

NBER WORKING PAPER SERIES

SOURCES OF LIFETIME INEQUALITY

Mark Huggett
Gustavo Ventura
Amir Yaron

Working Paper 13224
<http://www.nber.org/papers/w13224>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2007

We thank seminar participants at the SED, Minnesota Workshop in Macroeconomic Theory, NBER Summer Institute, Productivity over the Life Cycle Conference (Bank of Canada), Macroeconomics of Imperfect Risk Sharing (UC-Santa Barbara), PSU-Cornell Macro Theory Conference, Georgetown, NYU, Oslo, Penn, UCLA, Wharton, Wisconsin and Yale. We thank the National Science Foundation Grant SES-0550867 for research support. Ventura thanks the Research and Graduate Studies Office from The Pennsylvania State University for support. Yaron thanks the Rodney White Center at the The Wharton School for research support. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Mark Huggett, Gustavo Ventura, and Amir Yaron. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Sources of Lifetime Inequality

Mark Huggett, Gustavo Ventura, and Amir Yaron

NBER Working Paper No. 13224

July 2007

JEL No. D31,D91,E21

ABSTRACT

Is lifetime inequality mainly due to differences across people established early in life or to differences in luck experienced over the working lifetime? We answer this question within a model that features idiosyncratic shocks to human capital, estimated directly from data, as well as heterogeneity in ability to learn, initial human capital, and initial wealth -- features which are chosen to match observed properties of earnings dynamics by cohorts. We find that as of age 20, differences in initial conditions account for more of the variation in lifetime utility, lifetime earnings and lifetime wealth than do differences in shocks received over the lifetime. Among initial conditions, variation in initial human capital is substantially more important than variation in learning ability or initial wealth for determining how an agent fares in life. An increase in an agent's human capital affects expected lifetime utility by raising an agent's expected earnings profile, whereas an increase in learning ability affects expected utility by producing a steeper expected earnings profile.

Mark Huggett
Georgetown University
mh5@georgetown.edu

Gustavo Ventura
Department of Economics
University of Iowa
W358 PBB
Iowa City, IA 52242
gventura1967@gmail.com

Amir Yaron
The Wharton School
University of Pennsylvania
2256 Steinberg-Dietrich Hall
Philadelphia, PA 19104-6367
and NBER
yaron@wharton.upenn.edu

1 Introduction

To what degree is lifetime inequality due to differences across people established early in life as opposed to differences in luck experienced over the lifetime? Among initial conditions, individual differences established early in life, which ones are the most important?

A convincing answer to these questions is of fundamental importance. First, and most simply, an answer serves to contrast the potential importance of the myriad policies directed at modifying or at providing insurance for initial conditions (e.g. public education) against those directed at shocks over the lifetime (e.g., unemployment insurance programs). Second, a discussion of lifetime inequality cannot go too far before discussing which type of initial condition is the most critical for determining how one fares in life. Third, a useful framework for answering these questions should also be central in the analysis of a wide range of policies considered in macroeconomics, public finance and labor economics.

We view lifetime inequality through the lens of a risky human capital model. Agents differ in terms of three initial conditions: initial human capital, learning ability and financial wealth. As agents age, they accumulate human capital by optimally dividing their available time between market work and human capital accumulation. Human capital and labor earnings are risky as human capital is subject to uninsured, idiosyncratic shocks each period.

We ask the model to account for key features of the earnings distribution dynamics by cohorts. To this end, we document how mean earnings and measures of earnings dispersion and skewness evolve for U.S. males. We find that mean earnings are hump shaped and that earnings dispersion and skewness increase with age over most of the working lifetime.¹

Our model produces a hump-shaped mean earnings profile by a standard human capital channel. Early in life earnings are low as agents allocate time to accumulating human capital. Earnings rise as human capital accumulates and as a greater fraction of time is devoted to market work. Earnings fall later in life because human capital depreciates and little time is put into producing new human capital.

