Heterogeneity in Talent or in Tastes? Implications for Redistributive Taxation

Ian Fillmore
Trevor Gallen

Preliminary and Incomplete: Do not cite

November 2018

Abstract

Do differences in tastes for leisure play an important role in determining income inequality? Although in a static model tastes for leisure only affect work hours, in a lifecycle human capital model tastes can also affect wages via their effect on human capital investment. After focusing on highly-attached prime-age males in the NLSY79 and filtering out idiosyncratic shocks to work hours and earnings, we establish two important facts: college graduates work more hours than non-graduates, even conditional on wages, and the standard deviation of work hours is large and does not decline with age. We argue that both of these facts point to a lifecycle human capital model in which tastes play an important role in determining income variation. We fit a simple model in which workers are heterogeneous in (i) their ability to accumulate human capital (talent) (ii) their preferences over consumption vs. leisure (taste) and (iii) their initial human capital. We find that tastes play a large role: 68% of income variation at age 44 is due to tastes, rather than talent or initial human capital. These findings are driven by the high standard deviation in "permanent" labor hours and a large positive correlation between labor hours and earnings. Finally, we show that exchanging the sources of income variation between talent and tastes changes redistributive tax rates significantly, particularly when heterogeneity is due to differences in the marginal utility of consumption, rather than leisure.

1 Introduction

This paper observes two important stylized facts about labor hours and earnings among prime-age males. First, there is a large standard deviation of labor hours that is relatively constant throughout ages 30-44. Second, among the same group, the correlation of hours and earnings is also positive, relatively high, and not declining in age. We find that fitting these dynamic, lifecycle facts to a simple lifecycle model of labor supply drives a finding that heterogeneous preferences are important determinants of income inequality. We find that accounting for this heterogeneity in a simple optimal tax model has the capacity to dramatically change optimal tax rates.

Casual empiricism suggests that at least some of the differences in people's life outcomes are driven by differences in their life goals and preferences. Consistent with this observation, Kahneman (2011, p. 401) reports results from a study of students at elite colleges in 1976. At ages 17 or 18, the researchers asked students about the value they placed on becoming well off financially. The researchers followed these students over time and found that, among those who had placed a high value on financial well being, adult earnings was an important predictor of happiness. But for those who had placed a low value on financial well being, adult earnings was far less predictive of happiness. These results suggest the possibility that one's preferences may be important drivers of one's earnings. People who care a lot about money and the things it can buy will be more inclined to make choices that lead to high incomes. Those who care more about non-monetary goods will tend to pursue different paths.

Mirrlees (1971) spearheaded the study of optimal tax design in the presence of unobservable heterogeneity in worker skills. Following Mirrlees, this literature has traditionally assumed that individuals possess the same preferences and that their skills (and wages) are exogenously fixed rather than an endogenous product of human capital investment. In this paper, we provide evidence that tastes for leisure differ across individuals and that they can and do have a substantial effect, not just on earnings, but on wages themselves by affecting the incentives to invest in human capital earlier in life. In other words, tastes for leisure do not simply affect labor supply holding the wage constant, as in a static labor supply model. They also influence earlier human capital investments which then directly affect wages later in life. Moreover, this dynamic channel has much larger effects on income differences than the static channel, with important consequences for redistributive taxation.

This paper uses a simple lifecycle model to identify the relative importance of tastes versus talent. We find that tastes account for a large fraction of income inequality at age 44, and we identify which empirical moments drive our estimates. Finally, we show that accounting for these taste differences lowers the optimal redistributive income tax. We are not the first to consider the tax consequences of human capital investment (Stantcheva, 2015), nor are we the first to worry about taste heterogeneity (Lockwood and Weinzierl, 2015; Bergstrom and

Dodds, 2018). We follow Neal and Rosen (2005) in pointing out that combining human capital investment over the lifecycle with taste heterogeneity introduces a dynamic channel by which relatively small taste differences can be magnified into large wage and income differences later in life. Our primary contributions are to (1) advance the literature on static identification of preference heterogeneity by pointing to new dynamic (lifecycle) moments that drive our finding that tastes are important (2) to fit our dynamic model to panel data on work hours and earnings and explore the consequences of taste heterogeneity for redistributive taxation.

To understand the dynamic channel by which tastes for leisure can affect wages and earnings, consider the classic lifecycle model of Ben-Porath (1967). In this model, individuals differ in their ability to acquire new human capital, or talent. More talented individuals find it easier to acquire new skills and optimally choose to spend more time investing in their own human capital. Their initial earnings are low both because their initial human capital is low and because they devote more of their time to investment rather than work (either formal schooling or on-the-job training). As they accumulate human capital over time, their wages rise and they substitute away from human capital investment and toward work. Their low initial earnings are more than compensated for by high earnings later in life. In this manner, highly talented children become highly productive, and thus highly compensated, adults. Even relatively small differences in talent can compound into large differences in earnings later in life.

Neal and Rosen (2005) demonstrate a similar phenomenon regarding tastes for leisure by extending the model of Ben-Porath (1967) to include a labor/leisure choice. They emphasize that workers with low preference for leisure not only lose little utility by sacrificing leisure time to accumulate human capital, but, because they tend to work more in later life, amortize that cost over more labor hours. Consequently, ceteris paribus, low value of leisure when young will be correlated with higher wages when old. Taste will thereby doubly drive inequality when old, as it will contribute to both inequality in wages and inequality in hours, with higher-wage individuals working more. What Neal and Rosen point out is that greater human capital investment is driven by both an individual's talent *and* her taste for leisure. And just like talent, small differences in a person's taste for leisure can compound into large differences in human capital and wages later in life.

In Figure 1, we present some simple reduced form evidence that tastes for leisure are correlated with human capital investment. What does a classical model of labor supply, with homogenous preferences, predict about the work hours of college graduates relative to non-graduates? Because there are no differences in taste, 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 so only the income effect remains, we should expect college graduates to work fewer hours than non-graduates who have the same wage at age t, particularly at older ages when human capital investment is minimal.

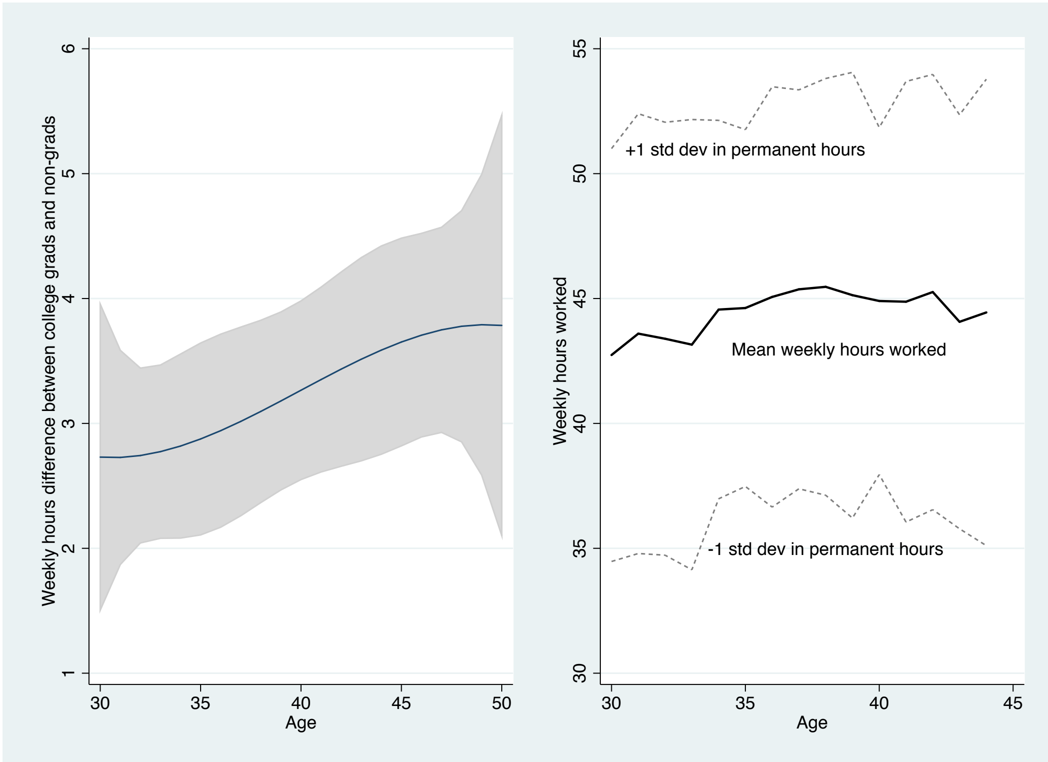

Figure 1: The left graph plots the estimated difference in weekly hours worked between college graduates and non-graduates. The gray bands indicate 95% confidence intervals. The graph is based on a regression of weekly hours at age t on a dummy for being a college graduate interacted with a cubic polynomial in age as well as a control for log wage at age t. The right graph plots the mean weekly work hours for ages 30 through 44. The dotted lines indicate one standard deviation in "permanent" hours where the variation due to transitory shocks has been removed (see section 5 for details).

Figure 1 shows the opposite is true. Using data on highly-attached, prime-age males from the National Longitudinal Survey of Youth 1979 (NLSY79), 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.[1] We find that college graduates actually work more hours than non-graduates with the same wage, and that this difference rises over the lifecycle. Although this evidence is inconsistent with a labor supply model with identical preferences, 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).

The right graph of Figure 1 presents additional evidence that is difficult to reconcile with a model of homogeneous preferences. In a simple lifecycle model with human capital accumulation and homogenous preferences, early-life work hours may differ for a variety of reasons, such as

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

differences in wages, differences in the marginal utility of lifetime income, and differences in human capital investment. Later in life, when human capital investment is minimal, work hours will only differ across individuals due to differences in wages and the marginal utility of lifetime income. However, an extensive literature in labor economics has established that substitution and income effects in labor supply mostly cancel out, especially for highly-attached prime-age men.[2] Thus, as workers age we would expect the variance of work hours to shrink.

Returning to our sample of highly-attached, prime-age males from the NLSY79, the right graph of Figure 1 plots mean weekly hours worked from ages 30 to 44. The dashed lines plot one standard deviation bands in weekly hours across individuals. In calculating these bands, we filter out "transitory" variation in hours and focus on the variance of "permanent" hours worked (see section 5 for details). Even in our sample of highly-attached, prime-age males, "permanent" work hours vary enormously across individuals, and this variance does not decline with age, both inconsistent with a model in which income and substitution effects largely (or completely) offset.

