Disciplining the Human Capital Model: Learning by Doing, Ben-Porath, and Policy Analysis

Adam Blandin*

Arizona State University

January, 2015

Abstract

The human capital literature is largely split between two models of human capital investment: Learning By Doing (LBD) and Ben-Porath (BP). Given the importance of human capital investment for a host of policy issues, I ask whether observable macroeconomic moments are informative about the relative importance of LBD investment versus BP investment. A life-cycle human capital model is constructed which nests both LBD and BP as extreme special cases.

I find: (1) Both the BP and LBD versions of the model are consistent with the aggregate distribution of earnings, hourly wages, and hours worked for men in the PSID. (2) Conditional on matching these aggregate *levels* facts, the BP version of the model is more consistent with the variance in earnings *growth rates* in the data. (3) Policies which decrease the return to human capital investment, such as a progressive earnings tax, decrease aggregate human capital investment and earnings substantially more in a BP world than in a LBD world. Taken together my findings suggest that within a plausibly parametrized model of human capital accumulation, government policies which reduce the return to human capital investment will generate large decreases in human capital investment and earnings.

*Contact Info: Department of Economics, Arizona State University, P.O. Box 879801, Tempe AZ 85287-9801. *Email*: ajblandi@asu.edu

1 Introduction

A large research agenda has studied the implications of government policies for environments with heterogeneity in the earnings profiles of individuals. Recognizing that earnings may react to changes in policy, many works have endogenized earnings within a human capital framework. However, once formal education is over with there is no consensus among economists as to how we should model human capital accumulation. This is unsettling because recent work, including Cossa, Heckman, and Lochner (2003); Hansen and Imrohoroglu (2009); and Peterman (2012), have demonstrated that policy implications can significantly depend on the assumed form of the human capital process.

Presently there are two human capital accumulation processes which dominate the literature: Learning By Doing (LBD) and Ben-Porath (BP). Each emphasizes a distinct form of human capital invesment. LBD assumes time spent producing a good for an employer is the sole source of human capital investment. I label this type of activity "production". BP, on the other hand, assumes that production and human capital investment are mutually exclusive activities. To accumulate human capital, then, BP requires that individuals spend time "studying". Studying may take place on or off the job; the defining feature is that while studying the individual is not producing a good on behalf of her employer, and so is not compensated for this time. A fundamental difference between the BP and LBD processes, then, is that in the BP case there is a tradeoff between earnings today and earning potential (human capital) tomorrow, while no such tradeoff exists in the LBD case[1].

This paper investigates the aggregate implications of a class of human capital processes which nest BP and LBD as extreme special cases. Specifically, I ask two questions. First, are observable macroeconomic moments informative about the relative importance of studying versus production? Second, for which government policies does the relative importance of studying versus production play a key role?

[1] Previous works have made points very similar to this; for example see Cossa, et al. (2003) and Wallenius (2011).

To answer these questions I construct a life-cycle model featuring a human capital process which weights two inputs to human capital formation: production time, during which the individual is compensated, and studying, which is uncompensated. When production receives full weight the process corresponds to a LBD view; when full weight is placed on studying the process corresponds to a BP view. Individuals in my model are born as young adults endowed with an initial human capital stock and a learning ability, which determines their proficiency at accumulating future human capital; heterogeneity in these initial endowments is the only way individuals differ from one another. Given their endowments, individuals maximize lifetime utility through decisions about study, production, leisure, and consumption in each period.

My findings are as follows. First, both the BP and LBD versions of the model are consistent with the aggregate distribution of life-cycle hours worked, annual earnings, and hourly wages for men in the PSID (1968-2009). This helps explain why so little is known about the relative importance of these two forms of investment. Second, conditional on matching these aggregate *levels* facts, the BP version of the model is more consistent with the variance of earnings and wage *growth rates* in the data. Finally, policies which affect the return to human capital investment have larger impacts on human capital investment and earnings in the BP model relative to the LBD model. Taken together my findings suggest that within a plausibly parametrized model of human capital accumulation, government policies which reduce the return to human capital investment, for example a progressive income tax, will generate large decreases in human capital investment and earnings.

I now briefly discuss the intuition behind my third finding. An increase in the return to human capital investment leads to increased investment in both models. However, the investment response is relatively small in the LBD model because individuals can only increase investment (production) by reducing leisure. In the BP model, by contrast, individuals can increase investment without reducing leisure, by allocating a greater fraction of non-leisure time to studying. When I analyze the impact of a revenue neutral change from a progressive income tax to a proportional income tax, I find that mean human capital investment increases 12% in the BP model relative to 2% in the LBD model. This leads to a 3.6% increase in aggregate earnings the BP model relative to a 2.3% increase in the LBD model. I argue that the size of these effects should be viewed as a distant lower bound on the magnitude of the effects from a more detailed policy analysis.

The tradeoff between production and investment in the BP model also drives my second finding. In the data the variance in earnings growth rates among 50 year olds is a third lower than the variance among 25 year olds. The BP model endogenously predicts this pattern: early in the life-cycle earnings grow at different rates due to heterogeneity in learning abilities, but late in the life-cycle earnings growth concentrates near zero because all individuals spend less time studying. In the LBD model, by contrast, investment is equivalent to production, which is roughly flat from ages 25 to 50. When the variance in earnings *growth* among young individuals in the LBD model is as large as in the BP model, then later in the life-cycle the variance in earnings *levels* in the LBD model is much larger than in the BP model, because investment continues throughout the life-cycle in the LBD model.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 presents the human capital model and compares the LBD and BP versions of the model. Section 4 documents the aggregate facts for US workers which I use to discipline my human capital model. In Section 5, alternately assuming a LBD and BP model, I calibrate the model parameters to match facts from Section 4. Section 6 then compares the LBD and BP model predictions for the variance of wage and earnings growth rates over the life-cycle. In Section 7 I analyze the effects of a change in the progressivity of the US federal income tax for both calibrated models, and Section 8 concludes.

2 Relevant Literature

Current human capital theory is founded largely on works by Mincer (1958), Schulz (1961), Becker (1964, 1965), Ben-Porath (1967), and Heckman (1976a,b). The theory of human capital that arose from this body of work posited that individuals own a stock of knowledge and abilities which augment the productivity of their labor. The size and composition of this stock at a given date is determined primarily by an initial endowment, and investments made prior to that date. One of the most important forms of investment is through the allocation of the time endowment.

As Schulz (1961) and others have pointed out, time investments can be difficult to measure because many activities are not pure investments. That is, if we think of investment in human capital as some fraction I of the time endowment, then often I does not only appear in the human capital production function. Instead, I may also show up in the individual's leisure preferences (e.g. an American History teacher who takes a vacation to Gettysburg) or in production time (e.g. a surgeon whose skill improves with each surgery). With regard to the coincidence of investment and production, the vast majority of previous studies have simply assumed one of two extremes: LBD, which assumes that investment is equivalent to production; and BP, which assumes investment and production are mutually exclusive.

One notable departure from this pattern is Killingsworth (1982), which is the only other study I have found that analyzes a model of human capital accumulation with both production and study as inputs. I build on Killingsworth's contribution in two ways. First, several questionable assumptions made for theoretical tractability are avoided; instead I opt for a more flexible model that is solved computationally. Second, I exploit these computational solutions to answer empirical questions which Killingsworth is unable to address.

Several recent papers which were cited in the introduction examine how policy implications vary depending on whether one assumes a BP or LBD process. This paper goes a step further by showing that observable macroeconomic moments are informative about the relative importance of these two forms of investment. To the extent that these observations can discipline the human capital process, they will allow for more precise policy analysis.

3 The Human Capital Model

This section lays out the human capital accumulation model. I first discuss the economic environment and the lifetime utility maximization problem of an individual in Section 3.1. In section 3.2 I highlight the human capital production function, paying special attention to the central model parameter: the relative weight of studying versus production. Section 3.3 demonstrates that the impact of an important class of government policies depends on this parameter.

3.1 Environment and the Individual's Utility Maximization Problem

Time is discrete and a model period corresponds to approximately 5 years. The model contains a unit mass of finite-lived individuals. In the first period of life each individual is endowed with initial human capital $\bar{h}_1$, as well as an ability a which determines their proficiency in the formation of new human capital. The distribution of endowment $(\bar{h}_1, a)$ is given by $\Lambda(\bar{h}_1, a)$; this is the only source of heterogeneity in the model.

Individuals are also endowed with a unit of time in each period. This time is split between three activities: leisure, $1-n$, production, $n(1-s)$, and studying, ns. The theoretical distinction between studying and production is that the individual is only paid a wage for time spent in production; time spent studying is only useful to the extent that it contributes to future human capital.

There is a single consumption good, c, whose price is normalized to one in each period. The individual has time-separable preferences over consumption and leisure. Utility from leisure only depends on n, and not the fraction of n spent studying or producing. The individual's problem is to maximize the discounted value of lifetime utility subject to constraints on spending and human capital:

$$V(\bar{h}_1, a) = \max_{\{n_j, s_j, h_{j+1}, c_j\}_{j=1}^{\bar{J}}} \sum_{j=1}^{\bar{J}} \beta^{j-1} u_j(c_j \ , \ 1 - n_j) \tag{1}$$

$$s.t. \quad \sum_{j=1}^{\bar{J}} R^{1-j} c_j \leq \sum_{j=1}^{J} R^{1-j} w h_j n_j (1 - s_j) \ ; \tag{2}$$

$$h_{j+1} = (1 - \delta) h_j + H(h_j, n_j, s_j; a) \ ; \tag{3}$$

$$n_j \in [0,1]; \quad n_j = 0 \ \forall j > J; \quad s_j \in [0,1]; \quad h_1 = \bar{h}_1. \tag{4}$$

The first constraint states that the discounted sum of lifetime consumption spending cannot exceed the discounted sum of lifetime labor earnings. Labor earnings in each period are given by $w_j h_j n_j (1 - s_j)$, where w_j is the period j wage rate; h_j is the individual's human capital in period j; and $n_j(1 - s_j)$ is the fraction of the time endowment spent in production in period j. $J \leq \bar{J}$ is a mandatory retirement date, *i.e.* individuals must set hours worked equal to zero beyond this date. From period $J + 1$ to $\bar{J}$, the individual simply eats the portion of his lifetime income which has not already been consumed.

The second set of constraints describes the human capital accumulation process. Human capital in period $j + 1$ is the sum of undepreciated human capital from the beginning of period j and human capital produced in period j. δ is the depreciation rate of existing human capital, and H is the human capital production function. The amount of newly produced human capital is a function of the individual's production and studying choices in period j, as well as his existing human capital and his learning ability.