Two forces within the model account for the increase in earnings dispersion. One force is that agents differ in learning ability. Agents with higher learning ability have steeper mean

¹Mincer (1974) documents related patterns in U.S. cross-section data. Deaton and Paxson (1994), Storesletten, Telmer and Yaron (2004), Heathcote, Storesletten and Violante (2005) and Huggett, Ventura and Yaron (2006) examine cohort patterns in U.S. repeated cross section or panel data.

earnings profiles than low ability agents, other things equal.² The other force is that agents differ in idiosyncratic human capital shocks received over the lifetime.

To identify the contribution of each of these forces, we exploit the fact that the model implies that late in life little or no new human capital is produced. As a result, moments of the change in wage rates for these agents are almost entirely determined by shocks, rather than by shocks and the endogenous response of investment in human capital to shocks and initial conditions. We estimate the shock process from U.S. data using precisely these moments. Given an estimate of the shock process and other model parameters, we choose the initial distribution of financial wealth, human capital and learning ability across agents to best match the earnings facts described above.³ We find that learning ability differences are important in that they produce much of the rise in earnings dispersion over the lifetime, given our estimates of the magnitude of human capital risk.

We use our estimates of shocks and initial conditions to quantify the importance of different proximate sources of lifetime inequality. We find that as of a real-life age of 20 differences in initial conditions are more important than are shocks received over the remaining lifetime as a source of variation in realized lifetime utility, lifetime earnings and lifetime wealth.⁴ We find that between 62 to 73 percent of the variation in lifetime utility and between 60 to 71 percent of the variation in lifetime earnings is due to variation in initial conditions. The higher estimate for each statistic applies when the magnitude of shocks is set to our lowest point estimate, whereas the lower estimate applies when the magnitude of shocks is set to our highest point estimate. Intuitively, the greater the shock variance the smaller is the role for initial conditions in accounting for the pattern of increasing earnings dispersion over the lifetime.

Among initial conditions, we find that, as of age 20, variation in initial human capital is substantially more important than variation in either learning ability or initial wealth for how an agent fares in life. This analysis is conducted for an agent with the median value of each initial condition. We find that a one standard deviation increase in initial wealth increases

²This mechanism is supported by the literature (see Card (1999)) on the shape of the mean age-earnings profiles by years of education. It is also supported by the work of Lillard and Weiss (1979), Baker (1997) and Guvenen (2006). They estimate a statistical model of earnings and find important permanent differences in individual earnings growth rates.

³Since a measure of financial wealth is observable, we choose the tri-variate initial distribution to be consistent with features of the distribution of wealth for young households.

⁴Lifetime earnings equals the present value of earnings, whereas lifetime wealth equals lifetime earnings plus initial wealth.

expected lifetime wealth by 3 to 4 percent. In contrast, a one standard deviation increase in learning ability or initial human capital increases expected lifetime wealth by 9 to 10 percent and 30 to 34 percent, respectively. We also analyze how an agent in the model values these changes in initial conditions. Specifically, we ask what is the permanent percentage change in consumption which is equivalent for an agent in expected utility terms to these changes in initial conditions. We find that the equivalent percentage changes in consumption are roughly in line with how a change in initial condition impacts, in percentage terms, expected lifetime wealth.

A leading and alternative view of lifetime inequality to the one analyzed in this paper is presented in Storesletten et. al. (2004). The model analyzed in that paper is a standard, incomplete-markets model in which labor earnings over the lifetime is exogenous.⁵ These authors estimate an earnings process from U.S. panel data to match features of earnings over the lifetime. Within their model, slightly less than half of the variation in realized lifetime utility is due to differences in initial conditions.⁶

