Heterogeneity in Talent or in Tastes? Implications for Redistributive Taxation

Ian Fillmore
Trevor Gallen

Preliminary and Incomplete: Do not cite

May 2018

Abstract

How much of income heterogeneity is due to taste vs. talent? We use a simple model of human capital accumulation inspired by Neal and Rosen (2005) 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 to estimate the joint distribution of taste and talent. Filtering out idiosyncratic hours shocks from the NLSY79 and focusing on highly-attached prime-age males, we find that even for this group, heterogeneneity in taste plays a large role: 71% of income variation at age 44 is due solely to taste, rather than talent or initial conditions. We find that the moments driving high heterogeneity in taste are the relatively high standard deviation in "permanent" labor hours and the positive correlation between labor hours and earnings that does not strongly decline by age. Finally, we show that exchanging the sources of income variation changes utilitarian optimal tax rates significantly, particularly when heterogeneity is due to differences in the marginal utility of consumption, rather than leisure.

1 Introduction

What are the sources of earnings inequality? A social planner facing a given distribution of income would behave differently if inequality was driven by heterogenous tastes for consumption instead of heterogenous talent. In both cases, he would move resources to agents with the highest marginal utility from consumption, but in the first case this may translate to a regressive tax system, while in the case of heterogenous talent, the most unlucky agents would receive positive transfers. An implicit assumption in most optimal tax models is that unchosen "luck" or "talent" is the sole determinant of earnings inequality (Mirrlees, 1971), and the problem of redistribution is one of maximizing utility while maintaining incentives to work. However, utilitarian calculus may change dramatically if an important determinant of earnings variation comes from taste: say, because the consumption bundle that some households prefer is heavy on expenditure, while for others it is heavy on non-labor time.

This paper asks two questions: first, how much of earnings inequality in the United States is due to tastes, rather than talent? Second, how do the sources of income heterogeneity affect results about optimal taxation? We find that a simple parameterized model suggests a significant role for taste, driven by a large difference in lifetime labor hours that is relatively constant in age, and a large positive correlation between hours and earnings that also does not decline in age. With a population calibrated to moments on the income and hours distribution by age, we find that small shifts in the sources of earnings heterogeneity can yield significant changes in a utilitarian optimal tax rate, particularly when agents are heterogeneous in the marginal utility of consumption, rather than of leisure.

We use a standard intertemporal model of labor supply from Neal and Rosen (2005), that allows for differences in the ability to accumulate human capital, which we call "talent," and the desirability of 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 variation.[1] Using data from the National Longitudinal Survey of Youth 1979 (hereafter NLSY) on the joint distribution of "permanent" earnings and hours by age, we estimate the distribution of taste and talent for a representative population. In our 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 71% of earnings variation in even highly-attached male workers.

What drives our results? A crucial moment that helps determine the role of taste and talent is the correlation of hours and earnings. Intuitively, when data display high earnings levels linked with a high hours of work, it suggests people who earn more do so because they are working more hours, and taste plays a large role in generating earnings inequality. When there is little

[1] See, for instance, Blandin (2016), Guvenen et al. (2014), Huggett et al. (2011).

correlation between earnings and hours of labor, it suggests that talent (and initial conditions) play a larger role. In addition to the level, our results also come from the time path of this correlation. A declining correlation between hours and earnings by age suggests significant differences in talent that accrue as increasing wages as agents age, while a steady correlation leaves less room for differences in talent, and suggests differences in taste are the primary factor driving earnings differences.

While the level and path of the correlation between hours and earnings are important, they are not the only drivers of our estimates. Even after isolating permanent components of earnings and hours by age from idiosyncratic year-to-year shocks, and focusing on highly-attached prime-age males, we still find significant variation in hours worked per year even among this group. This variation in hours, combined with a positive correlation of hours and earnings, leads our model to find that taste plays an important role.

This paper emphasizes that intertemporal choice and data helps identify taste vs. talent. Intuitively, the slope of earnings as a function of age and variation in the slope identifies variation in talent, while persistent differences in labor hours conditional on the same wage informs us about variation in taste. The correlation of hours and earnings informs us about the conditional distributions of talent conditional on taste (and taste conditional on talent). Of course, as emphasized in Neal and Rosen (2005), variation in taste alone may cause differences in the slope of earnings: a young person recognizing they want to live a high-consumption lifestyle should invest heavily in human capital while young, even if they have below-average talent. Similarly, a talented young person may be induced by his higher wage to work more hours over his lifetime than a less-talented peer even though they have the same taste for leisure.

However, tastes for leisure and talent can be identified under particular common functional form assumptions about preferences and the human capital production function. We use panel data from the NLSY79 on highly attached prime age males to fit the model and estimate the joint distribution of tastes and talent. We find that tastes account for 71% percent of the variation in earnings at age 44. 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.[2] Finally, we document that the magnitude of this result is dependent on where heterogeneity lies: if preference heterogeneity lies on consumption preferences, optimal

[2] 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.

taxes are strongly responsive to the source of income variation, while if it lies on labor preferences, optimal taxes are less responsive.

2 Literature Review

Mirrlees (1971) spearheaded the study of optimal tax design in the face of unobservable heterogeneity. and standard optimal taxation framework.