3.2 The Human Capital Production Function

The human capital production function H is given by the CES function

$$H(h,n,s) = ah^{\phi_h}[\ \alpha(ns)^{\rho} \ + \ (1-\alpha)(n(1-s))^{\rho}\ ]^{\phi_n/\rho} \ . \tag{5}$$

ϕ_h and ϕ_n determine the curvature of new human capital with respect to current human capital and time investment, respectively. a is the individual's learning ability: individuals with higher ability produce more human capital for a given vector (h,n,s) than individuals with lower ability. $\rho < 1$ is the elasticity of substitution between studying and production.

The central model parameter is $\alpha \in [0,1]$, the CES weight for studying. It is instructive to first examine the extreme cases $\alpha = 0$ and $\alpha = 1$. When $\alpha = 0$, (5) reduces to

$$H(h,n,s) = ah^{\phi_h}[n(1-s)]^{\phi_n} \ . \tag{6}$$

This represents a pure LBD human capital accumulation process because the only input to human capital investment is production, $n(1-s)$ [2]. Alternatively, when $\alpha = 1$, (5) becomes

$$H(h,n,s) = ah^{\phi_h}[ns]^{\phi_n} \ , \tag{7}$$

which assumes that studying is the sole investment channel for human capital. In the special case $\phi_h = \phi_n$ this is identical to the original Ben-Porath (1967) specification[3].

In general, for $\alpha \in (0,1)$, H is increasing in s over the range

$$\left[0\ ,\ \frac{1}{\left(\frac{1-\alpha}{\alpha}\right)^{\frac{1}{1-\rho}} + 1}\right) \ , \tag{8}$$

where the right hand side is increasing in α since $\rho < 1$.

[2]For examples of LBD functions previously considered in the literature see Shaw (1989); Chang, Gomes, and Schorfheide (2002); Imai, Keane (2004); or Gemici, Wiswall (2014).

[3]Papers using this functional form include Heckman, Lochner, and Taber (1998); and Huggett, Ventura, and Yaron (2011).

The two questions this paper asks can now be stated more concretely. First, for which government policies does α play a key role? Section 3.3 briefly illustrates that the impact of policies which affect the returns to human capital investment depends strongly on α. Section 7 quantitatively explores the impact of α in a more detailed policy analysis.

Second, given that policy outcomes may depend on α, are observable macroeconomic moments informative about the value of α? To answer this question, first note that one could easily put a lower bound on α if studying were observable and ρ was known. For example, suppose we knew $\rho \to 0$, in which case (8) becomes $[0, \alpha)$. Since studying is only useful to the extent that it produces more human capital, all model individuals will choose a value of s within the interval $[0, \alpha)$. Then if many individuals in the data spent at least a fraction $\underline{s}$ of their non-leisure time studying, it would follow that $\alpha \geq \underline{s}$.

However, studying is difficult to observe. For one thing, large time use surveys (e.g. the American Time Use Survey) typically do not ask detailed questions about specific activities while a respondent is at work, which is where a significant portion of studying is likely to occur for adults. For another, even if activities were closely tracked it is difficult to neatly place them into a "studying" or "production" or "leisure" cateogry. In light of these limitations I refrain from attempting to measure studying directly, and instead ask which values of α are consistent with macroeconomic facts from the US. That analysis is conducted in Sections 4-6. I now briefly motivate this investigation in Section 3.3.

3.3 Impact of a Change in the Return to Future Human Capital

This section demonstrates that the impact of a change in the return to future human capital can depend on the value of α. To clarify the analysis I impose several convenient assumptions for this section only. The simplified problem to be solved is

$$\max_{c_1,c_2,s} \quad u_1(c_1) + u_2(c_2)$$
$$s.t. \quad c_1 + c_2 \leq h_1(1-s) + (1-\tau)h_2 \ ;$$
$$h_2 = H(s;\alpha) = ah_1[\ \alpha(s)^\rho \ + \ (1-\alpha)(1-s)^\rho\ ]^{\phi_n/\rho} \ .$$

Note that I have assumed away the preference for leisure and human capital depreciation. I have also set $\phi_h = \beta = 1/R = w = 1$ and $J = \bar{J} = 2$. The tax rate τ determines the return to a unit of human capital in period 2 relative to period 1. It is without loss of generality that agents only make a studying decision in period 1, since individuals never study in the final period of the life-cycle. With no leisure choice the problem can be solved in two steps: first by choosing s to maximize the individual's lifetime earnings, second by deciding the optimal division of that income into consumption over the two periods. The second step is unaffected by α, so the analysis in this section focuses on the choice of s.

I solve the above problem for $\tau \in [0,1]$ and $\alpha \in \{0, .5, 1\}$. For each α a pair of parameters need to be set: $(h_1(\alpha), a(\alpha))$. The values are chosen so that when $\tau = 0$, income in the first period is $h_1(1-s) = 1$ and income in the second period is $(1-\tau)h_2 = h_2 = 3$. Solving this set of problems allows me to compare the endogenous response to an increase in τ for different α values which are observationally equivalent.

Figure 1(a) plots the investment decisions of individuals as a function of $(1-\tau)$ for the Cobb-Douglass case $\rho \to 0$. Investment is given by $s^\alpha(1-s)^{1-\alpha}$. Each profile is normalized to one at the value $(1-\tau) = 1$. For all values of α, a fall in $(1-\tau)$ reduces the return to human capital investment. However, there is no investment response to a change in $(1-\tau)$ in the LBD ($\alpha = 0$) case because the individual can invest in future human capital without forgoing any income today. When $\alpha > 0$, by contrast, lower values of $(1-\tau)$ generate endogenous reductions in investment, and the size of the reduction is increasing in the value

Figure 1: Impact of τ on Model Outcomes

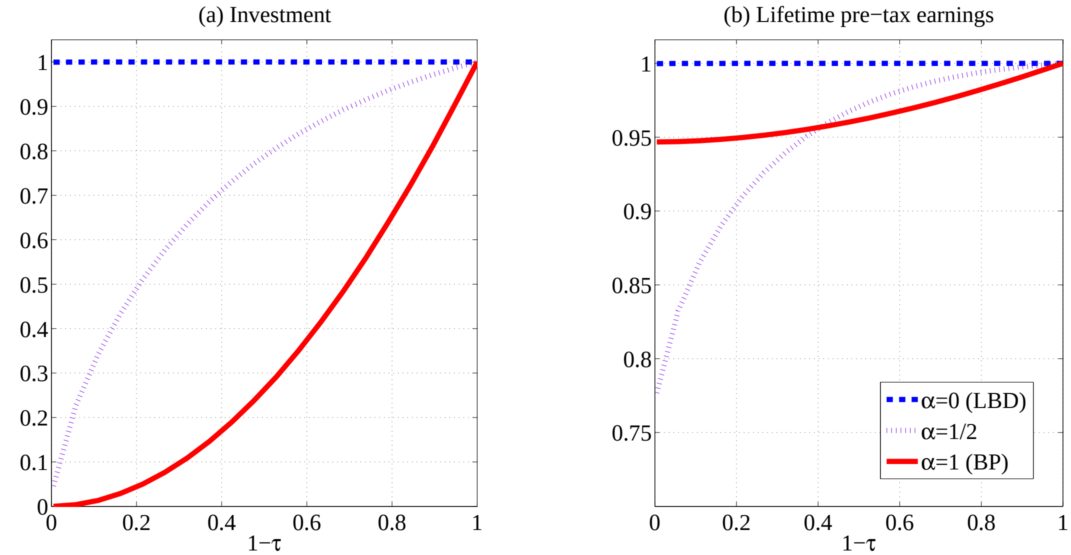

of α.

Figure 1(b) plots lifetime pre-tax earnings, $h_1(1-s)+h_2$, as a function of $(1-\tau)$. Each profile is normalized to one at the value $(1-\tau)$. For values of $(1-\tau)$ near one lifetime earnings are decreasing in α, and as with Figure 1(a) the LBD $(\alpha=0)$ profile is flat. Somewhat surprisingly, for values of $(1-\tau)$ near zero lifetime earnings are increasing in α.[4]

The main takeaway from this exercise is that the impact of government policies can strongly depend on the relative importance of studying and production. In particular, a BP and LBD model calibrated to match the same income facts may predict different policy outcomes. Therefore, empirical evidence of the relative importance of studying and production can reduce the ambiguity of policy analysis conducted within a human capital framework. A detailed search for such evidence is now conducted in Sections 4-6.

[4]This is simply because $h_1(\alpha)$ is increasing in α, and when $(1-\tau)\searrow 0$, $h_1(\alpha)(1-s(\alpha))+h_2(\alpha)\searrow 2h_1(\alpha)$. In words, for very low values of $(1-\tau)$ all individuals supply $h_1(\alpha)$ in both periods, because the returns to investment are trivial. Of course, this result is dependent on $\rho=0$. For larger values of ρ the size of the reduction is increasing in α for all values of $(1-\tau)$.

4 Aggregate Life-cycle Facts for US Workers

This section establishes the basic life-cycle facts for US workers which I use to discipline my human capital model. Section 4.1 summarizes the distribution of annual earnings, hourly wages, and hours worked over the life-cycle. I use these moments to set model parameters for both the LBD and BP versions of my model in Section 5. Section 4.2 documents a novel fact: the variance in wage and earnings growth among individuals of a given age falls 25-50% over the life-cycle. In Section 6 I show that the BP model can account for this entire decline, while the LBD model accounts for less than half.

I obtain data on earnings and hours from the 1968-2009 family-level files of the Panel Study of Income Dynamics (PSID). I restrict my sample to observations of male heads of households between 23-60 years of age who report working at least 500 hours in a given year[5] and whose labor earnings were between \$5,000 and \$1 million in 2009 dollars[6]. These qualifications result in a sample of $11,145$ individuals and $80,427$ year-individual observations.

I exclude males over 60 for two reasons. First, the PSID contains few observations of working men over 60. Second, decisions about wages and hours at this stage in life may be exceptionally influenced by factors which are not present in this model.

I restrict ages on the low end in an attempt to focus on "working age" human capital accumulation, as opposed to human capital accumulated during primarily "school age" years. Two comments are necessary here. First, this restriction does not imply that I ignore early-life human capital accumulation. Rather, I simply capture the *result* of this early accumulation in a reduced form manner, through the initial distribution of human capital and ability endowments. In fact, since these initial endowments are the only source of heterogeneity in

[5]Working 500 hours in a year averages to slightly less than 10 hours per week.