We note three difficulties related to this alternative incomplete-markets view. First, the importance of idiosyncratic earnings risk may be overstated. The reason is that all of the rise in earnings dispersion with age is attributed to shocks and none to initial conditions. In our model learning ability differences lead to systematic differences in earnings growth rates across agents. Lillard and Weiss (1979), Baker (1997) and Guvenen (2006) provide evidence for such differences in permanent earnings growth rates in male earnings data. Second, although the incomplete-markets model with exogenous earnings produces the rise in U.S. within cohort consumption dispersion over the period 1980-90 documented by Deaton and Paxson (1994), the rise in consumption dispersion is substantially smaller in U.S. data over a longer time period. Our model produces less of a rise in consumption dispersion than the exogenous-earnings model. A key reason for this is that part of the rise in earnings dispersion is due to initial conditions. This component is anticipated by agents and therefore reflected in consumption dispersion early in life. Finally, the standard incomplete-market, life-cycle model is not useful for some purposes. Specifically, since earnings are exogenous, the model gives up on theorizing about the underlying sources of earnings inequality. Thus, the model

⁵Similar models have been used in the macroeconomic literature on economic inequality. Recent papers in this literature include Huggett (1996), Castañeda, Diaz-Jimenez and Rios-Rull (2003), Krueger and Perri (2006), Guvenen (2006) and Heathcote, Storesletten and Violante (2006), among many others.

⁶In the context of a career-choice model, Keane and Wolpin (1997) find a more important role for initial conditions. They find that unobserved heterogeneity realized at age 16 accounts for about 90 percent of the variance in lifetime utility.

can not shed light on how policy may affect inequality in lifetime earnings or may affect welfare through earnings. Models with exogenous wage rates (e.g. Heathcote et. al. (2006)) face this criticism, but to a lesser extent, since most earnings variation is attributed to wage variation. In our view, it is worthwhile to pursue a more fundamental approach that in essence endogenizes wage rate differences via human capital theory.

The paper is organized as follows. Section 2 presents the model. Section 3 documents earnings distribution facts and estimates properties of shocks. Section 4 sets model parameters. Section 5 analyzes the model. Section 6 analyzes sources of lifetime inequality. Section 7 concludes.

2 The Model

We add risky human capital to the life-cycle, permanent-income framework.⁷ An agent's preferences over consumption allocations are determined by a calculation of expected utility as indicated below. Consumption $c_j(z^j)$ at age j is risky as it depends on the j -period history of human capital shocks z^j . The set of possible j -period histories is denoted $Z^j \equiv \{z^j = (z_1, \dots, z_j) : z_i \in Z, i = 1, \dots, j\}$, where Z is a finite set of possible shock realizations. $P(z^j)$ denotes the probability of history z^j .

$$E\left[\sum_{j=1}^J \beta^{j-1} u(c_j)\right] = \sum_{j=1}^J \sum_{z^j \in Z^j} \beta^{j-1} u(c_j(z^j)) P(z^j)$$

An agent solves the decision problem below, taking initial financial wealth $k_1(1+r)$, initial human capital h_1 and learning ability a as given.

$$\max_{\{c_j, l_j, k_{j+1}\}} E\left[\sum_{j=1}^J \beta^{j-1} u(c_j)\right]$$

subject to

$$(1) \quad c_j + k_{j+1} = k_j(1+r) + e_j, \forall j \text{ and } k_{J+1} \geq 0$$

⁷The model generalizes Ben-Porath (1967) to allow for risky human capital. Risky human capital is modeled by extending the two-period models of Levhari and Weiss (1974) and Eaton and Rosen (1980) to a multi-period setting. Krebs (2004) also analyzes a multi-period model of human capital with idiosyncratic risk. Our work differs by its focus on lifetime inequality, among other differences.

(2) $e_j = R_j h_j L_j$ if $j < J_R$, and $e_j = 0$ otherwise.

(3) $h_{j+1} = z_{j+1} F(h_j, l_j, a), \forall j$ and $L_j + l_j = 1, \forall j$

In this decision problem an agent faces a period budget constraint in which consumption c_j plus financial asset holding k_{j+1} equal earnings plus the value of assets brought into the period. Financial assets pay a risk-free, real return r . Earnings e_j before a retirement age J_R equal the product of a human capital rental rate R_j , an agent's human capital h_j and the fraction L_j of available time put into market work. Earnings are zero at and after the retirement age J_R . An agent's future human capital is determined by an idiosyncratic shock z_{j+1} multiplying the law of motion for human capital F . The law of motion F depends upon current human capital h_j , time devoted to human capital production l_j and an agent's learning ability a , and is increasing in its three arguments.