Finally, Table 1 provides a third piece of evidence that is consistent with heterogeneity in tastes for leisure. In most models of labor supply with separable labor and consumption preferences, the first order condition for leisure implies that in the absence of taste heterogeneity, work hours depend on (1) the current wage and (2) the marginal utility of lifetime income. Therefore, after controlling for wages and lifetime income, work hours today should not predict work hours in the future: there is no third "omitted variable" that joins past and present work hours. On the other hand, if work hours also depend on heterogeneous tastes then past work hours will be positively correlated with future work hours: the omitted variable is tastes that are persistent across time. In Table 1, we report estimates of the following regression

$$\log(h_{i,44}) = \beta \log(h_{i,34}) + \delta_1 \log(w_{i,44}) + \delta_2 \log(LifeInc_i) + u_{i,44}$$

where $\log(h_{i,34})$ and $\log(h_{i,44})$ are annual hours worked during ages 30–34 and 40–44, $\log(w_{i,44})$ is the (log) wage during ages 40–44, and $\log(LifeInc_i)$ is an estimate of (log) lifetime income.[3] In Table 1, β, the coefficient of interest, is positive and highly significant.[4] The estimates imply that, if you compared two workers with the same wage at age 44 and the same lifetime income

[2]The observation that long-run wages have not affected long-run hours in post-war U.S. data is a persistent one in macroeconomics, and is (re)documented in **?**. More recently, some attention has been paid to the idea that income effects may be slightly stronger than substitution effects (see, for instance, **?**) which would strengthen the puzzle shown in Figure 1.

[3]Lifetime income was calculated by taking the present value of all labor income earned between ages 18–50 using a discount rate of 5 percent. Wages at 40–44 were calculated by dividing total labor income reported between ages 40–44 by total hours worked between ages 40–44.

[4]The estimated coefficient on the (log) wage is negative 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.

but the first worker worked 10 percent more hours at age 34, then that worker will work 2.2 percent more hours at age 44. Although at odds with a model of homogeneous preferences, this result fits naturally within a model that includes differences in tastes for leisure.

Table 1: Regressing Work Hours at 40–44 on Work Hours at 30–34

	Log work hours at 40–44
Log work hours at 30–34	0.22***
	(0.036)
Log wage at 40–44	-0.25***
	(0.028)
Log lifetime income	0.35***
	(0.036)
Observations	1,449

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

Table 1: This table reports estimates from regressing (log) work hours at ages 40–44 on (log) work hours at ages 30–34, controlling for wages at ages 40–44 and lifetime income. Standard errors are reported in parentheses. Sample weights were used.

In light of these facts, we fit a stylized version of the model of Neal and Rosen (2005) to data from the NLSY79. The model allows for differences in the ability to accumulate human capital, which we call "talent," and preferences for leisure rather than consumption, which we call "taste." Allowing for idiosyncratic variation in the taste parameter distinguishes us from a variety of other papers examining the distribution of earnings.[5] Using data from the NLSY79 on the joint distribution of "permanent" earnings and hours by age, we estimate the distribution of taste and talent for strongly-attached, prime-age males. In our estimated model, taste plays a large role in explaining the variance of earnings even in advanced stages of work-life. For instance, at age 44, we find that taste alone explains 68% of earnings variation, even for strongly-attached male workers.

Based on our fitted model, we compute the welfare maximizing flat income tax policy and compare it with the optimal policy if tastes for leisure did not vary across individuals. In our baseline calibration, we find that in a simple flat tax, constant transfer regime, a reduction in taste variation that leads to a 1% reduction in earnings variation at 44, combined with an increase in talent variation that leads to a 1% increase in earnings variation (i.e. changing the causes of income variation while holding income variation constant) increases the optimal tax rate by 0.63% (0.3 percentage points), suggesting that the optimal tax rate is highly sensitive to the sources of income variation.[6] Finally, we find that the magnitude of this result depends on

[5] See, for instance, Blandin (2018), Guvenen et al. (2014), Huggett et al. (2011).

[6] Of the three "sufficient statistics" of Saez (2001), changes in the source of income variation primarily changes the redistributive tastes of government, by changing how much different households value consumption, rather than earnings elasticities or the shape of the income distribution, which are largely held constant.

where heterogeneity lies: if preference heterogeneity lies in the marginal utility of consumption, optimal taxes are strongly responsive to the source of income variation, while if it lies in the marginal utility of leisure, optimal taxes are less responsive.

2 Lifecycle Model of Human Capital Investment

In this section, we present a simple model of human capital investment. We do so in three parts. First, in section 3.1 we follow Neal and Rosen (2005) and extend the classic Ben-Porath model of human capital investment to include a labor/leisure tradeoff. In section 3.2 we discuss the strengths and limitations of the model. Finally, in section 3.3 we discuss identification of the model and the moments in the data that drive our results.

2.1 Review of Neal and Rosen (2005) Model

The Ben-Porath model of human capital investment is perhaps the canonical model for understanding earnings of agents 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 wages but increases future wages. Neal and Rosen (2005) extend the Ben-Porath model of human capital investment to include a leisure choice. Agents are characterized by the triple $(A, \phi, \overline{k})$. A captures the agent's ability to acquire new human capital or "talent," ϕ captures her relative taste for leisure, and $\overline{k}$ is her initial level of human capital. In each period, the agent allocates her 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 her human capital. Note that talent does not directly raise wages. Rather, talented individuals find it easier to increase their 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 show, one can incorporate a leisure decision into the Ben-Porath model. Period utility depends on consumption c_t, leisure ℓ_t, and a parameter ϕ determining the agent's relative taste for consumption vs leisure. We will assume that the marginal utility of leisure goes to infinity as leisure goes to zero, so that the agent always consumes a positive amount of leisure. The agent maximizes lifetime utility subject to a period time constraint and a lifetime budget constraint. In particular, the agent solves

$$
\begin{aligned}
\max_{\{c_t s_t \ell_t n_t\}} \quad & \sum_{t=1}^{T} \beta^{t-1} U(c_t, \ell_t; \phi) && (2) \\
s.t. \quad & s_t + \ell_t + n_t = 1 && (3) \\
& s_t,\ \ell_t,\ n_t \geq 0 && (4) \\
& \sum_{t=1}^{T} \left(\frac{1}{1+r}\right)^{t-1} c_t = \sum_{t=1}^{T} \left(\frac{1}{1+r}\right)^{t-1} w_t n_t && (5) \\
& w_t = R k_t && (6) \\
& k_{t+1} = (1-\delta) k_t + A(s_t k_t)^{\gamma} && (7) \\
& k_1 = \overline{k} && (8)
\end{aligned}
$$

The Kuhn-Tucker conditions for this problem are

$$
\begin{aligned}
[c_t]: \quad & \beta^{t-1} U_c(c_t, \ell_t; \phi) = \left(\frac{1}{1+r}\right)^{t-1} \lambda && (9) \\
[\ell_t]: \quad & U_\ell(c_t, \ell_t; \phi) = \mu_t && (10) \\
[n_t]: \quad & \left(\frac{1}{1+r}\right)^{t-1} \lambda w_t = \beta^{t-1}(\mu_t + \mu_t^n) && (11) \\
[n_t]: \quad & \mu_t^n n_t = 0;\ \mu_t^n \geq 0;\ n_t \geq 0 && (12) \\
[s_t]: \quad & A\gamma (s_t k_t)^{\gamma-1} k_t \lambda \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_\tau = \mu_t && (13)
\end{aligned}
$$

where λ is the Lagrange multiplier on the lifetime budget constraint, μ_t is the Lagrange multiplier on the period time constraint, and μ_t^n is the multiplier on the non-negativity constraint for labor. Note that leisure ℓ_t and human capital investment s_t will always be positive because the marginal utility of leisure and the marginal product of human capital investment are both infinite at zero. In the case where $n_t > 0$, so that we are at an interior solution ($\mu_t^n = 0$), the first order conditions for ℓ_t and s_t become

$$
\begin{aligned}
[\ell_t]: \quad & \beta^{t-1} U_\ell(c_t, \ell_t; \phi) = \left(\frac{1}{1+r}\right)^{t-1} \lambda w_t && (14) \\
[s_t]: \quad & \left(\frac{1}{\beta(1+r)}\right)^{t-1} \gamma \frac{R(s_t k_t)^{1-\gamma}}{A\gamma} = \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_\tau. && (15)
\end{aligned}
$$

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. The agent will also 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 (15) 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 (15) makes it clear that the agent's human capital investment today depends on his future labor supply. If the agent anticipates working a lot in the future, then he has a stronger incentive to invest in human capital today. Just as in the standard Ben-Porath model, agents experience concave lifecycle earnings paths. More talented agents invest more heavily in human capital causing them to experience lower earnings when young, followed by steeper earnings paths which lead to higher earnings when old.

2.2 Strengths and Limitations of the Model

The Ben-Porath model is a canonical model of lifecycle human capital investment and earnings. Following Neal and Rosen (2005), we extend the model to include a labor/leisure choice. Much more complex models of lifecycle labor supply and human capital investment have been developed and estimated in the labor literature. Unlike more sophisticated models in this 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) involuntary unemployment and search frictions (Low et al., 2010), and credit constraints (Rossi and Trucchi, 2016) while the cited papers do not.

Rather than attempt to explicitly model these factors, we have attempted to mitigate their impact through our choice of subject, age span, and data treatment. We abstract from transitory health shocks, involuntary unemployment, and search frictions by extracting the "permanent" level of hours and earnings from the NLSY data, in a manner reminiscent of how researchers interested in secular trends filter out high-frequency noise. Although credit constraints could be important, we note that in the context of our model, credit constraints drive hours to go the wrong direction over the course of the lifespan. Specifically, credit constrained individuals aren't able to effectively finance consumption when young through debt, and as a consequence both work and study more, leading to a declining trend in hours as a function of age, when the credit constraint binds. In the data, however, hours rise slightly as a function of age. Finally, we have attempted to minimize the importance of fertility shocks by focusing on prime age males, whose labor supply is less affected by fertility shocks (Angrist and Evans, 1998).

Our goal is to fit a model that is both canonical and transparent, while still allowing both talent and taste to determine earnings. Moreover, our model's transparency allows us to identify which empirical moments are key to determining the extent to which earnings are affected by tastes vs talent.