Contrary to the assumption usually made in the optimal taxation literature, casual empiricism suggests that at least some of the differences in people's life outcomes are driven by differences in their life goals, that is, their preferences. Kahneman (2011, p. 401) reports results from a study of students at elite colleges in 1976.

> When 17 or 18, the participants had filled out a questionnaire in which they rated the goal of "being very well-off financially" on a 4-point scale ranging from "not important" to "essential." The questionnaire they completed twenty years later included measures of their income in 1995, as well as a global measure of life satisfaction.
>
> ... Nineteen years after they stated their financial aspirations, many of the people who wanted a high income had achieved it. ...
>
> The importance that people attached to income at age 18 also anticipated their satisfaction with income as adults. We compared life satisfaction in a high-income group (more than $200,000 household income) to a low- to moderate-income group (less than $50,000). The effect of income on life satisfaction was larger for those who had listed being well-off financially as an essential goal: 0.57 points on a 5-point scale. The corresponding difference for those who had indicated that money was not important was only 0.12.

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, such as spending time with their families, will tend to pursue paths that allow them to do so.

Lifecycle human capital models provide a natural framework for thinking about how productivity (i.e. wage) differences later in life are driven by investment decisions early in life. For example, in the classic model of Ben-Porath (1967), an individual's ability to acquire new human capital, which we refer to as talent, provides the driving force behind differences in human capital investment, wages, and lifetime earnings. 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 most of their time to investment rather than work (either formal schooling or on-the-job training). As they get older, they accumulate human capital, 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.

Whereas in the static case a social planner would ideally redistribute on the basis of a worker's productivity, in the dynamic case the social planner would ideally redistribute on the basis of a worker's talent. It may very well be that a worker's talent is a much harder thing to measure than his productivity at a point in time and that proxying for talent with annual income, as governments must do in practice, presents an additional set of complications (Stantcheva, 2015). But the basic utilitarian logic continues to dictate redistribution away from those who are born with advantageous characteristics (productivity, talent, etc.) and toward those born without such characteristics. However, this logic breaks down if differences in wages arise due to differences in people's tastes.

Several economists have noted that, to the extent that observed income differences reflect tastes rather than talent, the utilitarian rationale for redistribution becomes weaker. Lockwood and Weinzierl (2015) formalize this intuition within a simple static economy. They argue that only income inequality driven by productivity differences justifies redistribution, and they show that heterogeneity in tastes will tend to dampen the optimal amount of redistribution. They provide suggestive evidence that such taste heterogeneity may be a significant factor in determining income inequality. Is it really plausible that tastes for leisure could generate wide dispersions in income? Moreover, wouldn't comparing hourly wages, rather than annual incomes, correct for any heterogeneity in tastes for leisure? The answers to these questions are, respectively, yes and no.

Neal and Rosen (2005) consider the consequences of adding a labor/leisure choice to the Ben-Porath model. They emphasize that

> ...in human capital models, earnings capacity is, in part, determined by tastes. Because workers with weak preferences for leisure forfeit relatively little in utility terms when they work, they have additional incentive to invest in market skills. In short, we expect future earnings capacity to vary inversely with current tastes for leisure.
>
> ...There are two components to the link between tastes and earnings. Given earnings capacity, workers with weak tastes for leisure work more. But earnings capacity is not independent of tastes. Weak tastes for leisure enhance investment and therefore raise earnings capacity. ...People with either weaker tastes for leisure (inferior opportunities for work outside the market) or lower costs of human capital production

[i.e. more talent] invest 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. Although both talent and tastes affect the path of human capital and earnings over the lifecycle, they contain very different implications for redistribution.

3 Lifecycle Model of Human Capital Investment

3.1 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 the agent's relative taste for leisure, and $\overline{k}$ is the agent's 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

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

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, with more human capital translating into higher wages. 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.

$$\max_{\{c_t s_t \ell_t n_t\}} \quad \sum_{t=1}^{T} \beta^{t-1} U(c_t, \ell_t; \phi) \tag{1}$$

$$s.t. \quad s_t + \ell_t + n_t = 1 \tag{2}$$

$$s_t,\ \ell_t,\ n_t \geq 0 \tag{3}$$

$$\sum_{t=1}^{T} \left(\frac{1}{1+r}\right)^{t-1} c_t = \sum_{t=1}^{T} \left(\frac{1}{1+r}\right)^{t-1} w_t n_t \tag{4}$$

$$w_t = R k_t \tag{5}$$

$$k_{t+1} = (1-\delta) k_t + A(s_t k_t)^{\gamma} \tag{6}$$

$$k_1 = \overline{k} \tag{7}$$

The Kuhn-Tucker conditions for this problem are

$$[c_t]: \quad \beta^{t-1} U_c(c_t, \ell_t; \phi) = \left(\frac{1}{1+r}\right)^{t-1} \lambda \tag{8}$$

$$[\ell_t]: \quad U_\ell(c_t, \ell_t; \phi) = \mu_t \tag{9}$$

$$[n_t]: \quad \left(\frac{1}{1+r}\right)^{t-1} \lambda w_t = \beta^{t-1}(\mu_t + \mu_t^n) \tag{10}$$

$$[n_t]: \quad \mu_t^n n_t = 0;\ \mu_t^n \geq 0;\ n_t \geq 0 \tag{11}$$

$$[s_t]: \quad A\gamma(s_t k_t)^{\gamma-1} k_t \lambda \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_\tau = \mu_t \tag{12}$$

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, the first order conditions for ℓ_t and s_t become

$$[\ell_t]: \quad \beta^{t-1} U_\ell(c_t, \ell_t; \phi) = \left(\frac{1}{1+r}\right)^{t-1} \lambda w_t \tag{13}$$

$$[s_t]: \quad \left(\frac{1}{\beta(1+r)}\right)^{t-1} \frac{R(s_t k_t)^{1-\gamma}}{A\gamma} = \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_\tau. \tag{14}$$

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 (14) captures the benefits of an additional unit of human capital today, in terms of increased future earnings. The left hand side represents the costs of investing in one more unit of human capital today, in terms of earnings in the current period. Equation (14) 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.