[6]The sources of labor income used were "wages and salary", "commission", "overtime", "professional practice income", "tips", "additional income", and "all other labor income". The minimum hours and earnings thresholds are designed to exclude individuals who were full time students or who experienced a long term unemployment spell. I throw out extremely high earners because these few outliers occasionally alter the measured moments.

my model, early human capital accumulation plays a crucial role in my analysis. Second, we know that young working age individuals work very little before their formal schooling is complete, and that individuals who attend school longer earn higher hourly wages in their first few years of employment. A LBD model of early human capital accumulation would struggle to generate these basic facts, while a simple BP process would do so easily. In contrast to this trivial analysis, we know relatively little about which human capital theories are consistent with hours and earnings facts after formal education has been completed.

I now discuss how the age profiles plotted in this section are obtained. The human capital model laid out in Section 3 is not able to replicate every aspect of time variation contained in the data, so I use a statistical model to extract the sources of variation my model model is designed to address. Following the approach of Hugget, Ventura, and Yaron (2011), I first group the data into 5-year centered age bins. I then assume that the statistics of interest are governed by the following fixed effects model:

$$stat_{a,t,c} = x_a + x_t + x_c + \epsilon_{a,t,c} \tag{9}$$

where $x_a,\ x_t,\ x_c$ denote age, year, and birth cohort fixed effects, respectively, and $\epsilon_{a,t,c}$ is an error term. Since $c \equiv t - a$, there is a well-known multicollinearity problem involved with estimating this model. In light of this, the standard approach is to estimate age effects in one of two ways, alternately assuming $x_t = 0\ \forall t$ (a "cohort-effects" view) and $x_c = 0\ \forall c$ (a "year-effects" view). This draft only analyzes the the year-effects results.

4.1 Life-cycle Profiles Used to Discipline the LBD and BP Models

Figures 2 and 3 depict the data moments used to calibrate the BP and LBD versions of the human capital model. Figure 2 plots the age effects for mean (a) annual earnings, (b) hourly wages, and (c) hours worked. Annual earnings double between 23 and 48 years of age, then decline 10% by age 60. Hourly wages double by age 50 and remain near this peak through age 60. Mean hours display a much flatter curve: they grow about 10% from age 23 to age

Figure 2: Life-cycle Means

35, then gradually fall back to their original level by age 60. In absolute terms mean annual hours start at 2,011 for 23 year olds and peak at 2,200 hours.

The age effects for the log variance of annual earnings are displayed in Figure 3. The variance in earnings among 23 year olds is approximately 45% of mean earnings, and increases to 65% of mean earnings by age 60. An important feature of this profile is that it is positively sloped at almost every age, which says that earnings are becoming more spread out over the life-cycle.

Figure 3: Life-cycle Variance of log Earnings

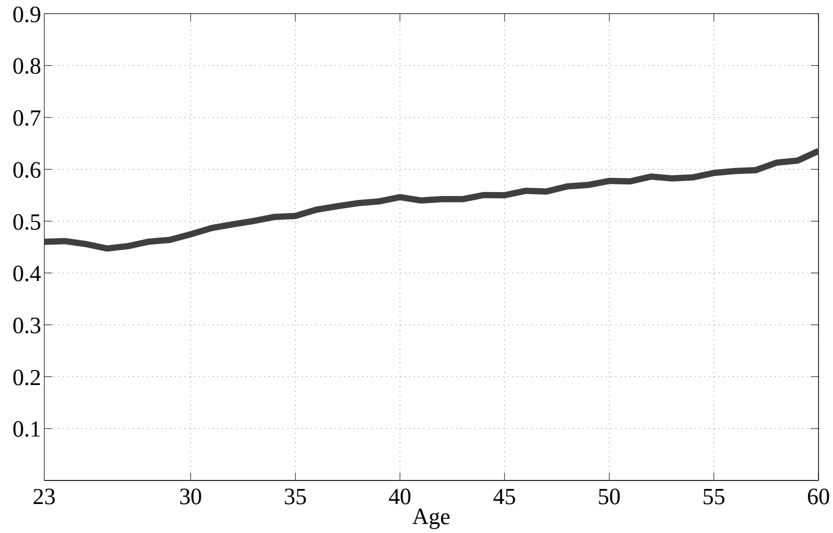

4.2 Life-cycle Profiles Used to Discipline α

Figure 4: Life-cycle Variance of log Earnings

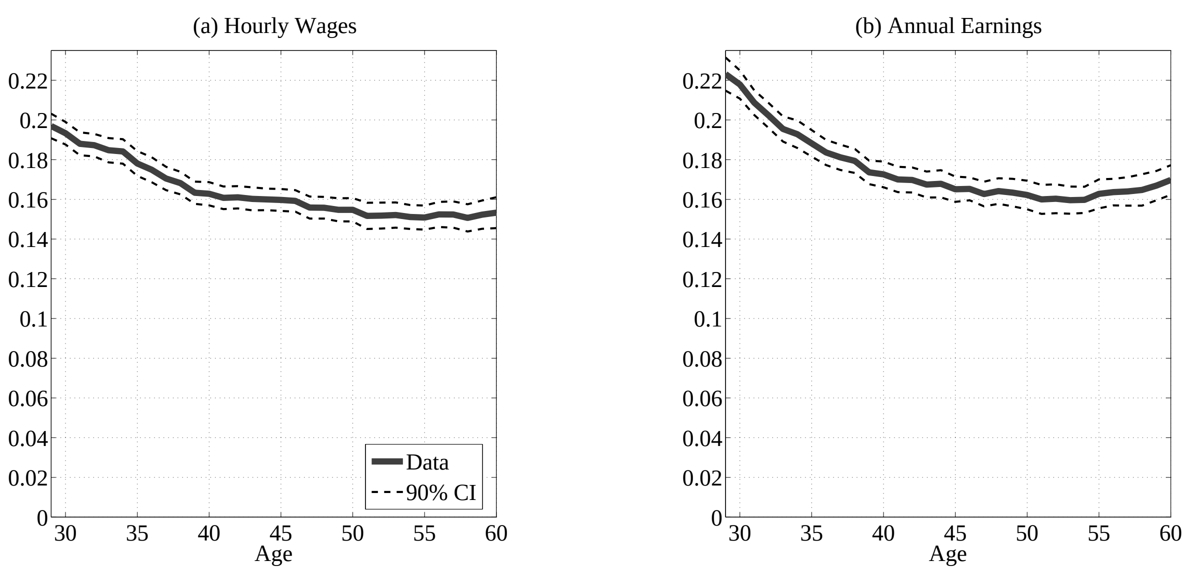

This section presents a novel fact: the variance of wage and earnings growth rates among individuals of a given age falls substantially over the life-cycle. In my human capital model a period will correspond to approximately 5 years, so I examine growth rates over 5 years in the data. The computation of the variance in wage growth merits explanation. To compute the wage growth statistic for age a and year t, I first group all individuals who are a years old in year t and who reported valid hourly wages at both year $t-5$ and year t. I compute the *log* difference of wages between years t and $t-5$ for each of these individuals, then find the variance in these *log* differences across all individuals in the group (a, t). The statistic for earnings growth is computed analagously. Age effects for this statistic are then estimated using (9). Figure 4 establishes that the variance of 5-year wage and earnings growth rates fall by 25-30% over the lifecycle[7]. Roughly half of this decline occurs after age 35.

Figures A1-A5 demonstrate that this finding is robust to several sample restrictions and controls. First, Figure A1 shows that there is a modest change in the level of the profiles, but no substantive change in the life-cycle decline of the profiles, when the minimum hours threshold is increased from 500 hours to 1,800 hours (equivalent to working 35 hours per week for 50 weeks). Therefore, the relatively high variance in growth rates early in the life-cycle is not due to some young workers earning higher wages as they make the transition from part time to full time employment, a phenomenon which does not occur in my model.

Some of the variance in wage growth could be due to differences in the labor markets faced by different education groups. Since my model has only one type of human capital and one labor market, it cannot address heterogeneity in wage growth which arises because of differences in labor markets. However, Figure A2 reveals that when I recompute the statistic within different education groups the initial level and life-cycle decline of the profiles remain broadly consistent with their values in Figure 4 [8]. Additionally, Figure A3 demonstrates that the the same qualitative pattern holds when I estimate potential experience effects within education groups instead of age effects.

[7] To remove the effect of outliers, I throw out any observations with growth rates less that 1/3, and greater than 3. I follow the same procedure for the model in Section 6 when I compare the model predictions to the data. The findings in that section are robust to changes in these cutoffs.

[8] The profile for individuals without a high school degree is somewhat erratic due to a low sample size.

Another possibility is that the high variance in wage growth for young adults could be due to more frequent job losses or changes in occupations relative to older adults, which often entail substantial declines in wages and earnings in the short run. To remove the component of growth variance which is due to large wage declines I throw out all observations in which wages or earnings fell by more than 10%, and plot the resulting age effects in Figure A4. The initial levels of the profiles are substantially lower than the initial levels in Figure 4, but the size of the declines over the life-cycle are similar: the variance in hourly wage growth falls by .034 in Figure A4(a) versus .046 in Figure 4(a); the variance in earnings growth falls by .044 in Figure A4(b) versus .063 in Figure 4(b). This establishes that the decline in the variance of wage and earnings growth is not driven by a reduction in the frequency of large wage/earning declines.

A final concern is that the life-cycle decline in these profiles is due to a decline in the variance of temporary idiosyncratic shocks to earnings, which are not present in my (deterministic) model. This suggests it may be important to distinguish between temporary wage growth and persistent wage growth. To remove purely temporary changes I smooth the wages and earnings of all individuals with a 3 year Moving Average process, then recompute the growth rates using these smoothed series. I also exclude observations in which earnings or wages decline by more than 10%, as was discussed in the preceding paragraph. The age effects obtained from equation for the variance in smoothed growth rates are plotted in Figure A5. The initial level of the profiles are roughly half the initial levels from Figure 4, but the life-cycle decline in the profiles are similar. The profile for hourly wages declines by .033 in Figure A5(a) compared to .046 in Figure 4(a); the profile for annual earnings declines by .045 in Figure A5(b) versus .063 in Figure 4(b).

The key features of the data in Sections 4.1-4.2 can be summarized as follows:

1. Mean hours worked are relatively flat over the life-cycle, growing only 10% from their initial level.

2. Mean annual earnings and hourly wages display significant curvature, doubling from age 23-50, then declining slightly until age 60.