2.3 Identification

In this section we illustrate how, within our model, the lifecycle paths of earnings and work hours (paid work time plus human capital investment) allow us to disentangle tastes and talent. In Figure 2 we plot the paths of earnings and work hours for an agent with triple $\{A, \phi, k_0\}$. We also plot paths for two alternative triples. In the first, we lower the disutility for labor ϕ; in the second, we choose talent A and initial human capital k_0 to match the earnings path of the low-disutility triple. 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 if we also observe lifecycle work hours, we can distinguish these two explanations. If the second agent is earning more over the lifecycle because he has a lower disutility of labor, then we should see him work more hours. On the other hand, if the second agent is earning more because he is more talented, then we should see him work less than than the first agent when young (due to an income effect).[7] Thus, despite producing the same lifecycle earnings paths, the low-disutility and high-talent triples display different paths for work hours. Lowering the distaste for labor shifts work hours up. In contrast, increasing talent actually lowers labor hours early in life due to an income effect.

3 Estimation

3.1 Data

We use data from the NLSY79 and restrict our sample to strongly attached males between the ages of 30 and 44 years old. We define strongly attached males to be those with complete or nearly complete data over the period who never report working zero hours in a year. For each respondent, we observe total labor income (across all jobs) as well as total number of hours worked (across all jobs) in the previous year.[8]

3.2 Model Specification for Estimation

In this section we adapt the model of Neal and Rosen (2005) for empirical estimation by specifying a functional form for the agent's period utility function, using the common preferences of MaCurdy (1981):

$$U_i(c, \ell; \phi_i) = \frac{c_t^{1-\sigma}}{1-\sigma} - \phi_i \frac{(1-\ell_t)^{1+\eta}}{1+\eta} \tag{16}$$

[7]Although it is not shown in Figure 2, high-talent work hours may exceed baseline work hours later in life as wages continue to rise.

[8]Reported work hours include both paid work time as well as human capital investment.

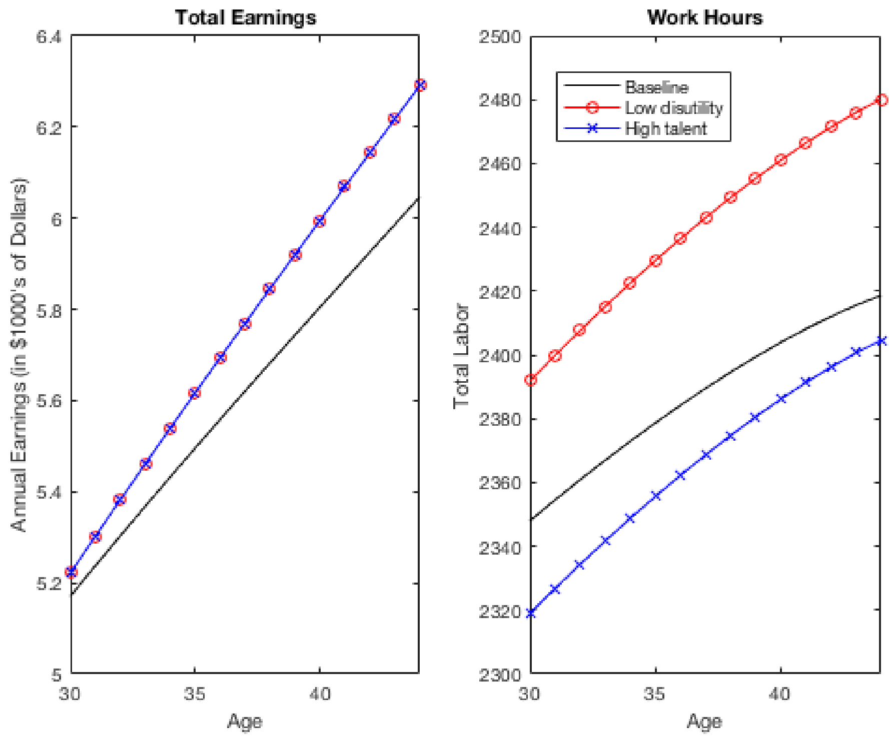

Figure 2: This figure solves the model for three separate triples (A, k_0, ϕ): a baseline triple (black lines), a low-disutility-for-labor triple (red lines) wherein ϕ is lower than in baseline, and a high-talent triple (blue lines). The high-talent triple was set by choosing talent A and initial human capital k_0 to match the earnings path of the low-disutility triple. Importantly, although the low-disutility and high-talent triples are observationally equivalent when looking at earnings alone, they can be distinguished by including data on work hours.

With this specification, the agent i's first order conditions become

$$[c_{it}]: \quad \beta^{t-1} c_{it}^{-\sigma} = \left(\frac{1}{1+r}\right)^{t-1} \lambda_i \tag{17}$$

$$[\ell_{it}]: \quad \beta^{t-1} \phi_i (1-\ell_{it})^{\eta} = \lambda_i \left(\frac{1}{1+r}\right)^{t-1} w_{it} \tag{18}$$

$$[s_{it}]: \quad \left(\frac{1}{\beta(1+r)}\right)^{t-1} \frac{R(s_{it}k_{it})^{1-\gamma}}{A_i \gamma} = \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_{i\tau}. \tag{19}$$

3.3 Estimating the Frisch Elasticity of Labor Supply

We estimate the Frisch elasticity of labor supply for men in our sample as follows. First, we follow Heckman et al. (1998) and restrict the sample to ages 48 to 55 and assume that $s_t \approx 0$ at these ages. Then, equation (18) becomes

$$
\begin{aligned}
n_{it}^{\eta} &= \left(\frac{1}{\beta(1+r)}\right)^{t-1} \frac{\lambda_i}{\phi_i} w_{it} \\
n_{it}^{\eta+1} &= \left(\frac{1}{\beta(1+r)}\right)^{t-1} \frac{\lambda_i}{\phi_i} y_{it}
\end{aligned}
$$

where y_{it} is annual earnings.[9] Writing this in logs gives us a simple fixed effects specification

$$\log n_{it} = \delta t + \alpha_i + \frac{1}{\eta+1} \log y_{it}$$

where δt is a common time trend and the α_i are fixed effects. Running this regression gives us an estimate for $\eta = 3.05$ with a standard error of 0.36, implying a Frisch elasticity of 0.33.[10]

3.4 Calibration

Although we estimate the "taste" disutility of labor parameter ϕ_{it}, the "talent" Ben-Porath efficiency parameter A_i, and initial human capital $k_{i,0}$ to match moments on leisure and earnings, we calibrate several other parameters directly. As described above, we set the Frisch elasticity of labor supply to be 0.33. In our baseline calibration, we choose $\gamma = 0.62$ and $\delta = 0.057$, consistent with Hendricks (2013).[11]

The elasticity of intertemporal substitution is relatively unimportant in our calibration, because conditional on σ differences in ψ regulate the static tradeoff between consumption and leisure, there are no shocks, and there are constant interest rates in our model.[12] We set σ so that the elasticity of inter temporal substitution is 0.5, consistent with both long-run labor supply (Basu and Kimball (2002) and micro-studies on the parameter (Havranek, Horvath, Irsova, and Rusnak 2013). The discount rate β is chosen to be 0.945, 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 1200. Thus, individual i's potential earnings in year t is given by $Rk_{it} = 1200k_{it}$.[13]

[9]Expressing the first order condition in terms of annual earnings rather than the wage avoids introducing so-called "division bias."

[10]In Appendix A, we calibrate η along with the joint distribution of (A, ϕ, k_0) to match our moments.

[11]While some literature suggests a Ben-Porath technology that's linear with depreciation near zero 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.

[12]In our model, the elasticity of intertemporal substitution will, ceteris paribus, help determine optimal tax rates by controlling utility function curvature.

[13]For convenience when interpreting results, we convert fraction of annual hours worked into annual hours

3.5 Estimation With Aggregate Moments

3.5.1 Main Calibration

We estimate the population of A_i, ϕ_i, and $\overline{k}_i$ to match the model to aggregate moments. We want to pin down the joint distribution of talent, taste, and initial human capital, choosing a population of triples so that their simulated aggregate moments match aggregate measures from the NLSY joint distribution of labor earnings by age and total labor by age. Specifically, we choose a population of individuals so that, given the solution to their individual problems, the simulated population matches our NLSY data on: 1) mean hours worked from 30-44 (inclusive of both human capital accumulation and labor hours) 2) standard deviation of hours worked from 30-44, 3) mean earnings path 4) standard deviation of log earnings path 5) 90/10 ratio of earnings from age 30-44, and 6) correlation of earnings and hours worked from 30-44.

3.5.2 Aggregate Moments

Our data consist of a panel of individuals from ages 30 to 44 with annual measures of total labor earnings and total work hours. Our first group of aggregate moments are simply the mean number of work hours and the mean of (log) annual income. We calculate these moments separately at each age. Our second group of aggregate moments measure the dispersion in work hours and annual income. We estimate the standard deviations of work hours and of (log) annual income as well as the 90/10 ratio of annual income. Since both annual income and work hours are subject to transitory shocks (as well as measurement error), the raw variances will be inflated. To deal with this, for each individual we regress both work hours and (log) earnings on age and store the fitted values and residuals. We use the variance of the residuals, at a given age, to estimate the variance of the transitory shocks. We then subtract the estimated variance of the transitory shocks from the total variance of observed work hours and annual income. The goal is to isolate the variance in hours and income that is *not* due to year-to-year transitory shocks. When we calculate the 90/10 ratio of annual income, we actual calculate the ratio of the 90th percentile of fitted annual income (from the individual specific regressions) to the 10th percentile of fitted annual income. Our third group of moments measures the correlation between work hours and annual income. We calculate the correlation between work hours and annual income, again adjusting for the additional covariance introduced by the transitory shocks to annual income and work hours. All moments were calculated using the NLSY79 sampling

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 hourly wage Rk_{it}, rather than raw human capital.

weights. In the end, we are left with 90 moments to target. We plot these moments in Figure 3.