We now comment on three key features of the model. First, while the earnings of an agent are stochastic, the earnings distribution for a large cohort of agents evolves deterministically. This occurs because the model has idiosyncratic but no aggregate risk.⁸ Second, the model has two sources of growth in earnings dispersion within cohort - agents have different learning abilities and different shock realizations. The next section characterizes empirically the rise in US earnings dispersion. Third, the model implies that the nature of human capital shocks can be identified from wage rate data, independently from all other model parameters. This holds, as an approximation, towards the end of the working life because the model implies that the production of human capital goes to zero. The next section develops the logic of this point.

3 Data and Empirical Analysis

In this section we use data to address two issues. First, we characterize how mean earnings and measures of earnings dispersion and skewness evolve with age for a cohort. Second, we estimate a process for human capital shocks from wage rate data.

⁸More specifically, $P(z^j)$ is both the probability that an agent receives a j-period shock history z^j and the fraction of the agents in a cohort that receive this shock history.

3.1 Age Profiles

The age profiles are based on earnings data from the Panel Study of Income Dynamics (PSID) 1969-2004 family files. We utilize earnings of males who are the head of the household, who work between 520 and 5820 hours per year and who earn at least 2000 dollars (in 1968 prices). We consider males between the ages of 21 and 62. These selection criteria are motivated by several considerations. First, the PSID has many observations in the middle but relatively fewer at the beginning or end of the working life cycle. By focusing on ages 21-62, we have at least 100 observations in each age-year bin with which to calculate age and year-specific earnings statistics. Our age bins are centered 5-year age bins. For each year we therefore have bins for ages 23- 60. Second, near the traditional retirement age there is a substantial fall in labor force participation that occurs for reasons that are abstracted from in the model. This suggests the use of a terminal age that is earlier than the traditional retirement age.

Let $e_{j,t}$ be the mean real earnings of agents who are age j at time t .⁹ The earnings data can be viewed as being generated by several factors that we name cohort, time, and age effects. Ultimately, we are interested in the age effect. However, as described in detail below, this measure depends on the identifying assumptions regarding cohort and time effects. To introduce notation, we denote a birth cohort as $s = t - j$ that is agents who were born in year $t - j$. We assume that $e_{j,t}$ is determined by cohort effects α_s , age effects β_j , time effects γ_t and shocks $\epsilon_{j,t}$. The relationship between these variables is given below both in levels and in logs, where the latter is denoted by a tilde. Cohort effects can be viewed as effects that are common to all agents who were born in a particular year (e.g., those who were born in the Great Depression may have suffered a permanent adverse shock). Time effects can be viewed as effects that are common to all individuals alive at a point in time. An example would be a temporary rise in the rental rate of human capital that increases the earnings of all individuals in the period.

$$e_{j,t} = \alpha_s \beta_j \gamma_t \epsilon_{j,t}$$

$$\tilde{e}_{j,t} = \tilde{\alpha}_s + \tilde{\beta}_j + \tilde{\gamma}_t + \tilde{\epsilon}_{j,t}$$

⁹Real values are calculated using the CPI. To calculate $e_{j,t}$ we use a 5 year bin centered at age j . For example, to calculate mean earnings of agents age $j = 30$ in year $t = 1980$ we use data on agents age 28 – 32 in 1980.

The linear relationship between time t , age j , and birth cohort $s = t - j$ limits the applicability of this regression specification. Specifically, without further restrictions the regressors in this system are co-linear and these effects cannot be estimated. This identification problem is well known.¹⁰ In effect any trend in the data can be arbitrarily reinterpreted as due to year (time) effects or alternatively as due to age or cohort effects.