3.2 Strengths and Weaknesses 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. We consider our model to be almost maximally parsimonious compared to alternative possibilities, while still allowing strong roles for talent and taste. We consider the illustrative strengths this gives the model to far outweigh the weaknesses. We fully acknowledge that we have not precisely estimated the extent to which variation in wages are driven by tastes for leisure as opposed to talent: instead, we see our contribution as highlighting important obstacles for papers considering redistribution in the context of a lifecycle model.

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 strong functional form assumptions and panel data on 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.

Like all models, our model differs from reality. Unlike some models with a lifecycle human capital framework, we miss 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. 2011) involuntary unemployment and search frictions (Low et al. 2010), and credit constraints (Rossi and Trucchi, 2016) while the cited papers do not. While each show the importance of modeling these components for lifecycle labor supply, we have attempted to mitigate their impact through our choice of subject, age span, and data treatment.

We attempt to 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. We have attempted to filter out important permanent fertility shocks to earnings by focusing on prime age males, whose labor supply is less affected by fertility shocks (Angrist and Evans (1998)). We acknowledge that credit constraints are a potentially important part of the model, but 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. Because our data actually displays strongly rising hours as a function of age, we conclude credit constraints may not be a first-order problem. Finally, we avoid many household decisions by reinterpreting our "taste" coefficient to reflect household production technology, rather than being a "pure" taste coefficient.

3.3 Identification

In order for our model to be identified, some observable behavior must change differently depending on whether or not ϕ or A is changed. For instance, suppose we wish to identify our model using earnings and labor behavior (where measured labor includes human capital accumulation). It must be the case that there exists no change in A (ΔA) and k_0 (Δk_0) such that for some change in ϕ $\Delta\phi$, equations 15 and 16 both hold.

$$\frac{\partial(L_t+s_t)}{\partial\phi}\Delta\phi = \frac{\partial(L_t+s_t)}{\partial A}\Delta A + \frac{\partial(L_t+s_t)}{\partial k_0}\Delta k_0 \quad \text{and} \tag{15}$$

$$\frac{\partial(w_tL_t)}{\partial\phi}\Delta\phi = \frac{\partial(w_tL_t)}{\partial A}\Delta A + \frac{\partial(w_tL_t)}{\partial k_0}\Delta k_0 \quad \forall t\in\{\underline{t},\bar{t}\} \tag{16}$$

Such a coincidence of effects is possible, with the right functional forms. However, we contend that it is actually quite difficult to write down such a model. While we discuss our identification at more length in decomposing our results, it is worth providing an illustrative exercise. Assume a triple $\{A, \phi, k_0\}$ that fits an individual agent's labor supply path and earnings path. Perturb ϕ downwards, and consider how we might counterbalance a reduction in ϕ with changes in A and k_0. In most macro models, a reduction in the distaste for labor primarily effects the level of hours throughout the lifetime by changing the "static" first order condition in equation (9). However, by increasing hours worked, it also increases the value of time invested in human capital (equation 12). Consequentially, we see the total profile of hours on the left hand side of equation 15 shifting up significantly, and as a second-order effect, the slope of hours and earnings also increases slightly. We may hold initial hourly earnings (left hand side of equation 16) constant by keeping initial human capital the same, leaving talent to both reduce the profile of hours and decrease the slope. Fortunately, talent plays a rather small role in determining total hours (via λ) but a large role in determining the slope, and is unable to achieve this goal. By reducing talent to cause the slope to return to its baseline level, hours are left elevated.[3].

More generally, we find that, given common functional form assumptions, talent affects the slope of earnings and hours very strongly, but does not strongly affect its level over the entire lifecycle. In contrast, taste affects the level of labor hours very strongly, but only weakly affects the slope of earnings and hours. Income variation driven by differences in talent would typically see hours and earnings variation increasing by age, as the correlation between hours and earnings decreases. However, income variation driven by differences in taste would typically see little change in either hours variation by age or the correlation between hours and earnings by age.

4 Estimation

4.1 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.[4]

[3]We depict this exercise in more detail in Section 5.1 and Figure 1.

[4]Reported work hours include both labor as well as on-the-job training.