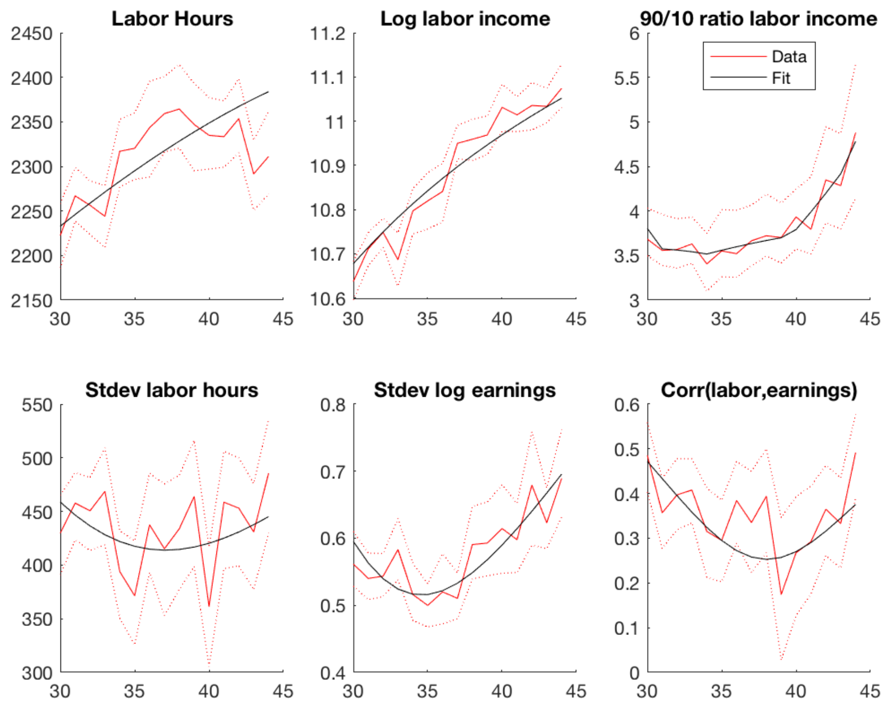

Figure 3: This figure depicts the six sets of empirical moments from the NLSY79 (red line) and model fits (black line) that describe the joint distribution of labor hours and earnings by age. Bootstrapped standard errors for data are depicted as red dotted lines.

Because our results depend on the moments in Figure 3, we also consider a series of "robustness" exercises, changing our methodology for selecting highly-attached agents and measuring their permanent choices for hours. We discuss these exercises, which largely do not change our results, in Appendix C.

3.5.3 Fits

Given the calibrated parameter values in Table 2, we find the population of talent, taste, and human capital that best fits our 90 empirical moments. Following Kennan (2006), we choose a population of 10 representative agents to fit the six moments of labor hours, log labor income, the 90/10 ratio of labor income, the standard deviation of log earnings, and the correlation between labor and earnings in each of the 15 years used in our model (ages 30-44). The moments and fits are depicted in Figure 3. While most of our simulated moments closely match their empirical counterparts, one failure stands out: the slope of our yearly labor hours "overshoots" the data, rising by nearly 151 hours per year where the data only rises by 89 hours per year. This is primarily being driven by our calibrated Frisch elasticity. Given a responsiveness of labor to changes in wages, our model could not generate a better fit for mean labor hours without lower wages, which would cause us to miss on the mean path of earnings. In Appendix A, we consider fitting the Frisch elasticity along with the distribution of talent, taste, and initial human capital.

Table 2: Directly-Calibrated Parameters

Parameter	Symbol	Value	Notes
Frisch elasticity of labor supply	η	-3.05	Frisch labor supply =0.33 (Estimated from NLSY)
Ben-Porath diminishing returns	γ	0.62	(Hendricks, 2013)
Human Capital depreciation	δ	0.057	(Hendricks, 2013)
Elas. of intertemporal subst.	σ	2	EIS=0.5 (Basu and Kimball, 2003)
Discount rate	β	0.945	(Gomme and Rupert, 2007)
Human capital wage	R	1200	Normalization

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

By matching the joint distribution of labor and earnings by age, we find ten triples of talent, taste, and initial human capital that minimize the sum of squared errors of the six sets of moments in the NLSY. We depict these values in Figure 4. We label the X-axis with talent A_i, the Y-axis with taste ψ_i, and display the initial hourly wage beside each of the ten points. Figure 4 indicates the presence of both low-talent, high-distaste for labor workers in the top right and high-talent, low-distaste workers in the bottom right of the figure. This type of heterogeneity suggests that a reasonable fraction of the population has lower earnings because of a taste for leisure combined with moderate talent, lowering the benefits of redistributing income toward them. We also find households with similar taste and talent but dramatically different initial human capital: this may represent differing initial conditions or financial frictions in the ages before thirty that cause agents that might otherwise have similar human capital to vary greatly. Finally, six of our agents are quite similar in their talent, taste, and human capital, denoting a core group that helps match the low standard deviation of earnings even as our extreme individuals help capture the larger 90/10 ratio.

4 Results

In order to decompose how much income variation is due to talent, taste, and initial human capital, we compare the effect of mean-preserving reductions in the standard deviation of each parameter of interest on the standard deviation of earnings. Denoting standard deviation of earnings at age t as $S_{E,t}$, and the standard deviation of earnings as S_x, $x \in \{A, \phi, \bar{k}\}$, we approximate the standard deviation of earnings by using a first-order taylor approximation with respect to each parameter:

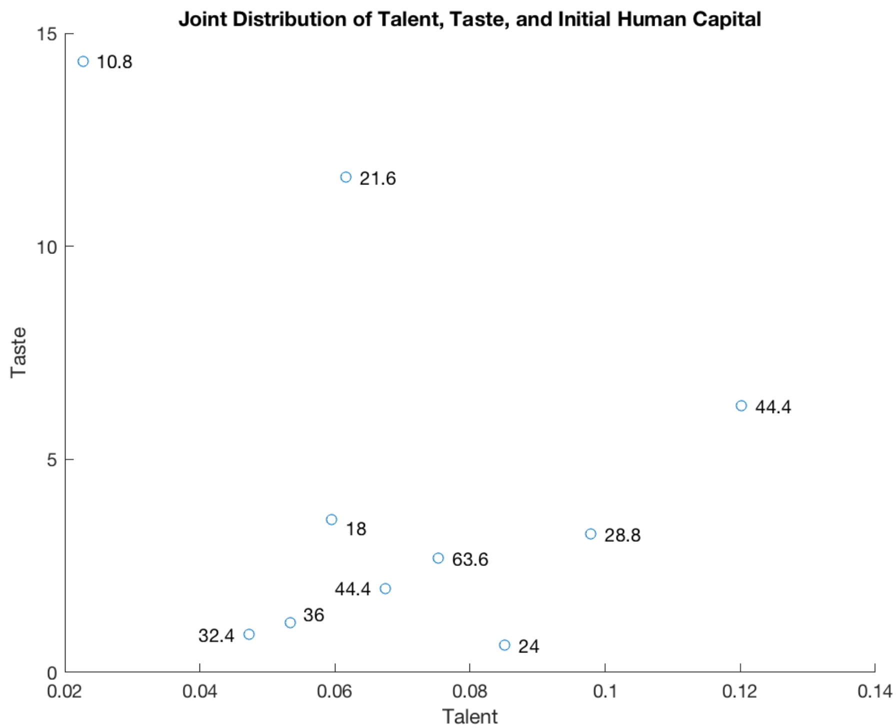

Figure 4: This figure depicts point estimates for the joint distribution of talent A_i, taste ψ_i, and initial human capital at age 30 (interpreted as an hourly wage) RK_i. Talent is on the x-axis, taste is on the y-axis, and the number next to each point is the initial human capital. Importantly for our results, our model moments put our fit to include people who work many hours (have a low distaste for labor) but are relatively untalented (have a reactively low wage), while there are people with high talent and high initial wage who work relatively little.

$$S_{E,t} \approx \overline{S_{E,t}} + \left.\frac{\partial S_{E,t}}{\partial S_A}\right|_{S_A=\overline{S_A}} \left(S_A - \overline{S_A}\right) + \left.\frac{\partial S_{E,t}}{\partial S_\phi}\right|_{S_\phi=\overline{S_\phi}} \left(S_\phi - \overline{S_\phi}\right) + \left.\frac{\partial S_{E,t}}{\partial S_k}\right|_{S_k=\overline{S_k}} \left(S_k - \overline{S_k}\right) + O(x^2)$$

Or, denoting the percentage in each variable with Δ, this becomes:

$$\Delta S_{E,t} \approx \epsilon_{A,t}\Delta S_A + \epsilon_{\phi,t}\Delta S_\phi + \epsilon_{k,t}\Delta S_k + \kappa \tag{20}$$

Where $\epsilon_{X,t} = \frac{S_X}{S_{E,t}} \left.\frac{\partial S_{E,t}}{\partial S_X}\right|_{S_X}$, the elasticity of the standard deviation of earnings with respect to the standard deviation of parameter $X \in \{A, \phi, k\}$, and κ is the residual category, representing the residual higher-order terms.

The relative importance of variation in each of talent, taste, and initial human capital is summarized in their corresponding elasticities. Because linearization fits small changes well, the value of an elasticity divided by the sum of all elasticities and residual well describes the marginal contribution to earnings variation of a component. That is, the contribution of taste is given by $\frac{\epsilon_{\phi,t}}{\epsilon_{A,t}+\epsilon_{\phi,t}+\epsilon_{k,t}}$. We numerically calculate the contribution of each concept at each age and depict them graphically in Figure 5.[14]

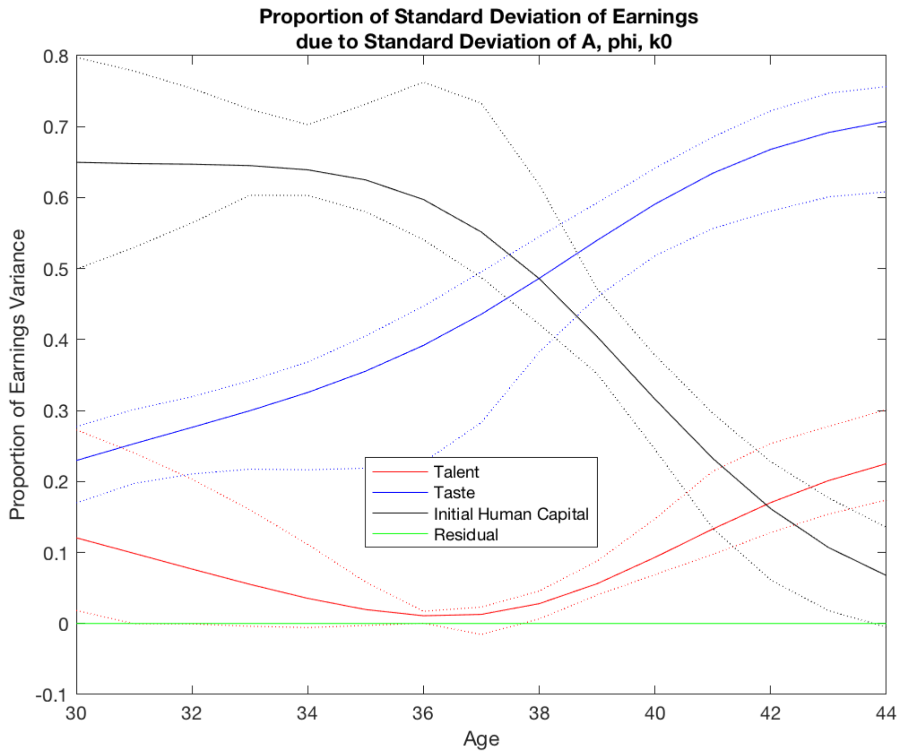

Figure 5: This figure summarizes our main decomposition results, depicting the three elasticities from equation 20. Each of the three visible lines (the residual category, a measure of how bad our linear approximation is, is not visible) indicates how much earnings variance at each age falls if variance in the corresponding parameter falls by 1%. These values are normalized by the total fall in earnings variance at each age.

