journal of International Economics*, 96(1):100–118, 2015.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. Optimal tax progressivity. *The Quarterly Journal of Economics*, 132:1693–1754, 2017.

Lutz Hendricks. The Ben-Porath model and age-wage profiles. Mimeo. University of North Carolina., 2013.

Charles Hokayem and James P. Ziliak. Health, Human Capital, and Life Cycle Labor Supply. *American Economic Review*, 104(5):127–131, 2014.

Susan Imai and Michael P. Keane. Intertemporal Labor Supply and Human Capital Accumulation. *International Economic Review*, 45(2):601–642, 2004.

Marek Kapička. Optimal Mirrleesean Taxation in a Ben-Porath Economy. *American Economic Journal: Macroeconomics*, 7(2):219–248, 2015.

Greg Kaplan. Inequality and the Lifecycle. *Quantitative Economics*, 3:471–525, 2012.

J Kennan. A Note on Discrete Approximations of Continuous Distributions. *Mimeo*, (September), 2006.

Benjamin B. Lockwood and Matthew Weinzierl. De Gustibus non est Taxandum: Heterogeneity in preferences and optimal redistribution. *Journal of Public Economics*, 124:74–80, 2015.

Hamish Low, Costas Meghir, and Luigi Pistaferri. Wage Risk and Employment Risk over the Life Cycle. *American Economic Review*, 100(4):1432–1467, 2010.

Thomas E MaCurdy. An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy*, 89(6):1059–1085, 1981.

J. A. Mirrlees. An Exploration in the Theory of Optimal Income Taxation. *The Review of Economic Studies*, 38(2):175–208, 1971.

Derek Neal and Sherwin Rosen. Theories of the Distribution of Earnings. In A. B. Atkinson and F. Bourguignon, editors, *Handbook of Income Distribution, Volume 1*, number 7, pages 379–427. Elsevier Science, 2005.

Sherwin Rosen. Learning and Experience in the Labor Market. *The Journal of Human Resources*, 7(3):326–342, 1972.

Mark R Rosenzweig and Kenneth I. Wolpin. Life-Cycle Labor Supply and Fertility: Causal Inferences from Household Models. *Journal of Political Economy*, 88(2):328–348, 1980.

Stefanie Stantcheva. Learning and (or) Doing: Human Capital Investments and Optimal Taxation. NBER Working Paper 21381, National Bureau of Economic Research, 2015.

Appendices

Appendix A Occupations

In the model in the paper, we abstract away from occupations. In this appendix section, we consider two ways in which occupations could be incorporated into the analysis.

First, as demonstrated by Rosen (1972), a Ben-Porath model of lifecycle human capital investment is equivalent to a model of learning-by-doing with a large number of occupations each offering different learning opportunities. Occupations with more learning opportunities tend to attract younger workers. In equilibrium, these occupations pay lower wages whereas occupations with fewer learning opportunities pay higher wages and tend to be filled by older workers. Thus, the Ben-Porath framework can be interpreted as a model of endogenous occupation choice, and occupation switching, over the lifecycle. In this interpretation, occupation choice is the mechanism by which workers make human capital investments and no correction for occupations is needed.

Second, it may be that wages and hours come in occupation-specific bundles. For example, it may simply be impossible to be a highly paid heart surgeon or CEO while only working a few hours per week, because the nature of the job demands a major time commitment, which is then reflected in wages. As we show in the paper, our finding of large taste differences is driven in large part by the positive correlation between hours and earnings. If earnings and hours come in bundles, where high hours come paired with high earnings, then this could explain why we find such a large positive correlation of earnings and work hours.