3. The variance of *log* earnings increases steadily over the life-cycle.

4. The variance of 5-year wage and earnings growth rates falls by 25-50% over the life-cycle, with half of this decline occurring after age 35.

The next sections investigate the extent to which the human capital model from Section 3 is consistent with these features.

5 Setting Remaining Model Parameters for a Given Value of α

This section demonstrates that, given a value of α, the remaining model parameters can be calibrated to be consistent with features 1-3 of the data. Section 6 then compares the predictions of the calibrated α-models to feature 4. In Section 5.1 I discuss parameters which are exogenously determined and identical for all α values; Section 5.2 discusses parameters which are set conditional on α so that model moments match the corresponding moments in the data. The calibration procedure in Section 5.2 is the same for all assumed values of α.

To compare hours worked and hourly wages in the model and the data, I need to take a stand on how model activities map into hours worked in the data. I assume reported hours worked are the sum of time in production and time spent studying. This is equivalent to assuming that all studying takes place on the job. While this is a strong assumption, it seems plausible that the majority of human capital investment takes place on the job after age 23. Additionally, in the calibrated BP model the mean fraction of the time endowment spent studying is below 5% for the majority of the life-cycle, suggesting that aggregate measures of hours worked and hourly wages will not vary much with the treatment of studying. Finally, the findings remain largely unchanged when I rerun the calibration assuming that one third of studying is not included in hours worked.

5.1 Parameters set independently of α

Solving my model for a large number of individuals requires considerable computational time, so I limit the number of working periods in my model to $J = 7$. In the data I observe individuals over a 38-year working life, from ages 23-60, so each model period represents $38/J \approx 5$ years. To allow comparisons between my model and the data I partition the 38 years of data into J periods, then average the yearly statistics within each period bin[9]. I set $\bar{J} = 10$, so that individuals work for seven periods then live off their accumulated assets in retirement for three. Since a period corresponds to roughly 5 years this implies that retirement roughly corresponds to ages 61-75. The period real interest rate is set to $R = 1.24$ which implies an annual interest rate of 4%, and the individual discount factor is $\beta = 1/R$.

I set $\phi_n = .7$, near the middle of the values used in the BP literature. Period utility over consumption and leisure takes the form

$$u_j(c, 1-n) = log(c) + \psi_j \frac{(1-n)^{1-\gamma}}{1-\gamma} \ . \qquad (10)$$

I set $\gamma = 2$ to imply a mean Frisch elasticity[10] for non-leisure time of .7. These parameter values are summarized in Table A1 in the Appendix.

5.2 Parameters set conditional on α

The remaining model parameters are set to bring model predictions in line with the aggregate facts for US workers presented in Section 4.1. These parameters include: the distribution of initial human capital and ability, $\Lambda(h_1, a)$; the depreciation rate, δ; the utility of leisure parameter, ψ; and the curvature of new human capital with respect to existing human capital, ϕ_h. All parameters are jointly determined, but each can be linked to a specific target.

[9] I partition ages into periods according to the following schedule: Period 1, 23-28; Period 2, 29-33; Period 3, 34-39; Period 4, 40-44; Period 5, 45-50; Period 6, 51-55; Period 7, 56-60.

[10] See Guvenen et. al (2013). Given the utility specification above, Frisch elasticity of non-leisure time is $\varepsilon_F = \frac{1-n}{n\gamma}$. In the data, mean annual hours worked over all ages are 2,141. With an assumed time endowment of 5,110 hours, this implies that $\bar{n}_3^{data} \approx .42$. Since the model is calibrated to match this moment, we are left with $\gamma = \frac{1-\bar{n}_3^{data}}{.7\bar{n}_3^{data}} \approx 2$.

The distribution of earnings over the life-cycle:

The distribution of initial human capital and learning ability, $\Lambda(h_1, a)$ is set to match the key features of the life-cycle earnings distribution observed in Figures 2 and 3. I assume $\Lambda(h_1, a)$ is log Normal, which leaves me with five distribution parameters to set: $(\mu_{h_1}, \mu_a, \sigma_{h_1}, \sigma_a, \rho_{h_1,a})$. The mean values of initial human capital and ability, μ_{h_1} and μ_a, are set to match mean earnings in period 1 and the peak of mean earnings, respectively. The variance of initial human capital and ability, σ_{h_1} and σ_a, are set to match the variance of *log* earnings in period 1 and the peak of earnings variance, respectively. The correlation of initial human capital and ability, $\rho_{h_1,a}$ is set to minimize the squared distance between the earnings variance profiles for the model and the data. Huggett, Ventura, and Yaron (2011) show that a low correlation between initial human capital and ability will cause the spread in earnings to decline over the beginning of the life-cycle as the earnings of high-a/low-h_1 individuals catch up to the earnings of low-a/high-h_1 individuals. To be consistent with the nearly monotonoic increase in the variance of earnings observed in the data the correlation parameter must be sufficiently large, so that individuals endowed with a high a also tend to have high endowments of h_1.

The path of wages at the end of the working life-cycle:

The depreciation rate δ is set to match the decline in the mean hourly wage rate observed over the final period of the data, $\bar{w}_J^D/\bar{w}_{J-1}^D$. Recalling the human capital process from Section 3.1, the ratio of an individual's hourly wage in period J relative to period $J-1$ is

$$\frac{wh_J n_J(1-s_J)/n_J}{wh_{J-1}n_{J-1}(1-s_{J-1})/n_{J-1}} = \frac{h_J}{h_{J-1}(1-s_{J-1})} \approx \frac{h_J}{h_{J-1}} = (1-\delta) + \frac{H(h,n,s,a;\alpha)}{h_{J-1}} \,. \quad (11)$$

The first equality in (11) follows from $s_J = 0$; the approximate equality follows from $s_{J-1} \approx 0$, which approximately holds because any newly produced human capital can only be used for one period before retirement hits. In the BP ($\alpha = 1$) specification, $H \approx 0$ since $s_{J-1} \approx 0$, which means (11) approximately reduces to $(1-\delta)$. Therefore δ can be approximated by

$$\delta^{BP} \approx 1 - \bar{w}_J^D / \bar{w}_{J-1}^D .$$

Alternatively, in the LBD ($\alpha = 0$) specification significant investment still occurs even in period $J-1$ because production is equivalent to investment. Denoting $x = (h_1, a)$, the calibrated value of δ in the LBD case is given by

$$\delta^{LBD} = 1 - \bar{w}_J^D / \bar{w}_{J-1}^D + \int a h_{J-1}^{\phi_h - 1}(x) n_{J-1}^{\phi_n}(x) d\Lambda(x) \; .$$

The life-cycle profile of hours worked:

I assume that individuals are endowed with 5110 hours per year, which translates to 14 hours per day. I initially set the utility of leisure parameter ψ to match the peak of mean hours worked in the data. When calibrated using this approach, period J hours in the BP model are 3% below their peak, and period J hours in the LBD model are 27% below their peak, relative to a 7% decline in the data.

To reduce the discrepancy in predicted hours worked between the BP and LBD models, I consider including three additional model parameters. I first allow the utility of leisure to evolve at a constant rate over the life-cycle, $\psi_1 = \psi$ and $\psi_{j+1} = \psi_j(1 + g_\psi)$. Second, I impose $\psi_j = \psi \; \forall j$, but allow the depreciation rate to evolve at a constant rate, $\delta_1 = \delta$ and $\delta_{j+1} = \delta_j(1 + g_\delta)$. Finally I impose $\psi_j = \psi$ and $\delta_j = \delta \; \forall j$, but allow the learning ability of individuals to evolve at a constant rate, $a_1 = a$ and $a_{j+1} = a_j(1+g_a)$. In each case I continue to target ψ to match the peak of mean hours in the data, and the growth parameter g is targeted to match mean hours worked in period J.

Remaining Parameters:

Two parameters have yet to be set: the curvature of new human capital with respect to existing human capital, ϕ_h, and the elasticity of substitution between studying and production, ρ. In fact, since I only consider the extreme cases $\alpha = 0, 1$, the value of ρ is irrelevant (see equations (6) and (7) in Section 3.2).

Lower values of ϕ_h generate larger growth rates in human capital in the first period of the life-cycle and lower growth rates in human capital later in the life-cycle. Therefore, lower values of ϕ_h cause the variance in earnings growth rates to decline more from the first to the second period of the life-cycle. In Figure 4 and Figures A1-A5 roughly half of the decline in growth variance profiles occur after age 35. To be consistent with this feature of the data, I restrict ϕ_h from below so that at least one third of the life-cycle decline in the variance of earnings growth occurs after period 1.

Calibration summary:

Table A2 summarizes the BP and LBD calibration outcomes. I do not include the BP calibrations in which g_δ or g_a are nonzero for two reasons. First, the BP calibration with $g_\psi > 0$ is consistent with the earnings and wage growth facts examined next in Section 6. Second, the implied value of g_ψ is both modest and intuitive: it says that individuals in their late fifties value leisure 22% more than do individuals in their early twenties. I also do not include the LBD calibration in which g_ψ is nonzero. This calibration implied a period depreciation rate over 60%, which is substantially higher than all estimates of human capital depreciation in the literature. Figures A6-A7 plot the distribution of hours, annual earnings, and hourly wages for the BP calibration with $g_\psi > 0$ and the LBD calibration with $g_a < 0$.

A key takeaway from the LBD calibrations is that to be consistent with the distribution of earnings, wage, and hours in the data, the LBD model must feature very large life-cycle changes in some parameter. In particular, from their early twenties to their late fifties individuals must experience either a 390% increase in their human capital depreciation rate or an 73% decline in their learning ability.

Another important observation is that the LBD calibrations feature larger depreciation rates and lower values of ϕ_h than the BP calibrations. In the calibrated BP model, wage and earnings growth slow dramtically over the life-cycle because individuals *endogenously* reduce the fraction of hours worked they spend studying: the mean fraction of time spent studying falls from 10% in period 1 to .08% by period 5. In the calibrated LBD model, however, mean investment (time in production) is basically constant over the life-cycle because all hours worked are spent on production and the life-cycle profile of hours worked is flat. Therefore, the LBD model requires that wages grow more slowly over the life-cycle due to *exogenous* curvature which is built into the human capital process.