Our decomposition gives a clear picture: first, initial human capital is important in explaining variance of income at early ages. This follows naturally as initial human capital is a mix of both ability and taste in earlier, unmodelled periods, and its importance naturally declines over time as both taste and talent become more important. Second, differences in talent at early ages contributes little to difference in earnings variation, because high-talent individuals spend much of their non-leisure time investing in human capital. Finally, by age 44 talent explains approximately 22% of income variation while taste explains 69%. In the next section, we discuss

[14]There are many ways to decompose these elasticities. For instance, to calculate $\epsilon_{A,t}$, we could hold V_ψ and $V_{\bar{k}}$ constant at the estimated mean, change S_A by 1% and calculate the change in income variance. Or we could do the same thing but evaluate S_ψ at 1% decreased variance for the entire exercise. Numerically, these make little difference (we take the average of all possible decompositions). This reference-dependence parallels the Oaxaca decomposition in labor economics.

how this breakdown is largely caused by the correlation of earnings and hours, and differences in "permanent" hours choices that do not decrease by age. Finally, the near-absence of the residual $O(x^2)$ category (the difference between actual change in variance of income and change predicted by the talent, taste, and human capital terms in equation 20) suggests our linear approximation is good.

4.1 Why is taste so important?

Our decomposition provides a clear statement: when we allow for differences in taste, they crowd out differences in talent in the estimation procedure of ages 30-44. Why is this the case? Implicitly, taste that causes one to be willing to spend more hours in human capital accumulation is quite similar to talent at human capital accumulation. To illustrate this, we depict three comparison paths. First, we plot a baseline path of a household working 2000 hours/year on average from age 30-44, with a \$30/hour initial hourly wage, and a \$45/hour hourly wage at age 44. We then plot two comparable counterfactual paths: one in which distaste for labor is lowered until the household works 2080 hours on average while talent and initial human capital are held constant, and the second in which talent is raised while distaste for labor and initial human capital are held constant until the household's final wage is equal to the final wage in the low-distaste for labor path. The two counterfactual paths have nearly identical hourly wage paths, and we can compare their labor and earnings behavior to understand the role of talent and taste in the estimation procedure. We depict the relevant paths in Figure 2, denoting human capital as "hourly wage."

Figure 2 makes clear the comparison: our two counterfactual paths are nearly identical in human capital accumulation: one from higher investment, the other from higher talent. A key difference between the two parameters is that lower disutility of labor increases total labor hours (both because of higher labor for income and for human capital accumulation) while higher talent reduces total labor hours and labor for income. Because much of it is taken as leisure, higher talent means that earnings don't increase by much, even as labor decreases: this lowers the correlation between labor and earnings. Lower disutility causes earnings and labor to move together, causing a higher correlation between labor and earnings. These paths make clear the importance of accurately hitting the correlation between total labor hours and earnings.

As an additional exercise, Figure 6 depicts the components of our decomposition in Figure 5 if we had fitted our moments to a correlation that was 0.2 lower (or higher) than the actual NLSY79 correlation. It confirms the importance of the correlation between hours and earnings. If the correlation was 0.2 lower than our baseline calibration, then the gap in the proportion of variation attributable to taste rather than talent would fall to 27% (from 46%), while if it was 0.2 higher, it would rise to 67%. This exercise makes clear that the targeting the age path of correlation between hours and earnings is extremely important target for a model describing preference heterogeneity.

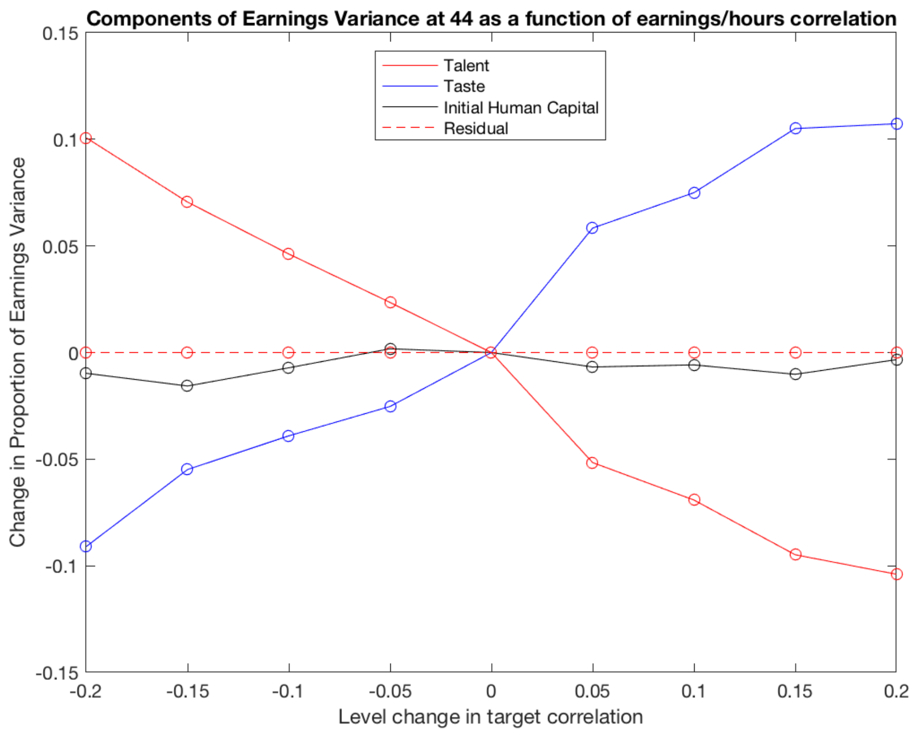

Figure 6: This figure depicts the change in the proportion of earnings variation at age 44 attributable to talent, taste and initial human capital for various level changes in the correlation between earnings and labor hours. Each elasticity is calculated using the elasticities of equation 20. As we increase the level of the correlation between hours and earnings throughout agents lifetimes, so that people who earn more typically work more, our model puts more emphasis on taste as the driving force behind earnings variation.

While the correlation between hours and earnings is an important sign of how much income variation comes from taste vs. talent, so is the standard deviation of hours. Intuitively, a large and persistent standard deviation in "permanent" hours can only be driven by taste, rather than talent. High levels of talent increase the slope of labor hours, but do not greatly change their level, particularly for ages 30-44. An examination of Figure 2 makes clear that in our model only differences in taste, rather than talent, can generate a relatively constant standard deviation in permanent labor hours. Figure 7 depicts a similar exercise as Figure 6, but changing the target standard deviation in hours, rather than the correlation between hours and earnings. The gap between taste and talent increases to as much as 65% (from 50%) when we increase the target level of the standard deviation in hours per year by 100, or falls to 36% when we decrease the target by 100 hours per year.

To confirm our intuition, we examine the results of one final exercise: we fit to the same

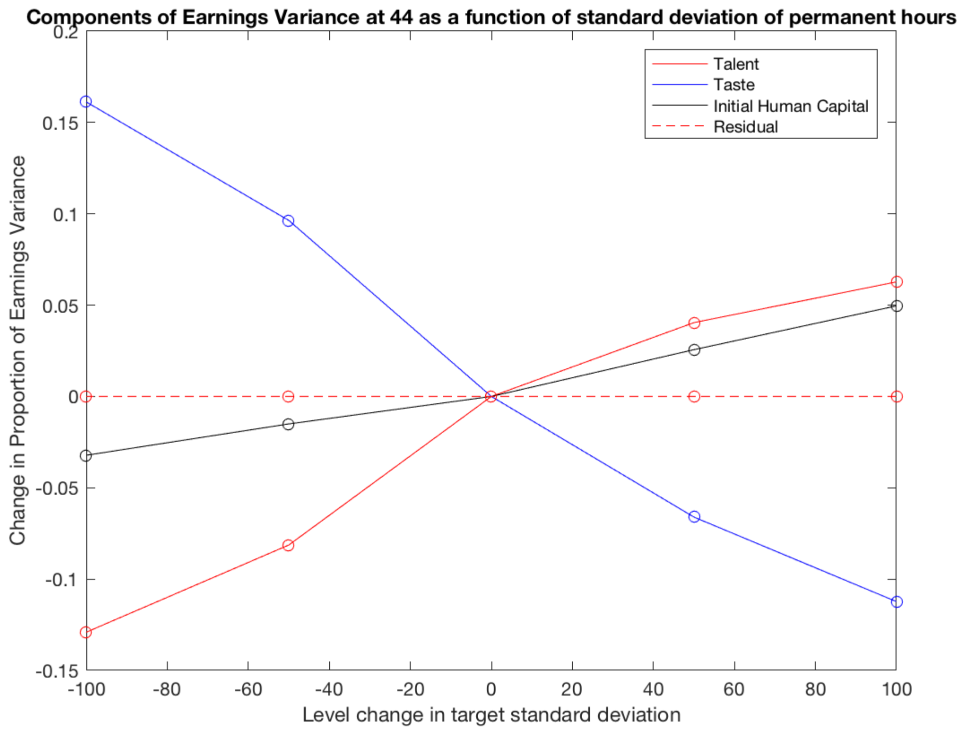

Figure 7: This figure depicts the change in the proportion of earnings variation at age 44 attributable to talent, taste and initial human capital for various level changes in the standard deviation of labor hours. Each elasticity is calculated using the elasticities of equation 20. As we increase the "permanent" variation in hours worked, so that there exist some people working many hours and some working few, our model puts more emphasis on taste as the driving force behind earnings variation.

set of six moments, but alter the age-path of the standard deviation of hours and the age-path of correlation of hours and earnings. Rather than having the standard deviation of hours stay relatively constant, around 430 hours/year, we reduce it to 345 at age 30 and have it fall to 255 by age 44. Additionally, rather than a u-shaped correlation between hours and earnings that does not display a long-run fall (starting at 0.49 and ending at 0.49) to a dramatic reduction, starting at 0.35 and ending at -0.35. Doing so would result in the proportion of income variation at age 44 attributable to taste alone falling from 68% to 26%, while variation attributable to talent alone increases to 48%. While this alternative calibration is not supported by the data, it highlights the type of data that a "talent-only" style model requires.