To address this second concern, we run a robustness check wherein we include occupation dummies as controls in our auxiliary statistical model. If high hours and earnings come in occupation-specific bundles, then these occupation dummies would absorb much of the covariance in hours and earnings (and the variances of hours and earnings). The statistical model then becomes

$$\log(earn_{ityo}) = \alpha_{i,earn} + \beta_{i,earn}t + \zeta_y + \varphi_o + u_{it} \tag{20}$$

$$work_{ityo} = \alpha_{i,work} + \beta_{i,work}t + \xi_y + \psi_o + v_{it} \tag{21}$$

where φ_o and ψ_o are occupation fixed effects. Our findings are not sensitive to including or excluding these occupation controls. While including occupation dummies does reduce Var[α_{earn}], Var[α_{work}], and Cov[$\alpha_{earn}, \alpha_{work}$] somewhat, the variance of tastes remains essentially unchanged.

Table 12: Estimated Means, Variances, and Covariances

	Talent (A)	Taste (ϕ)	Initial Human Capital ($R\bar{k}$)
Panel A: Baseline			
Means	0.251	0.725	33.652
		Covariances	
Talent (A)	0.011	-0.007	0.995
Taste (ϕ)		0.215	-4.548
Initial Human Capital ($R\bar{k}$)			563.409
Panel B: Occupation Controls			
Means	0.240	0.722	32.053
		Covariances	
Talent (A)	0.010	-0.006	0.684
Taste (ϕ)		0.223	-3.502
Initial Human Capital ($R\bar{k}$)			376.940

Table 12: This table reports means, variances, and covariances for our ten estimated triples. Panel A replicates our baseline results reported in Table 3 in the paper. Panel B reports the results from fitting to moments obtained from the auxiliary statistical model in equations (20) and (21), which include occupation controls (see Table 13).

Table 13: Estimated and Simulated Moments With Occupation Controls

$$\log(earn_{ity}) = \alpha_{i,earn} + \beta_{i,earn}t + \zeta_y + \varphi_o + u_{it}$$
$$work_{ity} = \alpha_{i,work} + \beta_{i,work}t + \xi_y + \psi_o + v_{it}$$

	No Occupation Controls		Occupation Controls	
	Estimated	Simulated	Estimated	Simulated
Mean(α_{earn})	10.63	10.63	10.64	10.64
	(0.02)	[0.00]	(0.02)	[0.16]
Mean(β_{earn})	0.024	0.024	0.022	0.023
	(0.007)	[−0.03]	(0.01)	[0.14]
Mean(α_{work})	43.83	43.83	43.76	43.75
	(0.26)	[0.00]	(0.25)	[−0.07]
Mean(β_{work})	-0.017	-0.008	-0.039	0.006
	(0.093)	[0.10]	(0.09)	[0.53]
Var(α_{earn})	0.29	0.29	0.24	0.24
	(0.01)	[0.00]	(0.01)	[−0.12]
Var(β_{earn})	0.00119	0.00119	0.00110	0.00110
	(0.00007)	[0.04]	(0.00)	[0.03]
Var(α_{work})	70.23	70.24	59.71	59.70
	(3.62)	[0.00]	(3.16)	[−0.00]
Var(β_{work})	0.439	0.437	0.402	0.397
	(0.034)	[−0.04]	(0.03)	[−0.16]
Cov(α_{earn},β_{earn})	-0.0039	-0.0039	-0.0040	-0.0046
	(0.0013)	[0.03]	(0.00)	[−0.55]
Cov(α_{earn},α_{work})	2.04	2.04	1.78	1.81
	(0.28)	[−0.00]	(0.25)	[0.10]
Cov(α_{earn},β_{work})	-0.11	-0.11	-0.10	-0.11
	(0.03)	[−0.01]	(0.02)	[−0.36]
Cov(β_{earn},α_{work})	-0.05	-0.05	-0.05	-0.05
	(0.02)	[−0.01]	(0.02)	[0.31]
Cov(β_{earn},β_{work})	0.010	0.010	0.010	0.009
	(0.002)	[−0.01]	(0.00)	[−0.23]
Cov(α_{work},β_{work})	-3.21	-3.21	-2.96	-2.96
	(0.02)	[−0.01]	(0.02)	[−0.05]
Median Average Tax Rate	0.07	0.07	0.07	0.07
Median Marginal Tax Rate	0.15	0.15	0.15	0.15
Objective Value		0.02		0.97

Table 13: This table reports estimated moments from the data along with their simulated counterparts. Column 1 reproduces the same moments reported in Table 4 in the paper, while column 3 reports moments after including occupation controls. Standard errors are reported in parentheses in columns 1 and 3. Columns 2 and 4 report the difference between each simulated and estimated moment, divided by the standard error, in square brackets. Data are from the NLSY79. Sample weights were used.