Given this problem, we provide two alternative measures of the age effects. These correspond to the cohort effects case where we set $\tilde{\gamma}_t = 0, \forall t$ and the time effects case where we set $\tilde{\alpha}_s = 0, \forall s$. We use ordinary least squares to estimate the coefficients. For the cohort effects case, the regression has $J \times T$ dependent variables regressed on $J + T$ cohort dummies and J age dummies. T and J denote the number of time periods in the panel and the number of distinct age groups, which in our case equal $J = 60 - 23$ and $T = 2004 - 1969$. For the time effects case the regression has $J \times T$ dependent variables regressed on T time dummies and J age dummies. This regression has J less regressors than the regression incorporating cohort effects.¹¹

In Figure 1 we graph the age effects of the levels of earnings implied by each regression. Figure 1 highlights the familiar hump-shaped profile of mean earnings. Figure 1 is constructed by plotting β_j from each regression above. The age effects β_j are scaled so that mean earnings equal 100 at the end of the working life cycle for the case of time effects.

A similar analysis is carried out to extract the age profile of measures of earnings dispersion and skewness. We consider two standard measures of dispersion: the variance of log earnings and the Gini coefficient of earnings. We measure skewness by the ratio of mean earnings to median earnings.

For each of these three statistics the procedure is the same. Let $stat_{j,t}$ denote the earnings statistic of interest calculated for age group j at time t .¹² We then estimate the dummy variable coefficients in the regression equation below, setting either $\alpha_s^{stat} = 0$ for the case of time effects or $\gamma_t^{stat} = 0$ for the case of cohort effects. Figure 2 (a)-(c) plots the estimated age dummy coefficients after normalization so that the age profile of each earnings statistic

¹⁰See Weiss and Lillard (1978) and Deaton and Paxson (1994) among others.

¹¹A third approach, discussed in more detail in Huggett et. al. (2006), allows for age, cohort and time effects but with the restriction that time effects are mean zero and are orthogonal to a time trend. That is $(1/T) \sum_{t=1}^T \tilde{\gamma}_t = 0$ and $(1/T) \sum_{t=1}^T \tilde{\gamma}_t t = 0$. Thus, trends over time are attributed to cohort and age effects rather than time effects. The results of this approach are effectively the same as those for cohort effects and we therefore omit them for brevity.

¹²We, again, use 5-year age bins centered at age j to compute the statistic of interest at age j .

runs through the mean value of this statistic across years at age $j = 38$.

$$stat_{j,t} = \alpha_s^{stat} + \beta_j^{stat} + \gamma_t^{stat} + \epsilon_{j,t}^{stat}$$

Figure 2 shows that both dispersion and skewness tend to increase with age in both the time and cohort effects views. The cohort effect view in Figure 2(a) implies a rise in the variance of log earnings of about 0.4 from age 23 to 60 whereas the time effects imply a smaller rise of only about 0.2. The same qualitative pattern can be seen for the Gini coefficient measure of dispersion in Figure 2(b). Figure 2(c) shows that the rise in earnings skewness with age is also greater for the cohort effect view than for the time effects view.

We will ask the model to best match the time effects view of the evolution of all three features (i.e mean, dispersion and skewness) of the earnings distribution. Heathcote et. al. (2005a) present an argument for stressing the time effects view. Their argument is based in part on the fact that within-cohort and within-age group changes in earnings and wage dispersion vary over time but that these two changes are of similar magnitude over the sub-periods they examine. They argue that this fact suggests that time effects are key.