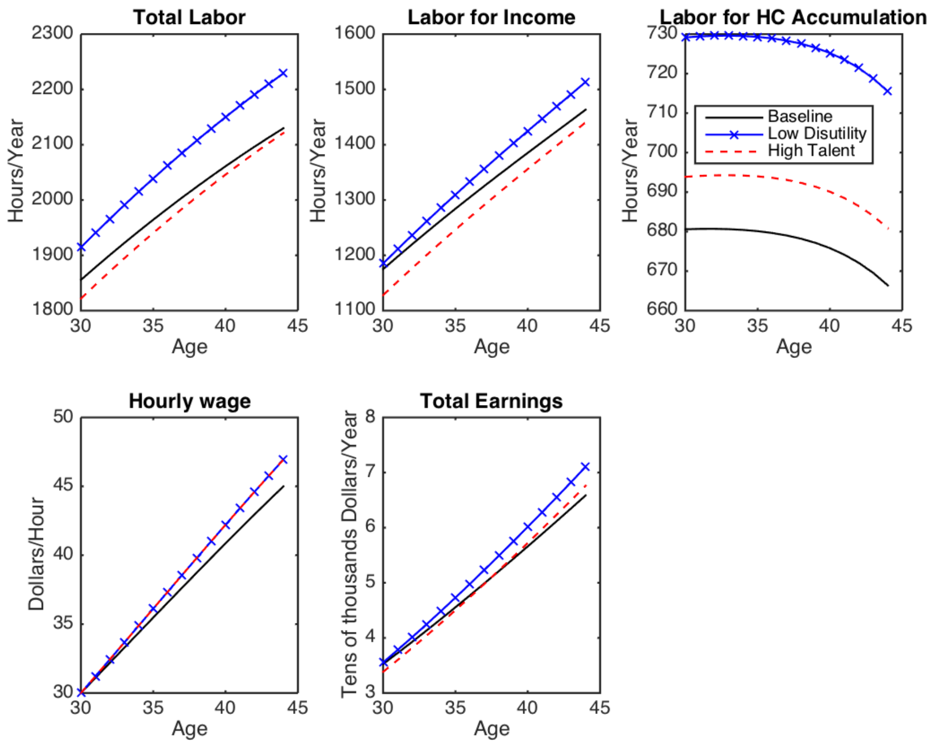

Figure 1: This figure depicts three comparison paths: a baseline path (black lines), a low-distaste for labor path (blue lines), and a high-ability path (red lines). Importantly, while the two counterfactual paths display similar hourly wages, an increase in talent decreases total labor hours while a decrease in distaste increases total labor hours. As a consequence, increasing earnings with higher ability typically lowers the correlation between hours and earnings, increasing earnings with a lower distaste for labor increases the correlation.

4.2 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 King-Plosser-Rebelo preferences (King et al., 1988):

$$U_i(c, \ell, \phi_i) = \frac{c_t^{1-\sigma}}{1-\sigma} - \phi_i \frac{(1-\ell_t)^{1+\eta}}{1+\eta} \tag{17}$$

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

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

$$[\ell_{it}]: \quad \beta^{t-1} \phi_i (1-\ell_{it})^{\eta} = \lambda_i \left(\frac{1}{1+r}\right)^{t-1} w_{it} \tag{19}$$

$$[s_{it}]: \quad \left(\frac{1}{\beta(1+r)}\right)^{t-1} \frac{R(s_{it}k_{it})^{1-\gamma}}{A_i \gamma} = \sum_{\tau=t+1}^{T} \left(\frac{1}{1+r}\right)^{\tau-t} (1-\delta)^{\tau-t-1} R n_{i\tau}. \tag{20}$$

4.3 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 (19) becomes

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

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

where y_{it} is annual earnings.[5] 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.[6]

4.4 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).[7]

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

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

[7] 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 (2010)), Hendricks (2013) models schooling choice closely and finds 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.

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.[8] 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 $Rk_{it} = 1200k_{it}$.[9]

4.5 Estimation With Aggregate Moments

4.5.1 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.

4.5.2 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

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

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

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 weights. In the end, we are left with 90 moments to target. We plot these moments in Figure 2.

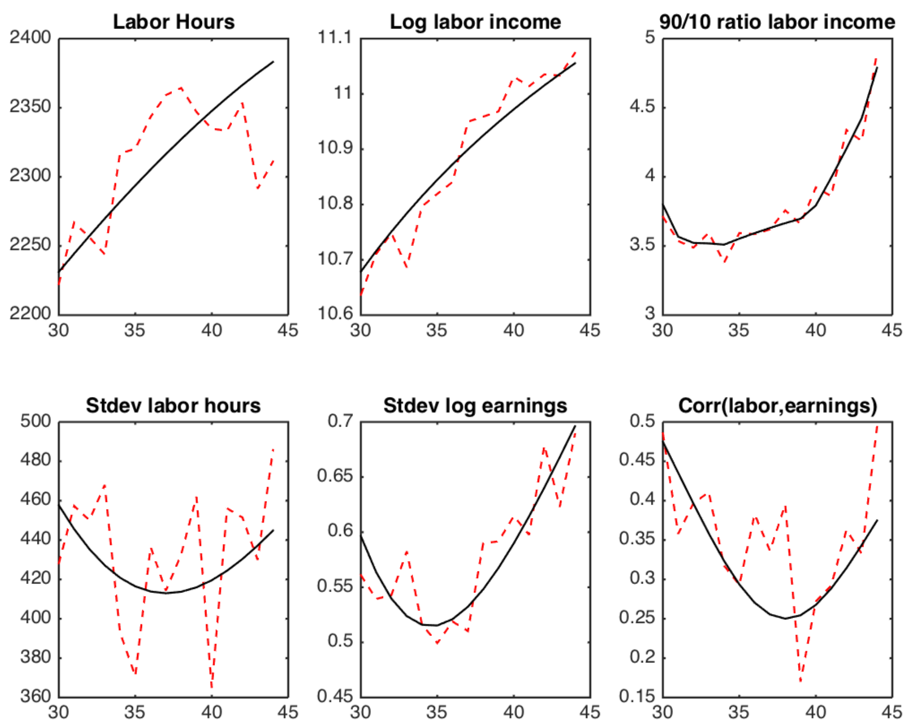

Figure 2: This figure depicts the six sets of empirical moments from the NLSY79 (red line) and model fits (black line) that describe the joint distribution of labor hours and earnings by age.