6 Ben-Porath and Learning By Doing: the Impact of α

The life-cycle model laid out in Section 3 featured a human capital model that weighted two time inputs to human capital formation: production and studying. The LBD ($\alpha = 0$) version assumed that production was the sole source of human capital investment; the BP ($\alpha = 1$) version assumed that studying was the sole source of human capital investment. Section 5 calibrated both versions of the model to match several aggregate life-cycle facts for US workers. In this section I show that the calibrated BP version of the model can account for a larger fraction of the life-cycle decline in the variance of wage and earning growth rates, relative to the calibrated LBD model.

Figure 5 plots the age profiles for the variance in earnings growth rates generated by the BP and LBD models. The dashed red line corresponds to the BP calibration for which $g_\psi > 0$. The dashed blue line corresponds to the LBD calibration for which $g_a < 0$; this is the LBD calibration whose profile is most consistent with the data. Panels (a), (b), and (c) also plot the data profiles corresponding to Figures 4, A4, and A5, respectively[11]. The sample restrictions which were applied to the data in panels (a)-(c) are also applied to the

[11] The 5-year wage growth rates in Figure 4 approximately correspond to one-period wage growth rates, since a period represents $38/7 \approx 5.4$ years. Therefore I compare age effects for these growth rates against 1-period growth rates produced by the model. The moments for 6-year wage growth rates are very similar to those in Figure 4.

model profiles plotted in the same panels (see Section 4.2). Several sources of variation in measured earnings growth, especially measurement error, are not present in my model, so I do not compare the raw variance profiles in the model and data. Instead, I normalize the age profiles in the model and data to zero for period 1-2 growth, and compare the change in the age profiles over the life-cycle.

Figure 5: Life-cycle Variance of 1-Period Earnings Growth Rates

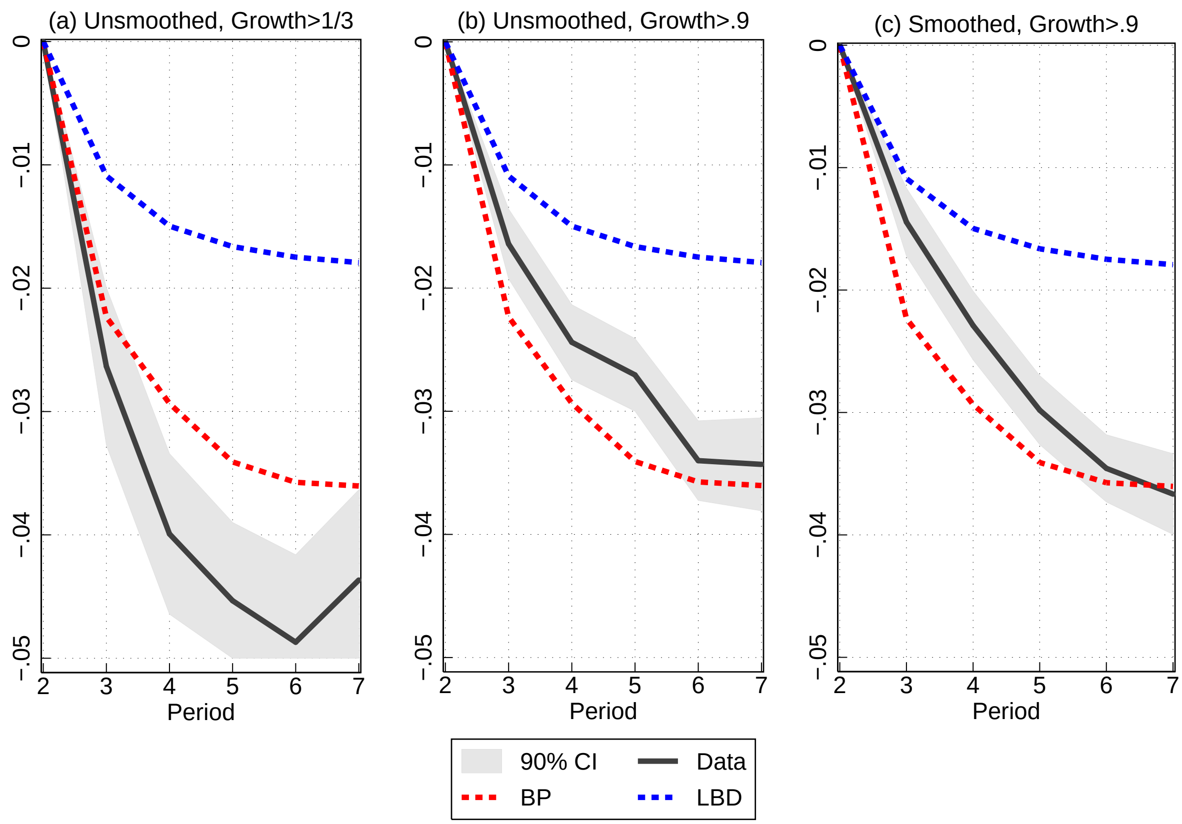

Across the three panels the BP profile accounts for between 82% and 105% of the life-cycle decline observed in the data, while the LBD profile accounts for between 40% and 52% of the decline. Said differently, increasing α from 0 to 1 doubles the fraction of the life-cycle decline accounted for by the human capital model.

It is straightforward to change the calibration procedure so that the LBD model explains a greater share of the life-cycle decline in the data profile. To see this, first note that heterogeneity in earnings growth rates is primarily due to heterogeneity in learning ability, and the variance of ability σ_a is calibrated to match the variance in *log* earnings at the end of the life-cycle (see figure A7(a)). One alternative calibration strategy would instead target σ_a to the life-cycle decline in the variance of earnings growth. I perform this alternate calibration for the LBD model, setting σ_a so that the life-cycle decline in the variance of earnings growth is equivalent to the decline for the BP model depicted in Figure 5(a). The correlation parameter $\rho_{h_1,a}$ is set so that the variance of earnings in period 2 matches that in the BP model. Figure 6 plots model outcomes for this calibration alongside the data and the profiles for the original BP calibration. When the life-cycle decline in the variance of earnings *growth* in the LBD model is equal to that of the BP model, the LBD model predicts that the life-cycle increase in the variance of earning *levels* is 65% larger than in the data.

Figure 6: Alternate LBD Calibration

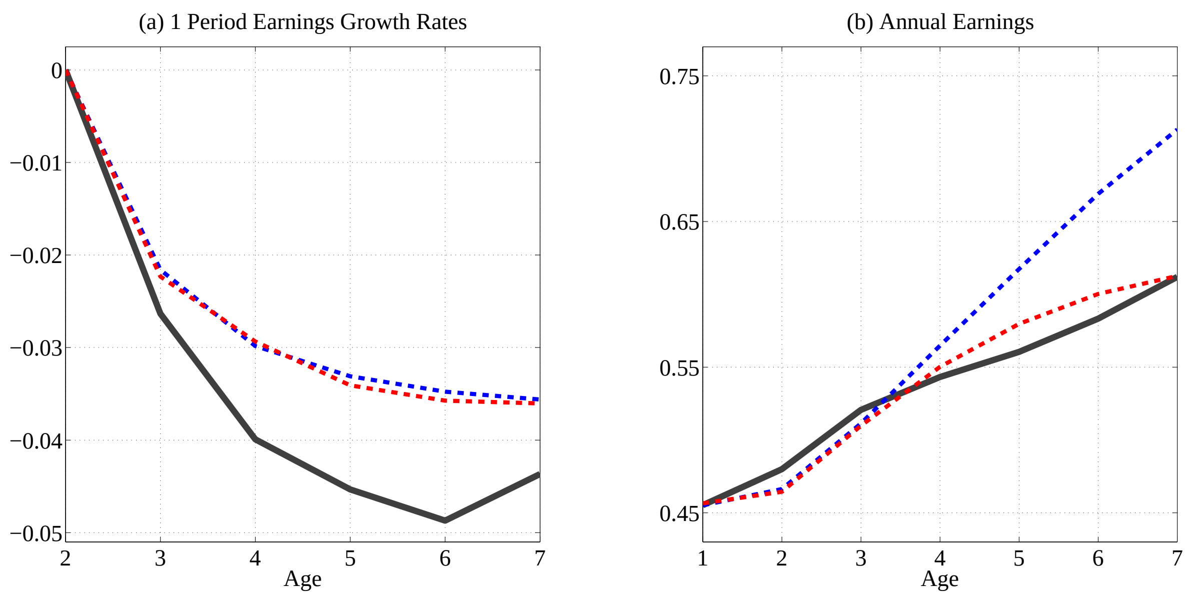

The question, then, is why the LBD model generates such a large variance in earnings among older workers when the variance in earnings growth rates among young workers is large. The answer is that in the LBD model individuals who experience high earnings growth early on *increase* investment in human capital at a faster rate than individuals whose earnings grow less, while in the BP model the opposite is true. Figure 7(a) plots the life-cycle investment differentials between the top quartile of period 1 earnings growers (the “High” group) and the bottom quartile of period 1 earnings growers (the “Low” group). In period 1 of the LBD model the High group spends roughly the same amount of time investing in human capital as the Low group, but by period 6 the High group is investing an additional 10% of the time endowment. In the BP model this pattern is reversed: the High group invests more than the Low group in period 1, and this gap disappears by period 6. The result, seen in Figure 7(b), is that the High-Low earnings gap increases more over the life-cycle in the LBD model than in the BP model.

The intuition is that earnings growth is correlated with human capital growth, which raises the marginal utility of production relative to studying. In the LBD model this increases the *marginal utility* of human capital investment (production) for the High group relative to the Low group, while in the BP model this increases the *opportunity cost* of investment (studying) for the High group relative to the Low group.