3.2 Human Capital Shocks

The model implies that an agent's wage rate, measured as market compensation per unit of work time, equals the product of the rental rate and an agent's human capital. The model also implies that late in the working life cycle human capital investments are approximately zero. This occurs as the number of working periods over which the agent can reap the returns to these investments falls as the agent approaches retirement. The upshot is that when there is no human capital investment over a period of time, then the change in an agent's wage rate is entirely determined by rental rates and the human capital shock process and not by any other model parameters.¹³

This logic is restated in the equations below. The first equation indicates how the wage w_{t+s} is determined by rental rates R_{t+s} and shocks z_{t+s} in the absence of human capital investment. Here it is assumed that there is no human capital investment from period t to $t + s$ so that $F(h, 0, a) = h$ in all periods with no investment. The second equation takes

¹³Heckman, Lochner and Taber (1998) use a similar line of reasoning to estimate differences in rental rates across skill groups within a model which abstracts from idiosyncratic risk.

logs of the first equation, where a hat denotes the log of a variable.

$$w_{t+s} \equiv R_{t+s} h_{t+s} = R_{t+s} z_{t+s} F(h_{t+s-1}, 0, a) = R_{t+s} z_{t+s} \times \dots \times z_{t+1} h_t$$

$$\hat{w}_{t+s} \equiv \log w_{t+s} = \hat{R}_{t+s} + \sum_{j=1}^s \hat{z}_{t+j} + \hat{h}_t$$

Now let measured s-period log wage differences (denoted $y_{t,s}$) be true differences plus measurement error differences $\epsilon_{t+s} - \epsilon_t$. This is the first equation below. We assume that human capital shocks and measurement errors (\hat{z}_t, ϵ_t) are jointly independent and are identically distributed over time and people. We also assume that $\hat{z}_t \sim N(\mu, \sigma^2)$ and that $Var(\epsilon_t) = \sigma_\epsilon^2$. These assumptions imply the three cross-sectional moment conditions below.

$$y_{t,s} \equiv \hat{w}_{t+s} - \hat{w}_t + \epsilon_{t+s} - \epsilon_t = \hat{R}_{t+s} - \hat{R}_t + \sum_{j=1}^s \hat{z}_{t+j} + \epsilon_{t+s} - \epsilon_t$$

$$E[y_{t,s}] = \hat{R}_{t+s} - \hat{R}_t + s\mu$$

$$Var(y_{t,s}) = s\sigma^2 + 2\sigma_\epsilon^2$$

$$Cov(y_{t,s}, y_{t,r}) = r\sigma^2 + \sigma_\epsilon^2 \text{ for } r < s$$

To make use of these moment restrictions, one needs to be able to measure the variable $y_{t,s}$ and to have individuals for which the assumption of no time spent accumulating human capital is a reasonable approximation. The focus on older workers addresses both issues. Wage data for younger workers are potentially problematic for both issues. Specifically, on the first issue it may be difficult to accurately measure the wage rates emphasized in the model when measured time at work is a mix of work time and learning time.

We calculate wages in PSID data as total male labor earnings divided by total hours for male head of household. We impose the same selection criteria as those presented in Section 3.1 for earnings. We follow males for either three years or four years. Thus, we calculate two log wage differences (i.e. $y_{t,s}$ for $s = 1, 2$) when males are followed for three years and three log wage differences when males are followed for four years.¹⁴ In estimation we use

¹⁴The PSID data is not available for the years 1997, 1999, 2001, and 2003. In the years preceding those

cross sectional variances and covariances aggregated across panel years. For each year we generate the sample analog to the moments: $\mu_{t,s} \equiv \frac{1}{N_t} \sum_{i=1}^{N_t} y_{t,s}^i$ and $\frac{1}{N_t} \sum_{i=1}^{N_t} (y_{t,s}^i - \mu_{t,s})^2$ and $\frac{1}{N_t} \sum_{i=1}^{N_t} (y_{t,s}^i - \mu_{t,s})(y_{t,r}^i - \mu_{t,r})$. We stack the moments across the panel years and use a 2-step GMM estimation with an identity matrix as the initial weighting matrix.