4.2 How does varying the importance of taste shift optimal tax rates?

Heterogeneity in preferences is a necessary but not sufficient condition for a utilitarian social planner's solution to change. If, as Mirrlees assumed, all heterogeneity is on labor preferences (as in equation 16) then in a model of inelastic labor supply, the normative distinction between taste and luck is not present. If however, all heterogeneity was on consumption preferences, then accurately assessing taste heterogeneity is crucial for the social planner. While we have identified taste heterogeneity, we are unable to assess whether or not it comes from heterogeneity in consumption's benefit or labor's cost: monotonic transformations of the utility function can yield the same labor, consumption, and study paths.

It is reasonable to assume that, given preferences differ, both consumption preferences and labor preferences differ. We therefore create a monotonic transformation, $\zeta_i(\alpha)$ that allows us to control the degree to which heterogeneity is on consumption rather than labor preferences

$$U_i(c, \ell, \phi_i) = \zeta_i(\alpha)\frac{c_t^{1-\sigma}}{1-\sigma} - \zeta_i(\alpha)\phi_i\frac{(1-\ell_t)^{1+\eta}}{1+\eta} \tag{21}$$

Where $\zeta_i(\alpha)$ is connected to our ϕ_i's by the monotonic transformation: $\zeta_i = \alpha + (1-\alpha)\frac{1}{\phi_i}$. When $\alpha = 0$, $\zeta_i = \frac{1}{\phi_i}$, and all heterogeneity is on consumption preferences, while when $\alpha = 1$, all heterogeneity is on labor preferences. We choose $\alpha = 0.5$ in our baseline optimal tax scenario to "split the difference" between consumption and labor heterogeneity, but examine the spectrum of α. To display how optimal tax rates can change when taste is made more or less important, we introduce simple tax scheme, in which the government levies a flat tax on labor income and imposes a uniform transfer, so that the net present value budget constraint becomes:

$$\sum_{a=1}^{A}\left(\frac{1}{1+r}\right)^{a-1} c_a = \sum_{a=1}^{A}\left(\frac{1}{1+r}\right)^{a-1}((1-\tau)w_a n_a + T)$$

Where, in equilibrium, T is the average tax payment recieved by the government across all N individuals and all A ages:

$$T = \frac{1}{N \cdot A}\sum_{a=1}^{A}\sum_{i=1}^{N} w_a n_a \tau$$

The government maximizes utilitarian welfare with equal pareto weights. Denoting the utility of individual i at age t as $U_{i,t}$, and assuming a population uniform across ages, the government's problem simplifies to:

$$\max_{\tau,T}\sum_{i=1}^{N}\sum_{a=1}^{A} U_{i,a}$$

The government's problem is a function of the joint distribution of A_i, ψ_i, and $\bar{k}_i$. By shifting the variation in income due to ψ_i and replacing it with variation in income due to A_i, we are able to answer the question "how much does the optimal tax rate vary as a function of the proportion of variation due to taste vs. talent?" Figure 8 depicts our results.

We find that in our baseline calibration ($\alpha = 0.5$), the percentage change in optimal tax rate

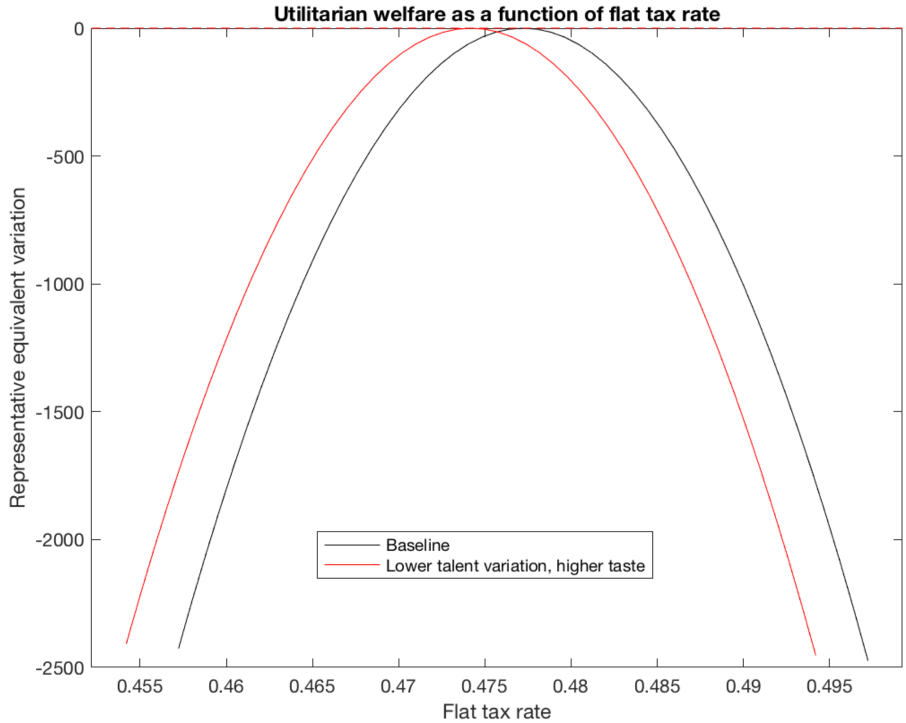

Figure 8: This figure depicts the equivalent variation of changing the tax rate from its optimum in our baseline setup and when we exchange sources of variation from talent to taste, keeping income variation at age 44 constant, with $\alpha = 0.5$ (heterogeneity on both consumption and labor preferences). The black line depicts the utility loss from changing the tax rate and resultant transfer (measured by the total change in utility of all agents, converted into dollars using the average marginal utility of money at the optimum). The red line depicts the same exercise after reducing population variation in talent and increasing it in taste. The shift from the black line's apex to the red line's apex shows the change in optimal taxation caused by changes in the sources of income variation.

τ^* generated by a 1% decrease in income variation at age 44 due to talent, combined with a 1% increase in income variation at age 44 due to taste is -0.63%.[15] With a baseline optimal tax rate of 48%, this represents a 0.30% percentage point decrease in the optimal tax rate: the optimal tax rate is highly sensitive to the sort of preference heterogeneity revealed by the data. We measure the equivalent variation generated by moving from suboptimal tax rates to the optimal tax rate for both the baseline calibration and the lower talent, higher taste calibration in Figure

[15]Specifically, we change the distribution of σ_ψ by $\frac{1}{\epsilon_{\sigma_\psi}}$, where ϵ_{σ_ψ} is the elasticity of income variation at age 44 with respect to σ_ψ, and the distribution of σ_A by $\frac{1}{\epsilon_{\sigma_A}}$, with the same notation. The mean-preserving spreads are calculated as:

$$A_i' = A_i + \omega(A_i - \bar{A})$$

where ω is the scaling factor and $\bar{A}$ is the average of all A_i's.

8.[16]

We emphasize that the sort of presence heterogeneity our model detects is a necessary, but not sufficient condition for large changes in the optimal tax rate. If all preference heterogeneity came from utility of consumption, the change would be larger, while if all preference heterogeneity came from disutility of labor, there would be less change in the optimal tax rate, as the utilitarian social planner would not take the normative distinction about differences in preferences into account. To illustrate this, we depict the optimal tax rate as a function of α for both our baseline calibration and one in which 1% of income variation at age 44 is exchanged between talent and taste in Figure 9.

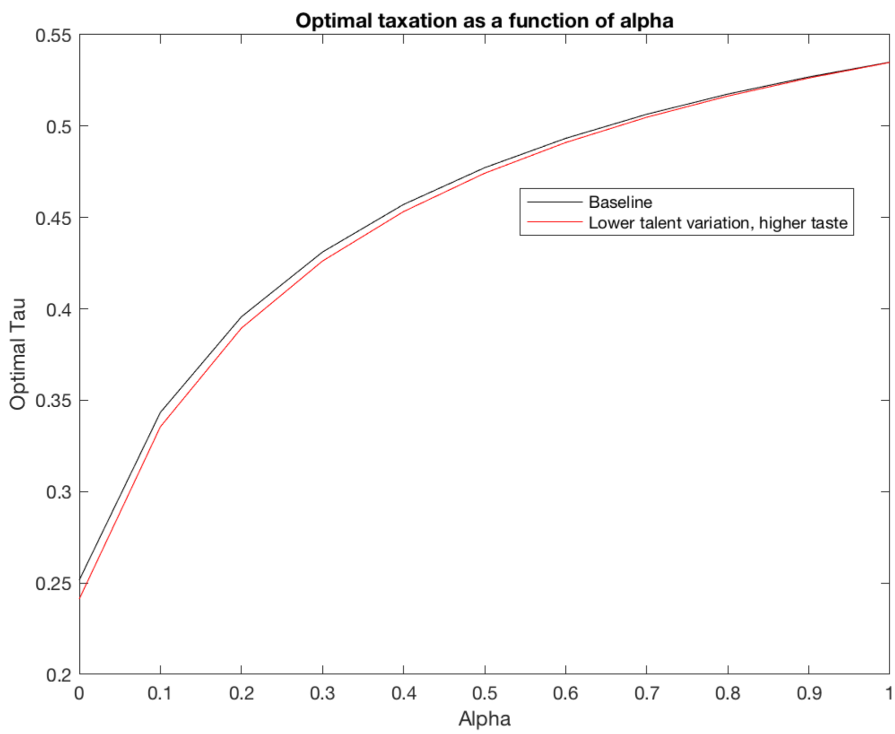

Figure 9: This figure depicts the two optimal tax rates as a function of α, which controls where heterogeneity in preference lies. When $\alpha = 0$, the heterogeneous term multiplies consumption preferences, and all heterogeneity lies on consumption. When $\alpha = 1$, heterogeneity shifts to labor preferences. When heterogeneity lies on labor preferences, the two tax rates are both higher and highly similar. When heterogeneity lies on consumption preferences, the impetus for redistributive taxation is lower, and the optimal tax rate is more sensitive (both in percentage and absolute terms) to the sources of income variation.