4.5.3 Fits

Given the calibrated parameter values in Table 1, 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 sets of moments. The moments and fits are depicted in Figure 2. 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 154 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 1: 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 & Kimball 2002)
Discount rate	β	0.945	Gomme and Rupert (2007)
Human capital wage	R	1200	Normalization

Table 1: 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 3. 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 3 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.

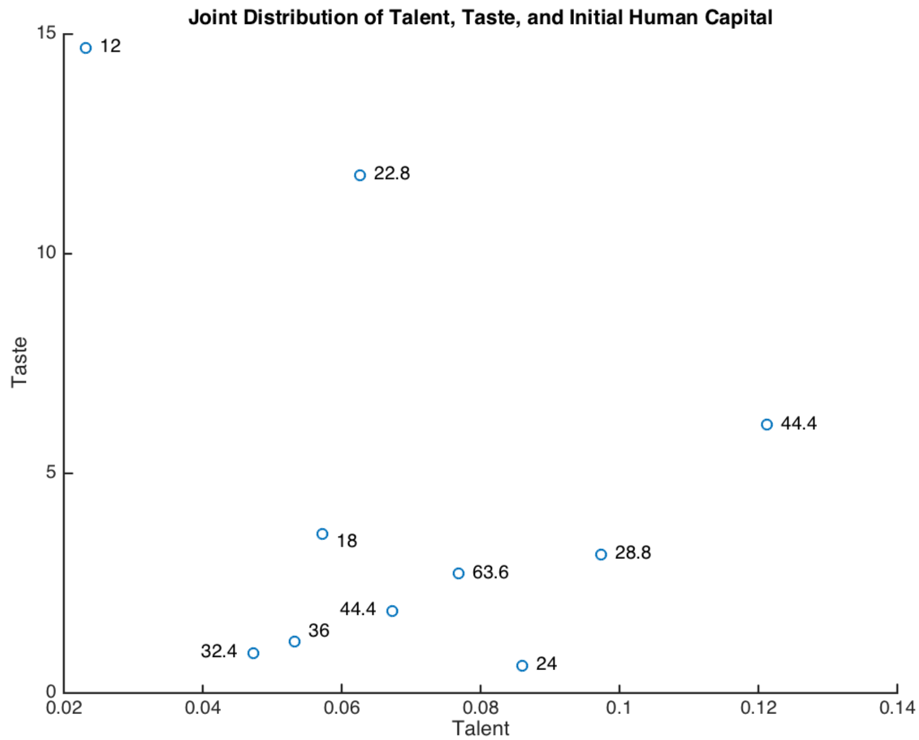

Figure 3: This figure depicts the joint distribution of talent A_i, taste ψ_i, and initial human capital at age 30 (interpreted as an hourly wage) RK_i. Importantly for our results, our model moments put our fit to include people who work many hours (have a low distaste for labor) but are relatively untalented (have a reactively low wage), while there are people with high talent and high initial wage who work relatively little.

5 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:

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

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

$$\Delta S_{E,t} \approx \epsilon_{A,t}\Delta S_A + \epsilon_{\phi,t}\Delta S_\phi + \epsilon_{k,t}\Delta S_k + \kappa \tag{21}$$

Where $\epsilon_{X,t} = \frac{S_X}{S_{E,t}} \frac{\partial S_{E,t}}{\partial S_X}\Big|_{S_X}$, 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 4.[10]

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 71%. In the next section, we discuss 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 21) suggests our linear approximation is good.

5.1 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

[10] 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.