Figure 7: Investment and Wage Growth

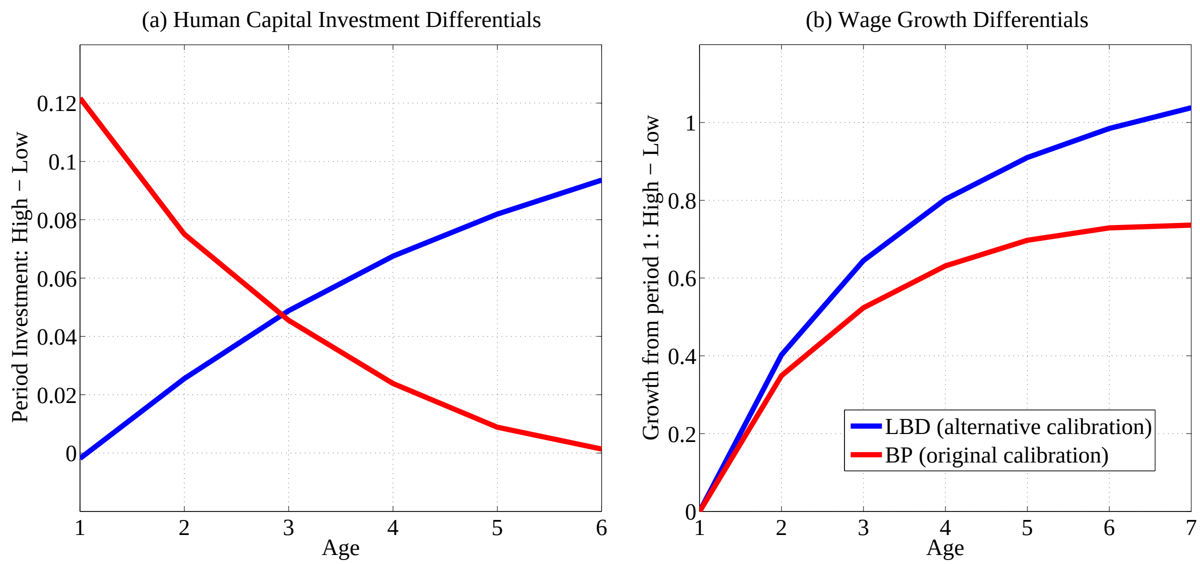

Note: The "High" group for each model is the top quartile of individuals in terms of hourly wage growth from period 1-2. The "Low" group is the bottom quartile. The red curve in plot (a) at age j is the mean investment in period j for the High group minus the mean investment in period j for the Low group, within the BP model. The red curve in plot (b) at age j is the mean hourly wage growth from period 1 to period j for the High group minus the mean hourly wage growth from period 1 to period j for the Low group, within the BP model. The blue curves plot these same statistics for the LBD model under the alternative calibration.

7 Policy Analysis

The previous section found that the human capital model from Section 3 can account for the entire life-cycle decline in the variance of earnings growth among US workers when α is close to one. In this section I show that α also has implications for government policy. Specifically, changes to the progressivity of income taxes generate larger impacts on economic aggregates in a BP environment relative to a LBD environment. This result is complementary to work by Guvenen, Kuruscu, and Ozkan (2014); and Badel and Huggett (2014), which quantitatively show that a progressive tax on earnings can generate large reductions in human capital investment. A key difference between my analysis and theirs is that I compare results for a LBD and BP process, while they assume a BP process outright.

Of course, even in a model with exogenous wages a revenue neutral decrease in the pro-

Figure 8: Income Tax Rates

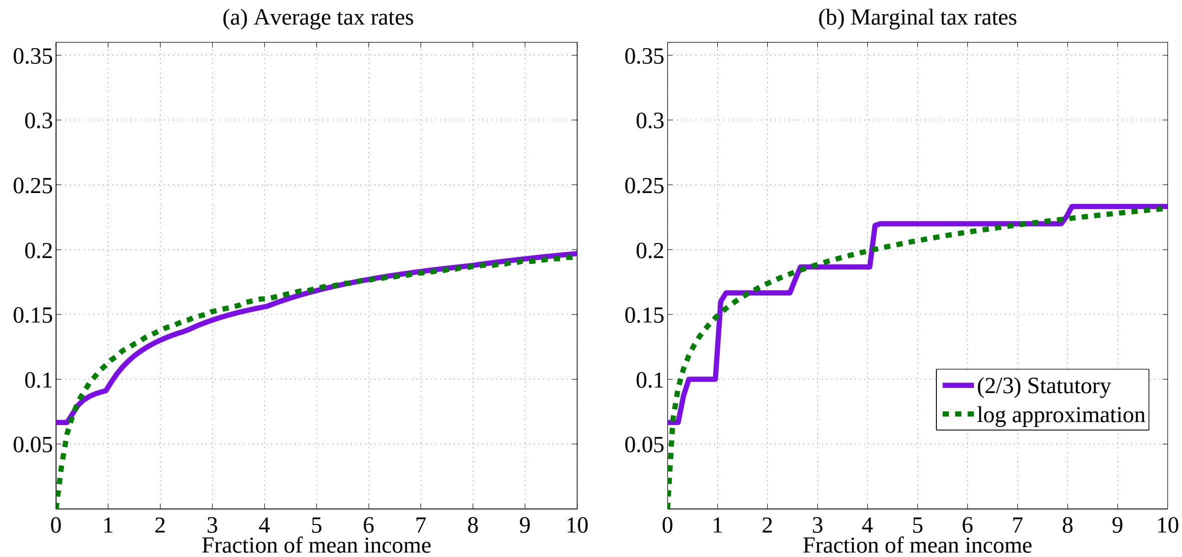

gressivity of earnings taxes will increase aggregate hours and earnings. However, when wage processes are endogenized within a human capital framework aggregate earnings will grow not only because hours increase, but because earnings per hour increase as well. The reason is that lowering the slope of marginal tax rates increases the returns to human capital investment. This induces individuals to invest more, which increases the stock of human capital in the economy. An open question is: do aggregate variables respond more given a BP process or a LBD process? The stripped down analysis in Section 3.3 found that investment and pre-tax earnings respond more to a change in the return to human capital in a BP model. I now conduct a more detailed policy analysis, using the calibrated BP and LBD models discussed in Sections 5-6.

I analyze the steady state effects of a revenue-neutral change from a progressive earnings tax modeled after the United State's 2009 federal income tax to a revenue-equivalent proportional earnings tax. I first write down the 2009 U.S. statutory average federal income tax rate schedule, ignoring all deductions and assuming that all income is in the form of labor earnings, as a function of the mean income in my 2009 sample ($47,000). To provide a conservative estimate for the effects of a transition to a proportional income tax I assume that all

Table 1: Equilibrium Impact of the Tax Change

	% Change in Investment Time	**% Change in Human Capital**	**% Change in Hours Worked**	**% Change in Earnings**
BP	11.80%	1.89%	2.17%	3.58%
LBD	1.63%	.45%	1.63%	2.27%
Ratio	7.2	4.2	1.3	1.6

households are able to avoid one third of the progressive tax; the resulting average tax rates are plotted in Figure 8(a). This curve can be approximated by the logarithmic function[12] $\tau(y/\bar{y}) = .1129 + .0355\ log(y/\bar{y})$. I subject all labor earnings in the model economies to this average tax rate schedule. The implied marginal tax rates are graphed in Figure 8(b), and grow monotonically from 7% for households with one tenth the mean income to 23% for individuals with ten times the mean income in the economy.

All tax revenue collected by the government is dumped in the ocean. The baseline BP and LBD models were calibrated assuming zero taxes and transfers in the economy. Since utility from consumption is *log* and additively separable from leisure utility, a proportional tax on labor income does not affect the decisions of any individual in the economy when the revenue is wasted. Therefore, the model outcomes from Section 6 are consistent with any proportional labor tax rate.

To compute aggregate figures I assume no population growth. The results from this exercise are summarized in Table 1. The most striking effects of the tax change are with respect to the first column: switching from a modestly progressive labor income tax to a proportional labor income tax increases mean investment time by over 11% in a BP economy. The size of this reduction is seven times larger than the 1.6% change in the LBD economy. The increase in investment translates to a nearly 2% addition to the human capital stock in the BP economy, quadruple the change observed in the LBD economy. Aggregate earnings

[12] Guner, Kaygusuz, and Ventura (2013) provide parametric estimates of tax functions which approximate federal income tax liabilities for US households. Their estimates correspond to liabilities as a function of *total* household income. To isolate the effects of a change in the *labor income* tax I opt to use statutory federal tax rates for labor income, but make convenient use of their tax function.

in the BP economy grow by 3.6%, which is roughly equivalent to two years of per capita income growth along the U.S.'s postwar trend, and is 60% larger than the change in the LBD economy.

In summary, the aggregate effects of a revenue neutral change from a modestly progressive labor income tax to a proportional labor income tax are substantially larger in a BP environment than in a LBD environment. It is important to emphasize that these findings should be viewed as a very conservative floor for the changes that would result from a more detailed policy analysis, for several reasons. First, both the size and slope (with respect to labor income) of effective marginal tax rates are likely larger than those that were used for this exercise. Second, the Frisch elasticity of hours worked used in the calibration is on the low end of estimates when human capital accumulation is considered[13]. Adopting a larger elasticity would likely generate larger responses to the tax change. Finally, this exercise assumes that human capital investment prior to age 23 is policy invariant. Just as human capital investment after age 23 increased in response to a flattening of tax rates, one would expect early human capital investment to increase as well, which would result in an even larger response in wages and earnings[14].

8 Conclusion

This paper constructs a life-cycle human capital model which nests both BP and LBD as extreme special cases. The fundamental difference between BP and LBD is that in the BP case individuals can only invest in their future earning potential by forgoing earnings or leisure today, while in the LBD model investment in future earning potential and earning today are inseparable. I use the model to highlight a class of policies whose implications differ across the BP and LBD cases, and to ask whether observable macroeconomic moments

[13]Wallenius (2011) finds that estimates of the intertemporal elasticity of substiution of hours worked which assume an exogenous wage process should be multiplied by 2.1 if a LBD process governs wages, or 2.6 if a BP process governs wages. Since estimates assuming an exogenous process are generally clustered around .3, this implies a true parameter value of .63-.78. Imai and Keane (2004) estimate an even larger parameter value around 3.8 when they assume a LBD technology.

[14]Krueger and Ludwig (2013) study the effect of taxation on education decisions.

are informative about the relative importance of BP investment versus LBD investment.

My findings are as follows. First, I show that aggregate human capital investment and earnings respond more to changes in the return to human capital investment in a BP environment than in a LBD environment, and that the difference in these responses is quantitatively important. Second, both the BP and LBD versions of the model are consistent with the aggregate distribution of life-cycle earnings, hourly wages, and hours worked for men in the PSID (1968-2009). Finally, conditional on matching these aggregate *levels* facts, the BP version of the model is more consistent with the large variance of earnings and wage *growth rates* observed in the data.

In sum, I demonstrate that a simple economic story in which workers respond optimally to a tradeoff between earnings today and earning potential tomorrow can explain a broad set of aggregate facts for US workers. I also show that when this tradeoff exists, human capital investment decisions and earnings will be sensitive to changes in the returns to human capital investment. This suggests that policies which affect the return to human capital investment—such as changes to tax progressivity, education subsidies, or pension reform—will generate significant endogenous responses in key macroeconomic variables.

References

Andalfatto, D., Gomme, P., and C. Ferrall, (2000). "Human Capital Theory and the Life-Cycle Pattern of Learning and Earning, Income and Wealth", mimeo.