Table 14: Moment Elasticities

Age	30-44	45-65	30-65
Panel A: Mean Income			
Talent: A	0.000	-0.032	-0.018
Taste: ϕ	0.017	0.026	0.022
Initial Human Capital: $\overline{k}$	-0.006	-0.004	-0.005
Panel B: Variance of Income			
Talent: A	0.25	1.50	1.26
Taste: ϕ	1.24	1.18	1.17
Initial Human Capital: $\overline{k}$	0.95	-0.68	-0.32
Panel C: Skewness of Income			
Talent: A	-0.32	1.42	3.24
Taste: ϕ	0.16	-2.01	-2.57
Initial Human Capital: $\overline{k}$	0.67	-0.17	-1.89

Table 14: Panels A, B, and C report the estimated elasticities from equations (16), (17), and (18) in the paper. Column 1 reports the elasticities when pooling together ages 30–44, column 2 reports them for ages 45–65, and column 3 reports them for all ages 30–65.

Appendix B Moment Robustness

In this appendix section, we explore alternative approaches to calculating the empirical moments using the NLSY79. Specifically, we try all permutations of the following:

1. Restrict the sample to observations with a minimum number of hours equal to either 1 or 100.

2. Censor work hours from below at 0 (no censoring) or 200.

3. Censor work hours from above at 4000 hours or no censoring.

Together, these three binary choices yields eight possible sample restrictions. The main estimates in the paper correspond to (i) minimum work hours equal to 1, (ii) no truncation from below, and (iii) no truncation from above (see the first row of Table 15).

Table 15 below summarizes the effect of each of these eight possible sample restrictions on the standard deviation of work hours and the correlation between work hours and log earnings. Rather than plot the moments, we simply report the average standard deviation across ages 30–44, and the correlation of log earnings and work hours at age 37 along with its first two derivatives at age 37. In the final column of table 15, we report the percentage of the standard deviation in log earnings at age 44 that our model attributes to tastes. As table 15 illustrates, our results are relatively consistent across all 16 data treatments; for instance, the percentage of variation in earnings that is attributed to tastes fluctuates between 0.559 and 0.588.

Table 15: Robustness to Different Treatments of the Data

Min. hours for inclusion	Left cens. of work hours	Right cens. of work hours	No occupation controls			Occupation Controls		
			Var(A)	Var(ϕ)	Var(k)	Var(A)	Var(ϕ)	Var(k)
1	None	None	0.011	0.215	1.409	0.010	0.221	0.979
1	None	4000	0.011	0.278	1.268	0.011	0.299	1.037
1	200	None	0.012	0.313	1.281	0.010	0.267	1.009
1	200	4000	0.012	0.331	1.472	0.009	0.220	1.087
100	None	None	0.012	0.362	1.250	0.011	0.274	1.035
100	None	4000	0.011	0.266	1.185	0.010	0.256	1.029
100	200	None	0.012	0.315	1.302	0.011	0.277	1.031
100	200	1000	0.011	0.240	1.443	0.009	0.224	0.959

Table 15: This table reports the variances of talent (A), tastes (ϕ), and initial human capital ($R\bar{k}$) for eight different treatments of the data. The first column, labeled "Min. hours for inclusion" indicates whether workers were required to work at least 1 hour per year or 100 hours per year to be in the sample. The second column, labeled "Left cens. of work hours" indicates whether annual work hours were censored from below at 200 hours or not. The column labeled "Right cens. of work hours" indicates whether annual work hours were censored from above at 4000 hours or not. The remaining columns report the variances of talent, tastes, and initial human capital, both with and without occupation controls.