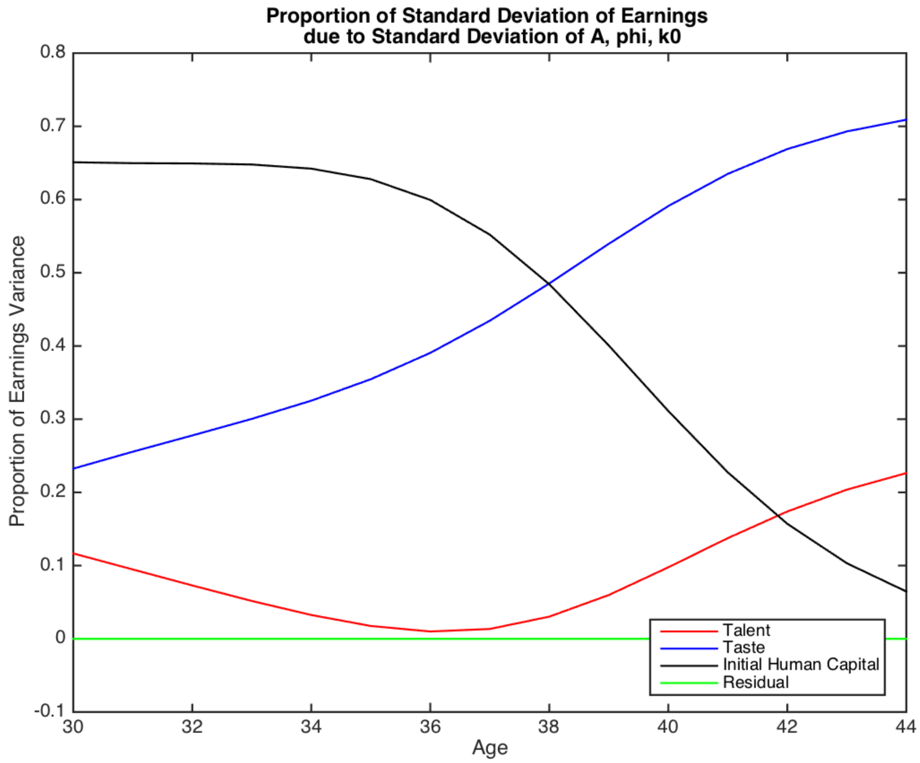

Figure 4: This figure summarizes our main decomposition results, depicting the three elasticities from equation 21. Each of the three visible lines (the residual category, a measure of how bad our linear approximation is, is not visible) indicates how much earnings variance at each age falls if variance in the corresponding parameter falls by 1%. These values are normalized by the total fall in earnings variance at each age.

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 1, denoting human capital as "hourly wage."

Figure 1 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 5 depicts the components of our decomposition in Figure 4 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 34% (from 50%), while if it was 0.2 higher, it would rise to 70%. 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.

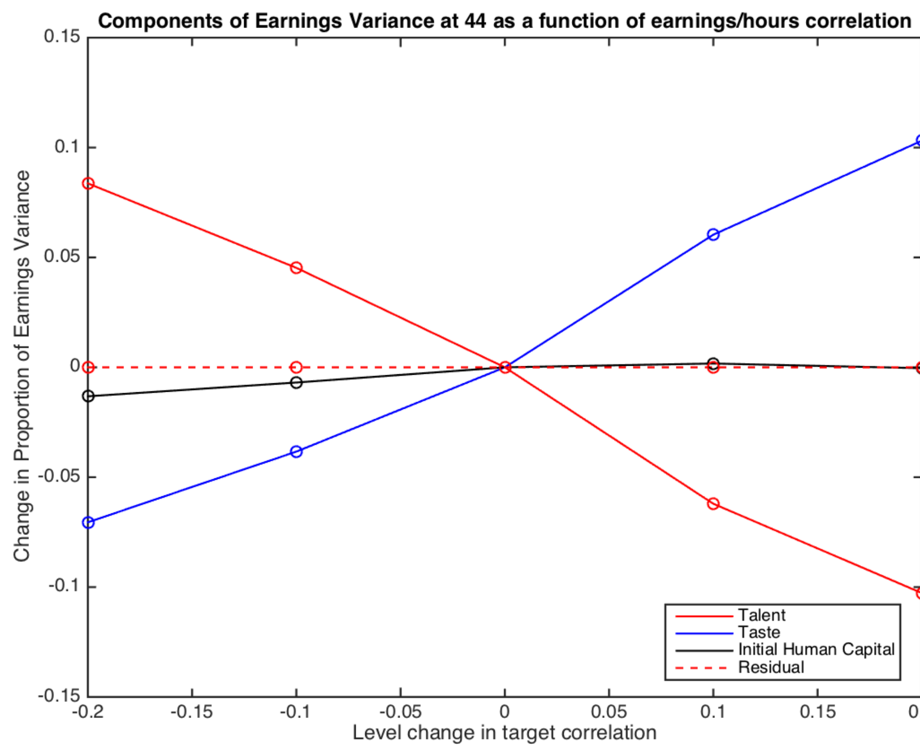

Figure 5: This figure depicts the change in the proportion of earnings variation at age 44 attributable to talent, taste and initial human capital for various level changes in the correlation between earnings and labor hours. Each elasticity is calculated using the elasticities of equation 21. As we 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 1 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 6 depicts a similar exercise as Figure 5, 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.