Badel, A. and M. Huggett (2014). "Taxing Top Earners: A Human Capital Perspective," Working Paper 2014-017B, posted July 2014, updated October 2014. http://research.stlouisfed.org/wp/more/2014-017

Becker, Gary S. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education.* Chicago, University of Chicago Press.

Becker, Gary S. (1965) "A Theory of the Allocation of Time," Economic Journal 75 (299), 493-517.

Ben-Porath, Yoram (1967). "The Production of Human Capital and the Life Cycle of Earnings," Journal of Political Economy, University of Chicago Press, 75(4), 352-365.

Belley, Philippe (2012). "Understanding Wage Growth: Estimating and Testing Learning-by-Doing," mimeo. http://www.k-state.edu/economics/staff/websites/belley/LBD.pdf

Browning, M., Hansen, L.P., and J. Heckman (1999). "Micro data and General Equilibrium Models," in J.B. Taylor and M. Woodford, eds., Handbook of Macroeconomics.

Chang, Y., Gomes, J.F., and F. Schorfheide (2002). "Learning-by-Doing as a Propagation Mechanism," American Economic Review, 92, 1498-1520.

Cossa, R., Heckman, James, and L. Lochner (2003). "Wage Subsidies and Skill Formation: A Study of the Earned Income Tax Credit," in E. Phelps, ed., Designing Inclusion, Cambridge University Press, 2003.

Deaton, Angus (1997). "The Analysis of Household Surveys: A Microeconomic Approach to Development Policy," World Bank Publications, 123-127.

Diamond, P. and E. Saez (2011). "The Case for a Progressive Tax: From Basic Research to Policy Recommendations," Journal of Economic Perspectives, 25, 165-90.

Gemici, A. and M. Wiswall (2014). "Evolution Of Gender Differences In Post-Secondary Human Capital Investments: College Majors," International Economic Review, vol. 55, pages 23-56, 02.

Guvenen, F., Kuruscu, B., and Serdar Ozkan. (2014). "Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis," Review of Economic Studies, 81, 818-850.

Hansen, G. D. and S.Imrohoroglu, (2009) "Business Cycle Fluctuations and the Life Cycle: How Important is On-The-Job Skill Accumulation?," Journal of Economic Theory, 144 (6), 2293-2309.

a. Heckman, James (1976). "Estimates of a Human Capital Production Function Embedded in a Life-Cycle Model of Labor Supply," Household Production and Consumption, ed. Nestor E. Terleckyj, NBER, 225-264.

b. Heckman, James J. (1976). "A Life-Cycle Model of Earnings, Learning, and Consumption," Journal of Political Economy, University of Chicago Press, 84(4), S11-S44.

Heckman, J., Lochner, L., and C. Taber (1998). "Tax-Policy and Human Capital Formation", American Economic Review, 88, 293-297.

Huggett, M., Ventura, G., and A. Yaron (2006). “Human Capital and Earnings Distribution Dynamics”, Journal of Monetary Economics, 53, 265-90.

Huggett, M., Ventura, G., and A. Yaron (2011). “Sources of Lifetime Inequality,” American Economic Review, 101(7): 2923-54.

Imai, S., and M. P. Keane (2004). “Intertemporal Labor Supply and Human Capital Accumulation,” International Economic Review, 45(2), 601-641.

Killingsworth, Mark R. (1982). “ ‘Learning by Doing’ and ‘Investment in Training’: A Synthesis of Two ‘Rival’ Models of the Life Cycle,” The Review of Economic Studies, 49(2), 263-271.

Kindermann, F., and D. Krueger (2014). “The Redistributive Benefits of Progressive Labor Income Taxation: How Large Should Top Marginal Tax Rates Be?” mimeo.

Krueger, Dirk, and Alexander Ludwig. 2013. ”Optimal Progressive Labor Income Taxation and Education Subsidies When Education Decisions and Intergenerational Transfers Are Endogenous.” American Economic Review, 103(3): 496-501.

Mincer,Jacob (1958).“Investment in Human Capital and Personal Income Distribution,” Journal of Political Economy, University of Chicago Press, 66, 281.

Peterman, William B. “The Effect of Endogenous Human Capital Accumulation on Optimal Taxation,” FEDS Working Paper 2012-03, March 2012.

Schultz, Theodore W. (1961). “Investment in Human Capital,” American Economic Review, 51(1):1-17.

Shaw, K. L. (1989). “Life-Cycle Labor Supply with Human Capital Accumulation,” International Economic Review , 30, 431-56.

Wallenius, Johanna (2011). “Human Capital Accumulation and the Intertemporal Elasticity of Substitution of Labor: How Large is the Bias?” Review of Economic Dynamics, 14, 577-591.

Appendix

A Robustness for Figure 4

Figure A1: Life-cycle Variance of 5-year Growth Rates

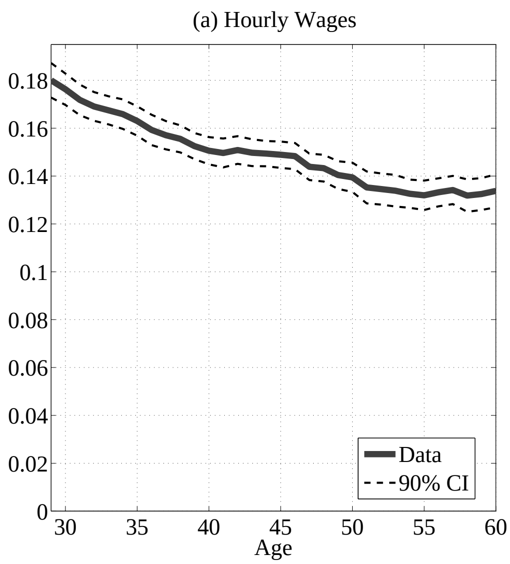

Restrictions on hours, wages, and earnings are identical to those listen in Section 4, except that the minimum hours requirement is raised from 500 to 1,800 annual hours. The number of valid observations is 38,464, down from 42,667 for the data corresponding to Figure 4.

Figure A2: Variance of 5-year Growth Rates by Age, Education

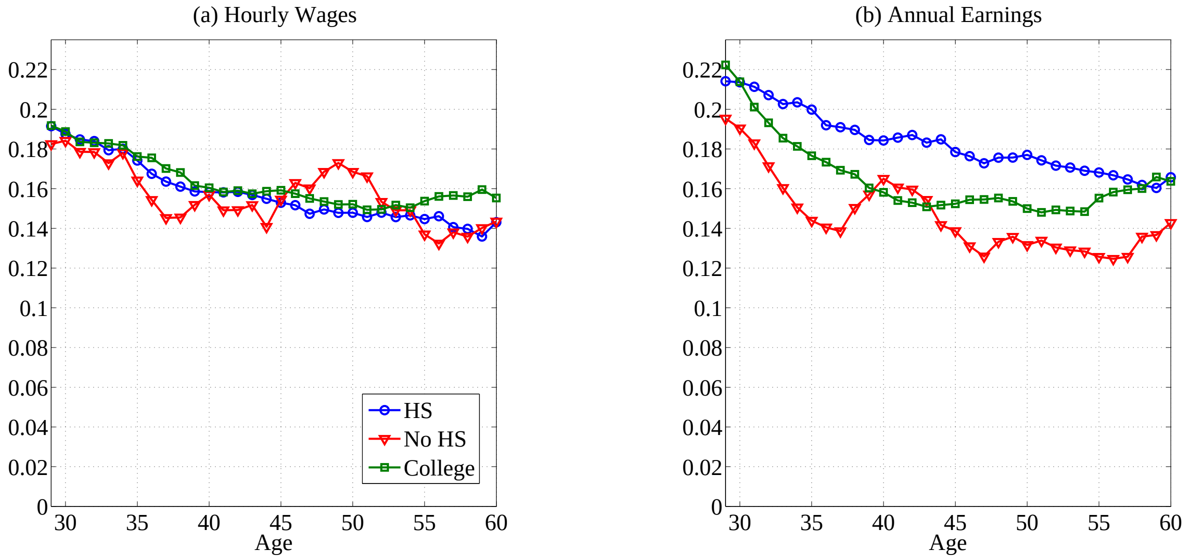

Restrictions on hours, wages, and earnings are identical to those listen in Section 4. Education is measured as the number of grades completed (see the variable series which has label 30010 for year 1968). The "No HS" category includes individuals who have completed 0-11 grades; the "HS" category includes individuals who have completed 12-13 grades; the "College" category includes individuals who have completed more than 13 grades. The number of valid observations is 4,269 for the "No HS" category, 16,389 for the "HS" category, and 22,009 for the "College" category.

Figure A3: Variance of 5-year Growth Rates by Experience, Education

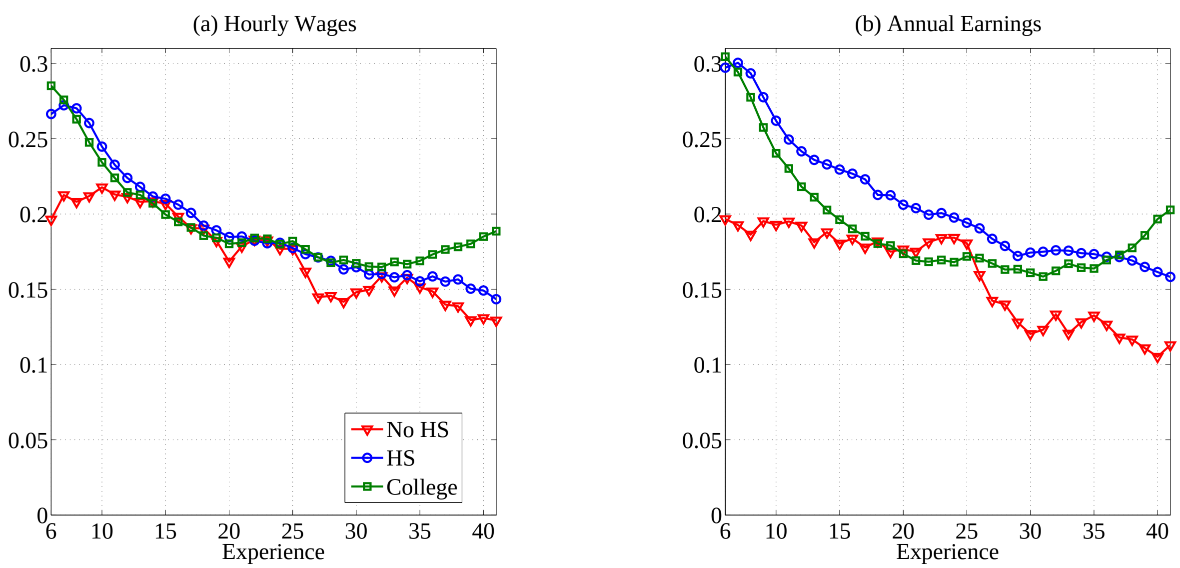

Restrictions on hours, wages, earnings, and education are equivalent to those in Figure A3. Experience is measured as potential experience: age minus grades completed minus 6. The number of valid observations is 5,954 for the "No HS" category, 25,422 for the "HS" category, and 30,518 for the "College" category.