Figure 9 shows that when α is near zero, we attribute all heterogeneity in taste to con-

[16]We measure the utilitarian's equivalent variation by taking the overall increase in utility and dividing by the average marginal utility of income.

sumption preferences: people who consume a lot may do so because they have higher taste for it, and the optimal tax is low. However, as α increases, and heterogeneity is placed on labor preferences, rather than consumption preferences, the optimal tax rate rises, as homogeneity in consumption preferences, combined with the strong diminishing marginal utility of consumption, yields a large benefit to redistribution. Consistent with intuition, an income variation preserving increase in talent variation increases the tax rate more when consumption preferences are the source of heterogeneity. When consumption preferences are the sole source of preference heterogeneity ($\alpha = 0$), the optimal tax rate has an elasticity of 4 with respect to talent (rather than taste) as the source of income variation. When labor preferences are the sole source of preference heterogeneity ($\alpha = 1$), the same elasticity falls by two orders of magnitude, to 0.03. In our baseline calibration ($\alpha = 0.5$), the elasticity is 0.63.

5 Conclusion

Our paper makes four 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, the optimal level of redistribution depends not just on talent, but also on the extent to which individuals differ in their preferences for consumption and leisure. Third, using standard panel data sources we cannot identify tastes for consumption and leisure with an arbitrary utility function. However, we can identify them using reasonable functional form assumptions and panel data on both earnings and hours of leisure. If we make these assumptions and estimate the model using data from the NLSY79, we find a substantial degree of variation in relative tastes for consumption versus leisure. Finally, the optimal tax policy depends crucially on whether the variation in relative tastes is due to variation in the marginal utility of consumption or in the marginal utility of leisure.

Expanding on the last point, optimal redistribution depends on the scaling of individuals' utility functions. Since differently scaled utility functions lead to observationally equivalent behavior, we cannot know how to apportion variation in relative tastes between the marginal utility of consumption and the marginal utility of leisure. Therefore, even if we are willing to accept the strong assumptions required the estimate relative tastes, optimal tax policy still depends crucially on untestable assumptions about the preferences of individuals. None of this is a concern if preferences do not vary across individuals, but we find evidence suggesting that they do vary, perhaps substantially.

Our results are driven by two empirical findings. First, earnings and hours worked are highly correlated and this correlation does not decline over the lifecycle. Second, even highly-attached prime age males display significant variation in permanent labor hours in the cross-section that does not decline with age. When these two facts are viewed through the lens of our model, we find that tastes are an important driver of income inequality. Moreover, the optimal tax rate is highly sensitive to the relative importance of tastes versus talent in a simple flat tax constant transfer regime.

We acknowledge that our model is simple and abstracts from many important factors that have been identified in the literature on lifecycle human capital investment and labor supply. And we acknowledge the likelihood that alternative models exist which may find a smaller role for tastes. But any such model must be able to generate persistent variation in permanent labor hours using talent alone. This might be done, for instance, with persistent shocks that reduce lifecycle labor hours on the intensive margin. Moreover, any alternative model must be able to generate a high and persistent correlation between hours worked and earnings. One possiblility might be to incorporate capital imperfections which force otherwise identical households 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 standard deviation of hours and the correlation of hours and earnings over the lifecycle are crucial moments to target for any model of optimal taxation.

References

J M Abowd and D Card. On the Covariance Structure of Earnings and Hours Changes. *Econometrica*, 57(2):411–445, 1989. ISSN 0012-9682. doi: 10.2307/1912561.

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. ISSN 0033-5533. doi: 10.1162/qjec.122.3.969. URL http://www.nber.org/papers/w12082.

Mark Aguiar, Erik Hurst, and Loukas Karabarbounis. Recent developments in the economics of time use. *Annual Review of Economics*, 4, 2012. ISSN 1057-9230. URL http://faculty.chicagobooth.edu/erik.hurst/research/recent_developments_time.pdf.

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. ISSN 0021-8499.

Regis Barnichon and Andrew Figura. Declining Desire to Work and Downward Trends in Unemployment and Participation. 2015.

Susanto Basu and Miles S. Kimball. Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption. 2003.

Raymond C Battalio, John H Kagel, and Morgan O Reynolds. Income Distributions in Two Experimental Economies. *Journal of Political Economy*, 85(6):1259–1271, 1977.

Yoram Ben-Porath. The Production of Human Capital and the Life Cycle of Earnings. *Journal of Political Economy*, 75(4):352–365, 1967. URL http://www.jstor.org/stable/1828596 papers2://publication/uuid/A3865742-682F-409D-BAF6-2D44B10F6053.

Katy Bergstrom and William Dodds. Sources of Income Inequality: Productivities vs. Preferences. 2018.

Adam Blandin. Learning by Doing and Ben-Porath: Life-cycle Predictions and Policy Implications. *Journal of Economic Dynamics and Control*, 90:220–235, 2018. ISSN 01651889. doi: 10.1016/j.jedc.2018.03.002.

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

Richard Blundell, Luigi Pistaferri, and Itay Saporta-Eksten. Consumption Inequality and Family Labor Supply. *American Economic Review*, 106(2):387–435, 2016. doi: 10.1257/aer.20121549.

Robin Boadway, Maurice Marchand, Pierre Pestieau, and Maria Del Mar Racionero. Optimal Redistribution with Heterogeneous Preferences for Leisure. *Journal of Public Economic Theory*, 4(4):475–498, 2002.

Audra J. Bowlus and Chris Robinson. Human capital prices, productivity, and growth. *American Economic Review*, 102(7):3483–3515, 2012. ISSN 00028282. doi: 10.1257/aer.102.7.3483.

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.

Raj Chetty. Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica*, 80(3):969–1018, 2012. ISSN 0012-9682. doi: 10.3982/ECTA9043. URL `http://doi.wiley.com/10.3982/ECTA9043`.

Raj Chetty, Adam Guren, Day Manoli, and Andrea Weber. Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review*, 101:471–475, 2011. ISSN 00028282. doi: 10.1257/aer.101.3.471.

Flavio Cunha and James Heckman. Decomposing Trends in Inequality in Earnings into Forecastable and Uncertain Components. *Journal of Labor Economics*, 34(2), 2016.

Paul Gomme and Peter Rupert. Theory, Measurement and Calibration of Macroeconomic Models. *Journal of Monetary Economics*, 54:460–497, 2007. ISSN 03043932. doi: 10.1016/j.jmoneco.2005.09.005.

Peter Gottschalk and Robert Moffitt. The Rising Instability of U.S. Earnings. *Journal of Economic Perspectives*, 23(4):3–24, 2009. ISSN 0895-3309. doi: 10.1257/jep.23.4.3.

Peter Gottschalk, Robert Moffitt, Lawrence F Katz, and William T Dickens. The Growth of Earnings Instability in the U . S . Labor Market. *Brookings Papers on Economic Activity*, 1994(2):217–272, 1994. ISSN 0007-2303. doi: 10.2307/2534657. URL `http://www.brookings.edu/ /media/Projects/BPEA/1994-2/1994b_bpea_gottschalk_moffitt_katz_di`

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. ISSN 15424766. doi: 10.1111/j.1542-4774.2011.01061.x.

Fatih Guvenen, Burhanettin Kuruscu, and Serdar Ozkan. Taxation of human capital and wage inequality: A cross-country analysis. *Review of Economic Studies*, 81(2):818–850, 2014. ISSN 1467937X. doi: 10.1093/restud/rdt042.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*, 104(7): 2075–2126, 2014. ISSN 00028282. doi: 10.1257/aer.104.7.2075.

James J. Heckman. A Life-Cycle Model of Earnings, Learning, and Consumption. *Journal of Political Economy*, 84(4):S11–S44, 1976a. URL `http://www.jstor.org/stable/1831101`.

James J. Heckman. *Estimates of a Human Capital Production Function Embedded in a Life-Cycle Model of Labor Supply.* 1976b. ISBN 0870145150.

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

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.

Mark Huggett, Gustavo Ventura, and Amir Yaron. Sources of lifetime inequality. *American Economic Review*, 101(7):2923–2954, 2011. ISSN 00028282. doi: 10.1257/aer.101.7.2923.

Daniel Kahneman. *Thinking, Fast and Slow.* Farrar, Straus and Giroux, 2011. ISBN 1429969350. URL `https://books.google.com/books?id=ZuKTvERuPG8C&pgis=1`.

Michael P. Keane and Kenneth I. Wolpin. The Role of Labor and Marriage Markets, Preference Heterogeneity, and the Welfare System in the Life Cycle Decisions of Black, Hispanic, and White Women. *International Economic Review*, 51(3):851–892, 2010.

J Kennan. A Note on Discrete Approximations of Continuous Distributions. *Mimeo*, 1(September), 2006. URL `http://www.ssc.wisc.edu/ jkennan/research/DiscreteApprox.pdf%5Cnpapers3://publication/uuid`

Robert G. King, Charles I. Plosser, and Sergio T. Rebelo. Production, Growth and Business Cycles. *Journal of Monetary Economics*, 21:195–232, 1988. ISSN 03043932. doi: 10.1016/0304-3932(88)90034-7.

Donghoon Lee and Kenneth I Wolpin. Intersectoral Labor Mobility and the Growth of the Service Sector. *Econometrica*, 74(1):1–46, 2006. URL `http://www.jstor.org/stable/3598922`.

Donghoon Lee and Kenneth I. Wolpin. Accounting for wage and employment changes in the US from 1968-2000: A dynamic model of labor market equilibrium. *Journal of Econometrics*, 156(1):68–85, 2010. ISSN 03044076. doi: 10.1016/j.jeconom.2009.09.008. URL http://dx.doi.org/10.1016/j.jeconom.2009.09.008.

Benjamin B. Lockwood and Matthew Weinzierl. De Gustibus non est Taxandum: Heterogeneity in preferences and optimal redistribution. *Journal of Public Economics*, 124:74–80, apr 2015. ISSN 00472727. doi: 10.1016/j.jpubeco.2015.01.002. URL http://www.sciencedirect.com/science/article/pii/S0047272715000134.

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. ISSN 03085147 (ISSN). doi: 10.1080/03085140120089063.