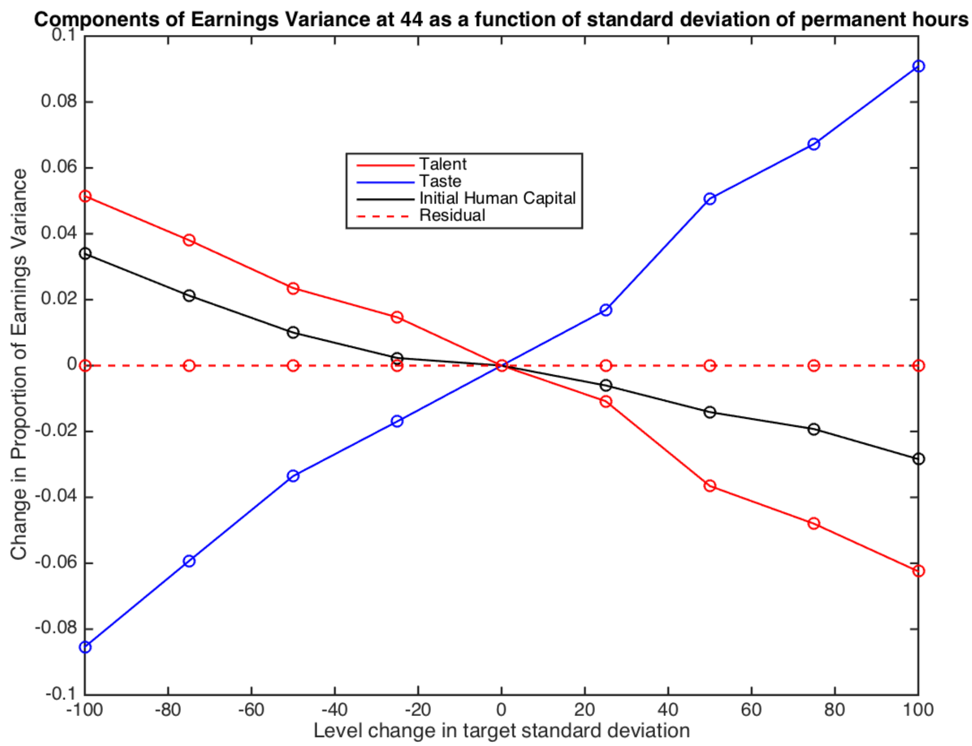

Figure 6: This figure depicts the change in the proportion of earnings variation at age 44 attributable to talent, taste and initial human capital for various level changes in the standard deviation of labor hours. Each elasticity is calculated using the elasticities of equation 21. As we increase the "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.

To confirm our intuition, we examine the results of one final exercise: we fit to the same 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 340 at age 30 and have it fall to 256 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 0.33 and ending at -0.37. Doing so would result in the proportion of income variation at age 44 attributable to taste alone falling from 71% to 16%. While this alternative calibration is not supported by the data, it highlights the data that generates our results.

5.2 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 17) 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

$$U_i(c, \ell, \phi_i) = \zeta_i(\alpha)\frac{c_t^{1-\sigma}}{1-\sigma} - \zeta_i(\alpha)\phi_i\frac{(1-\ell_t)^{1+\eta}}{1+\eta} \tag{22}$$

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 7 depicts our results.

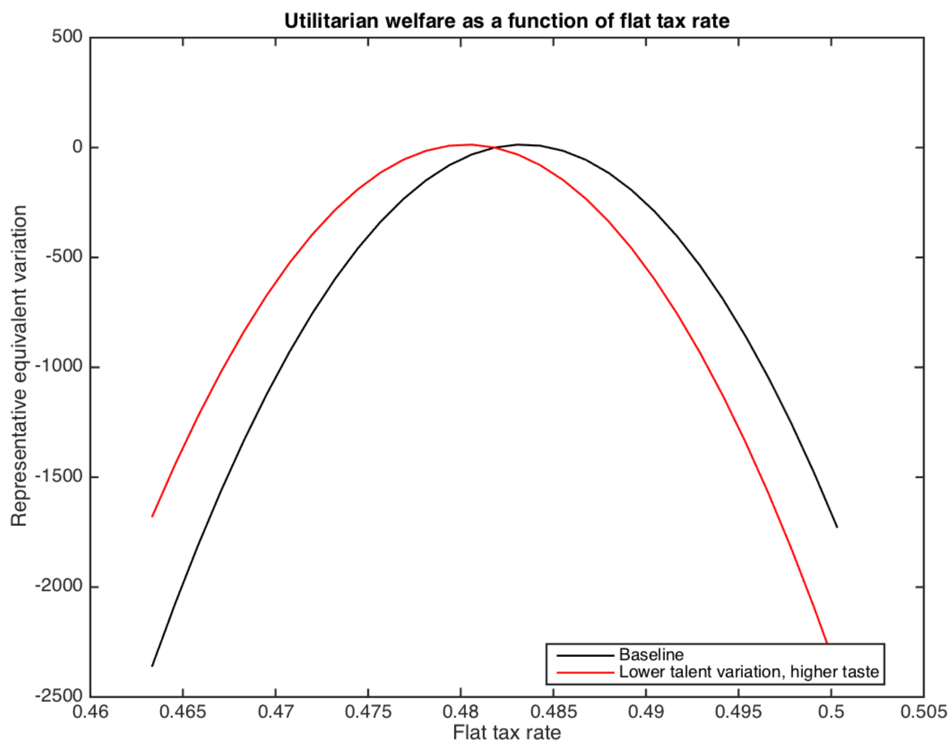

Figure 7: This figure depicts the equivalent variation of changing the tax rate from its optimum in our baseline setup and when we exchange sources of variation from talent to taste, keeping income variation at age 44 constant, with $\alpha = 0.5$ (heterogeneity on both consumption and labor preferences). 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.

We find that in our baseline calibration ($\alpha = 0.5$), the percentage change in optimal tax rate 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%.[11] 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 7.[12]

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 8.

Figure 8 shows that when α is near zero, we attribute all heterogeneity in taste to consumption 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.

6 Conclusion

Papers examining redistribution need to be concerned about the sources of income inequality. When income inequality is generated by taste differences on consumption, even a utilitarian might even desire to redistribute from the poor, with revealed low taste for consumption, to the

[11]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.