Figure A4: Life-cycle Variance of 5-year Growth Rates, $w_j/w_{j-5} \geq .9$

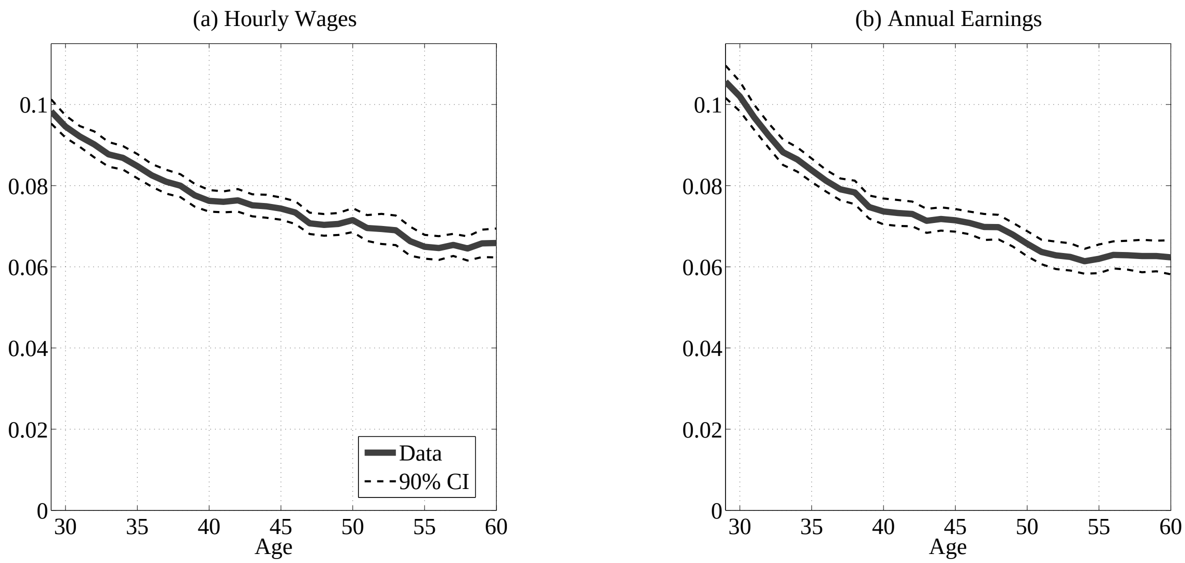

Restrictions on hours, wages, and earnings are identical to those listen in Section 4, except that the minimum growth cutoff is raised from .33 to .9. The number of valid observations is 33,803, down from 42,667 for the data corresponding to Figure 4.

Figure A5: Life-cycle Variance of 5-year Smoothed Growth Rates, $w_j/w_{j-5} \geq .9$

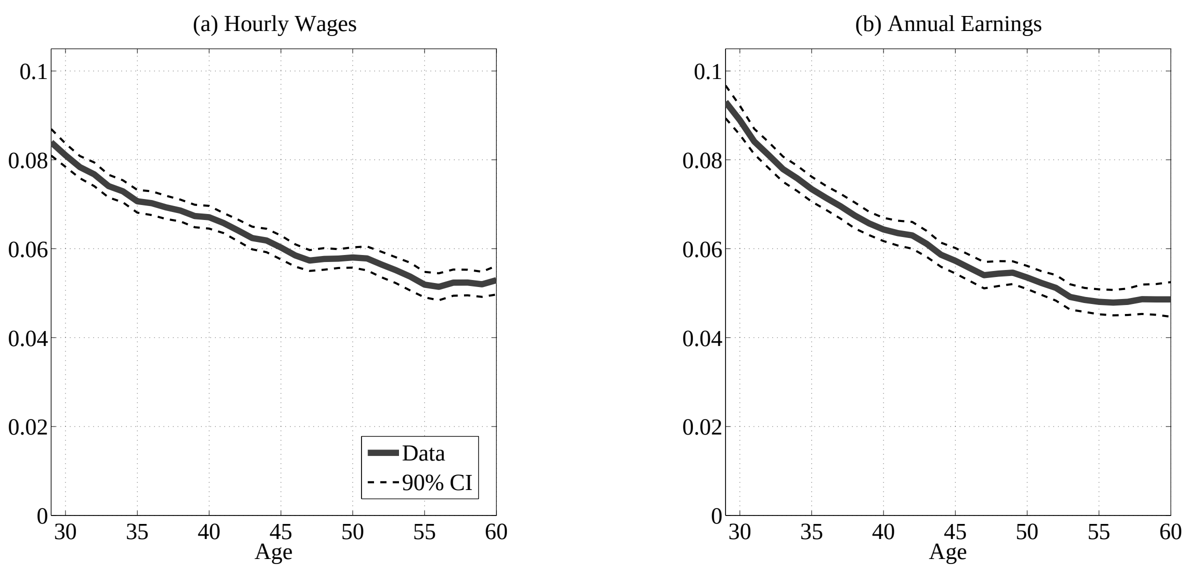

Hourly wages and earnings are smoothed with a 3 year centered MA process. Restrictions on hours, smoothed wages, and smoothed earnings are identical to those listen in Section 4, except that the cutoff for wage and earnings growth is raised from .33 to .9. The number of valid observations is 36,103, down from 42,667 for the data corresponding to Figure 4.

B Parameter Values for Different Model Calibrations

Table A1: Parameters set independently of α

Parameter	Description	Value
γ	Curvature of utility of leisure	2
β	Time discount factor	.81
$1 - R$	Interest rate	.24
$\bar{J}$	Periods in life-cycle	10
J	Periods in working life-cycle	7

Table A2: Parameters set conditional on α

Parameter	**Description**	**BP: Basic**	**BP:** g_ψ	**LBD: Basic**	**LBD:** g_δ	**LBD:** g_a
(μ_{h_1}, μ_a)	Population means for $log(h_1, a)$	(5.0,1.4)	(5.0,1.4)	(4.8,4.2)	(4.8,4.0)	(4.7,2.3)
(σ_{h_1}, σ_a)	Variance of $log(h_1, a)$	(.42,.11)	(.42,.11)	(.41,.29)	(.42,.52)	(.41,.30)
$\rho_{\bar{h}_1 a}$	Correlation between (h_1, a)	.95	.93	.99	.94	.66
ϕ_h	Curvature of $H(h, n, s)$ w.r.t. h	.6	.6	0.2	.0	.4
ψ	Weight on utility of leisure	1.04	.94	1.57	1.21	1.16
δ	Human capital depreciation rate	.022	.022	.559	.035	.079
$(1 + g_\psi)^{J-1}$	Lifetime growth in leisure utility	1	1.22	1	1	1
$(1 + g_\delta)^{J-2}$	Lifetime growth in depreciation	1	1	1	4.9	1
$(1 + g_a)^{J-2}$	Lifetime growth in ability	1	1	1	1	0.27

C Comparing Calibrated Model Predictions to Targeted Data

This appendix compares the outcomes of the calibration procedure for the LBD and BP models to the moments obtained from the PSID. The BP profile corresponds to column 4 of Table A2; the LBD model corresponds to column 7. Figure A6 displays life-cycle means for (a) annual earnings, (b) hourly wages, and (c) hours worked. Recall that the ceiling for annual hours worked in my sample is 5,110. Therefore, hours can converted into a fraction of the time endowment by dividing them by 5,110. Both LBD and BP generate predictions broadly in line with the data. In Figure A7 I plot (a) the log variance profiles for earnings.

Figure A6: Life-cycle Means

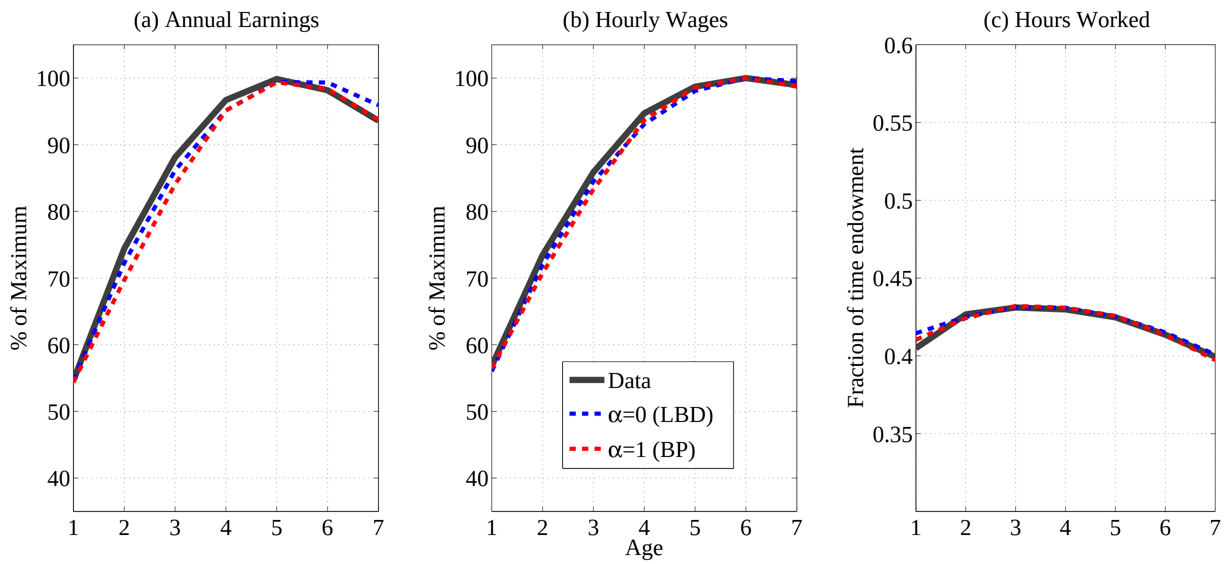

Figure A7: Life-cycle Variance of *log*'s