Rodolfo E Manuelli, Ananth Seshadri, and Yongseok Shin. Lifetime Labor Supply and Human Capital Investment. 2012. URL http://research.stlouisfed.org/wp/2012/2012-004.pdf February.

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 BT - Handbook of Income Distribution, Volume 1. In A B Atkinson and F Bourguignon, editors, *Handbook of Income Distribution, Volume 1*, volume 1, pages 379–427. Elsevier Science, 2005. URL http://www.sciencedirect.com/science?_ob=MImg&_imagekey=B7RKR-4FF32Y6-B-1&_cdi=25729&_user=

Paolo Giovanni Piacquadio. A Fairness Justification of Utilitarianism. *Econometrica*, 85(4):1261–1276, 2017. ISSN 0012-9682. doi: 10.3982/ECTA14151. URL https://www.econometricsociety.org/doi/10.3982/ECTA14151.

Solomon W Polachek, Tirthatanmoy Das, and Rewat Thamma-Apiroam. Micro- and Macroeconomic Implications of Heterogeneity in the Production of Human Capital. *Journal of Political Economy*, 123(6):1410–1455, 2015.

Martin Ravallion. Inequality when Effort Matters. jul 2015. URL http://www.nber.org/papers/w21394.

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.

Mariacristina Rossi and Serena Trucchi. Liquidity Constraints and Labor Supply. *European Economic Review*, 87:176–193, 2016. ISSN 00142921. doi: 10.1016/j.euroecorev.2016.05.001. URL http://dx.doi.org/10.1016/j.euroecorev.2016.05.001.

Emmanuel Saez and Stefanie Stantcheva. Generalized Social Marginal Welfare Weights for Optimal Tax Theory. *American Economic Review*, 106(1):24–45, 2016. ISSN 0002-8282. doi: 10.1257/aer.20141362.

Stefanie Stantcheva. Learning and (or) Doing: Human Capital Investments and Optimal Taxation. jul 2015. URL http://www.nber.org/papers/w21381?utm_campaign=ntw&utm_medium=email&utm_source=ntw.

Appendix A: Flexible Frisch Elasticity of Labor Supply

In this Appendix, we consider our main exercise of allowing the Frisch elasticity ϵ to be calibrated, along with the population of triples $\{A_i, \phi_i, \overline{k}_i\}$, to best fit our moments. Because the Frisch elasticity controls the responsiveness of labor conditional on wages, it plays a role in determining the slope of hours given the slope of wages (or earnings), or the slope of earnings given the slope of hours. While the majority of the age-paths of our baseline calibration moments fit the data well, our agents hours responded "too much" as wages rose over the lifecycle, rising 154 hours rather than 89 hours from ages 30 to 44.

One way of better fitting the lifecycle path of labor conditional on earnings is to reduce the elasticity of labor supply, making labor's path flatter (ceteris paribus). In the context of our preferences, it reduces both the substitution and income effect of a wage change. Because the way differences in talent change the lifetime level (as opposed to the slope) of labor is through the income effect, taste becomes more varied in order to explain lifetime differences in labor levels.

Consistent with this intuition, when we jointly fit the Frisch elasticity along with the joint distribution of $\{A_i, \phi_i, \overline{k}_i\}$, the Frisch elasticity falls from 0.33 to 0.22, better fitting the hours data (simulated hours rise by 83 by the age of 44, less than the long-differenced values). Our fitted moments and paths are depicted in Figure A.1 below. While estimated variation in talent increases, a lower elasticity means that changes in talent play a smaller role in hours choices, even as changes in taste have the same effect on total hours. Consequently, its contribution to earnings variation declines slightly, so that heterogeneity in taste explains 74% of earnings variation at age 44, rather than 68%.

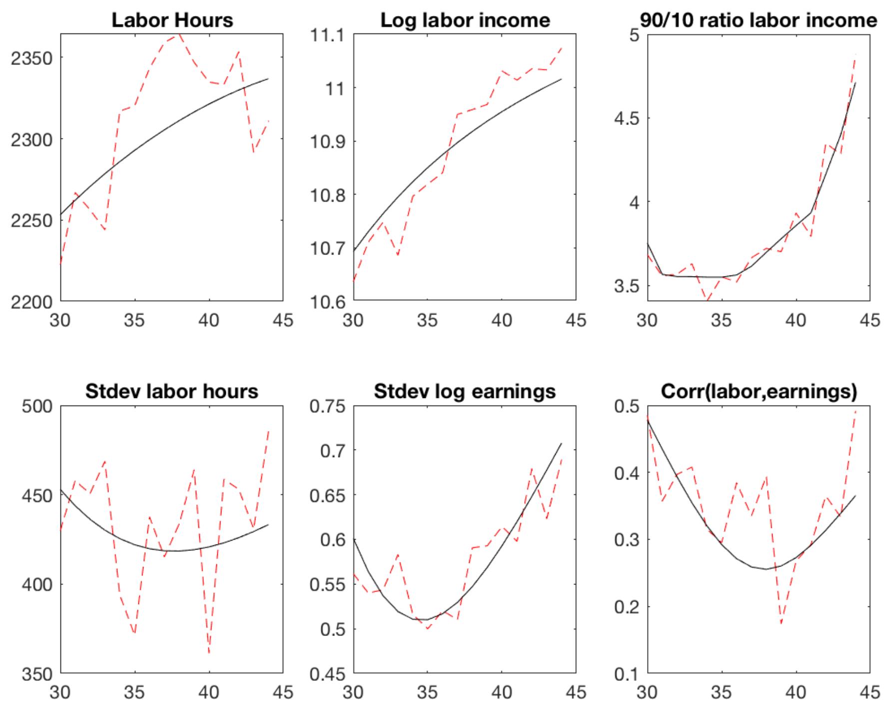

Figure A.1: This figure depicts the six sets of empirical moments from the NLSY79 (red line) and model fits with flexible epsilon. It may be compared to figure 3, which fixes the Frisch elasticity at 0.33, rather than using it to fit our moments, which results in an elasticity of 0.22.

Appendix B: Calculation of Moments

In this Appendix, we describe the method we used to calculate moments using the NLSY79. As is common in the structural labor literature, we interpreted deviations from the model at the individual level as coming from either measurement error or unforeseen shocks. That is, individuals are subject to possibly correlated shocks to both their (log) annual earnings and their annual work hours (which include both paid work time and human capital investment). Following the literature, we assume that individuals have access to complete markets so that they can insure against these shocks and their optimization problem reverts to the perfect foresight model in the paper. The presence of these shocks does not affect the estimates of mean earnings or work hours, but it does affect estimates of their variances and covariance.

We estimate the variances and covariance of log annual earnings and annual hours of work as follows.[17] First, for all men in our sample, ages 30 to 44, we regress log annual earnings on age separately for each individual. Then we calculate the variance of the residuals from these regressions across all men in our sample. This gives us an estimate of the variance of the shock to earnings. We do the same thing for annual hours of work. To estimate the variances of log annual earnings and annual hours, we simply calculate the raw variances and subtract the estimated variance of the shocks. To calculate the covariance of log annual earnings and annual hours, we calculate the raw covariance and subtract the estimated covariance of the shocks.

We estimate the 90-10 ratio of annual earnings in a slightly different way. For all men in our sample, ages 30 to 44, we regress log annual earnings on age separately for each individual and store the fitted values at each age. Then we calculate the 90-10 ratio of the (exponentiated) fitted values at each age.

[17] In calculating all moments, we use the sampling weights provided in the NLSY79.

Appendix C: Moment Robustness

In this Appendix, we describe alternative approaches to calculating the moments using the NLSY79. Specifically, we:

1. Calculate the moments using either a simple linear slope or a random slope model (with Bayes shrinkage).
2. Restrict the sample to observations with a minimum number of hours equal to either 1 or 100.
3. Truncate the minimum number of hours from below at 0 (no truncation) or 200.
4. Truncate the maximum number of hours from above at 4000 hours or no truncation.

Together, these four potential binary choices yields sixteen possible treatments of the data.

For parsimony, Table A.1 below summarizes the effect of each of these sixteen possible treatments on the level of the standard deviation of hours from ages 30-44, the level of the correlation between hours and earnings from ages 30-44, and the linear and quadratic trend of the correlation of hours and earnings from those same ages.

Table A.1: Data Robustness

Rand. slope	Min. hours	L. trunc.	R. trunc.	Std. Dev.	Corr Level.	Corr. slope	Corr quad.	% variation due to taste at age 40
0	0	0	0	415.92	0.32	-0.0032	0.0026	0.683
0	0	0	1	400.26	0.32	-0.0031	0.0025	0.677
0	0	1	0	434.14	0.30	-0.0039	0.0029	0.692
0	0	1	1	417.13	0.31	-0.0039	0.0027	0.674
0	1	0	0	415.62	0.32	-0.0032	0.0027	0.684
0	1	0	1	400.03	0.32	-0.0032	0.0025	0.683
0	1	1	0	433.86	0.30	-0.0039	0.0029	0.693
0	1	1	1	416.92	0.30	-0.0039	0.0028	0.673
1	0	0	0	412.54	0.32	-0.0036	0.0024	0.674
1	0	0	1	397.43	0.33	-0.0035	0.0023	0.671
1	0	1	0	430.87	0.30	-0.0044	0.0027	0.678
1	0	1	1	414.39	0.31	-0.0043	0.0025	0.666
1	1	0	0	412.45	0.32	-0.0036	0.0024	0.676
1	1	0	1	397.34	0.32	-0.0035	0.0023	0.669
1	1	1	0	430.78	0.30	-0.0044	0.0027	0.679
1	1	1	1	414.31	0.31	-0.0043	0.0025	0.666

Table A.1: This table displays the data and results robustness of the two particularly important moments that drive the importance of preferences. We offer sixteen treatments of the data. “Rand slope” calculates data moments using a random slope method or a linear slope. When “Min hours” is on, it requires workers to work at least 100 hours a year to be in the data, rather than one. “L. trunc” truncates the hours distribution we use to estimate our moments at 200 when on (zero otherwise). “R. trunc” truncates the hours distribution we use to estimate our moments at 4000 when on (no bound otherwise). “Std Dev.” gives the resultant average level of the standard deviation of hours, while the next columns describe the resultant correlation level, slope, and quadratic terms as a function of age minus 37, so that the intercept represents the average age, and the slope is the slope of the correlation at 37. The last column gives the values for the proportion of earnings variation attributable to taste at age 44.