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

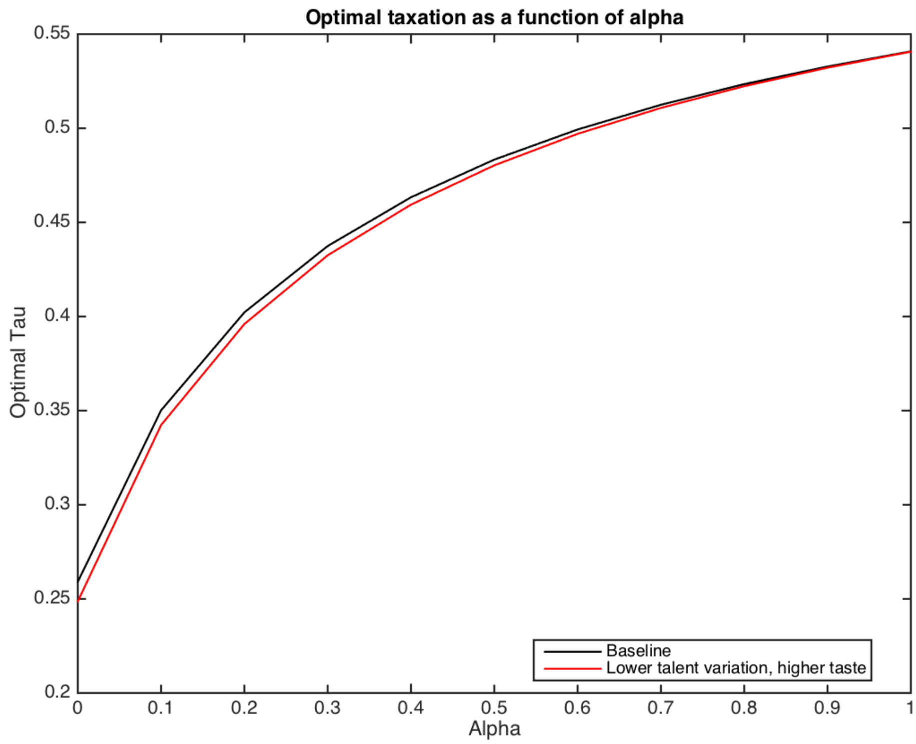

Figure 8: This figure depicts the two optimal tax rates as a function of α, which controls where heterogeneity in preference lies. When $\alpha = 0$, the heterogeneous term multiplies consumption preferences, and all heterogeneity lies on consumption. When $\alpha = 1$, heterogeneity shifts to labor 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.

rich, with revealed high taste. While heterogeneity in the relative preference for consumption vs. leisure is revealed by behavior, where heterogeneity lies (on consumption or leisure preferences) is not. However, this paper establishes that there is substantial taste heterogeneity, and that even small amounts of it on the consumption side curtail the redistributionary motive of a utilitarian social planner.

More broadly, this paper establishes that measurements of talent and taste, while mixed together in behavior, typically have quite different behavioral implications, and are consequentially identified for a typical labor supply model. Having identified substantial heterogeneity in taste in our model, we examine the causes. We find that income variation due primarily to talent would generate strongly declining correlation of hours and earnings, and would be unable to generate the relatively constant and substantial standard deviation of permanent hours displayed in the data. We find that the substantial and relatively nondeclining correlation of hours and earnings and standard deviation of permanent hours drives our results.

While we do not suggest our model is the final word on the sources of inequality, we do believe that these two dynamic targets presents a strong challenge for any model that generates income heterogeneity with “luck” or talent alone. Moreover, we note that even in the presence of preference heterogeneity, the implications for redistributionary motives are not necessarily dramatically threatened, if heterogeneity is on labor preferences, rather than consumption preferences, which individual behavior cannot offer information on.

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.

Mark Aguiar, Erik Hurst, and Loukas Karabarbounis. Recent developments in the economics of time use. *Annual Review of Economics*, 4, 2012. ISSN 1057-9230.

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

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.

Adam Blandin. Learning by doing and ben-porath: Life-cycle predictions and policy implications. Working Papers 1604, VCU School of Business, Department of Economics, 2016. URL `https://EconPapers.repec.org/RePEc:vcu:wpaper:1604`.

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. *The American Economic Review*, 102(7):3483–3515, 2012.

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

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.

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. URL `https://ideas.repec.org/a/oup/restud/v81y2014i2p818-850.html`.

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

James J. Heckman. A Life-Cycle Model of Earnings, Learning, and Consumption. *Journal of Political Economy*, 84(4):S11–S44, 1976a.

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

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.

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.

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.

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.

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.

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

Rodolfo E Manuelli, Ananth Seshadri, and Yongseok Shin. Lifetime Labor Supply and Human Capital Investment. 2012.

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.

Paolo Giovanni Piacquadio. A Fairness Justification of Utilitarianism. *Econometrica*, 85(4): 1261–1276, 2017. ISSN 0012-9682. 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.

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.

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 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.

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 0.25, better fitting the hours data (simulated hours rise by 120 by the age of 44). In order to keep earnings rising, wages rise by more, generated by a higher level of human capital. This lowers the fraction of income variation at age 44 attributable to taste from 71% to 58%.

Appendix B: Flexible Frisch Elasticity of Labor Supply

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 hours of work (including on-the-job training). 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 means of earnings and hours, but it does affect the variances and the covariance.

We estimate the variances and covariance of log annual earnings and annual hours of work as follows.[13] 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

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

and store the fitted values at each age. Then we calculate the 90-10 ratio of the (exponentiated) fitted values at each age.