Table 1 provides the estimation results. Over the entire sample period the point estimate of the standard deviation σ of the log shock to human capital is between 0.10 and 0.11 for both the age group 50 – 60 and 55 – 65. This holds when we follow males for three years ($\bar{s} = 2$) or for four years ($\bar{s} = 3$).¹⁵ This is smaller than the estimate which obtains when all males 23 – 60 are pooled together. From the point of view of human capital theory, pooling younger and older workers will mean that the change in log wages will be determined by shocks and by the endogenous response of human capital decisions to shocks and initial conditions rather than by shocks alone. We note that the magnitude of the persistent wage shocks estimated by Heathcote et. al. (2006) in PSID data when various age groups are pooled is between our estimates for the older age groups and our estimate for the pooled sample.

The lower part of Table 1 contains results for the period 1969-1981. The point estimates of the shock variance is lower for each age group in this period than the estimates obtained using the entire time period 1969-2004. This is consistent with the fact that the increase in cross-sectional earnings inequality in the U.S. over this period occurred mainly after 1981. Over the period 1969-1981 we find that the estimated variances for the 50-60 and the 55-65 age groups are smaller than the estimated variance when all age groups are pooled. This is the same pattern that was found over the entire sample period.

4 Setting Model Parameters

The strategy for setting model parameters is in three steps. First, we estimate the parameters governing human capital shocks directly. This was done in the previous section. Second, we choose parameters governing the utility function and interest rates based upon previous studies. Third, we set the parameters governing the distribution of initial conditions and the parameter governing the elasticity of the human capital production function so that

years we impose that the agent is available for three consecutive years and use a two year growth rate.

¹⁵We have also analyzed several other age-panel year configurations in order to gauge the potential sensitivity to proximity to retirement years and have found no material difference in the point estimates.

the model best matches the age profiles of the evolution of the male earnings distribution estimated in the previous section. In choosing the initial distribution and elasticity, we take all other model parameters as given.

Model parameter values are summarized in Table 2. We set the model period to be a year. Agents live $J = 56$ model periods or from a real-life age of 20 to 75. We set a retirement age at $J_R = 42$ or at a real-life age of 61. At the retirement period an agent can no longer engage in market work. The real interest rate in the model is set to $r = 0.042$. This is the average of the annual return to stock and long-term bonds over the period 1946-2001 (see Siegel (2002, Table 1-1 and 1-2)). The discount factor is set to $\beta = 1/(1+r)$ so that absent risk the consumption profile solving the model is flat.

The utility function is of the constant relative risk aversion class. The parameter ρ governing risk aversion and intertemporal substitution is set to $\rho = 2$. This value lies in the middle of the range of estimates based upon micro-level data which are surveyed by Browning, Hansen and Heckman (1999, Table 3.1).

We set the value g governing the growth in the rental rate of human capital in the model equal to the average growth rate of mean male earnings in US cross-section data. We calculate that in the PSID over the period 1968-2001 the mean arithmetic growth rate of mean male earnings equals 0.19 percent. The benchmark model, with homothetic preferences, implies that the earnings distribution of different cohorts is proportional to the initial level of this rental rate, other things equal. Thus, with stable demographics the average cross-sectional earnings in the model grows at rate g .

We set the standard deviation σ of the log human capital shocks to be consistent with the estimates in Table 1. We analyze $\sigma = 0.088$ and $\sigma = 0.108$. These are respectively the lowest and the highest point estimates from Table 1 for the 55 to 65 age group. We set μ , governing the mean log human capital shock, so that the model matches the average rate of decline of mean earnings for the cohorts of older workers in US data that we documented earlier in Figure 1. The fall in mean earnings in the model equals $(1+g)e^{\mu+\sigma^2/2}$ when agents make no human capital investments. Thus, μ is set, given the value g and σ , so that this holds.

We assume that the human capital production function is $F(h, l, a) = h + a(hl)^\alpha$, which is the functional form analyzed in Ben-Porath (1967). We set the elasticity parameter α and the parameters governing the initial distribution G of human capital and learning ability to

best match the three central features of the U.S. earnings distribution documented in section 3. The Appendix describes the distance metric between data and model statistics that we minimize. We restrict the distribution G to lie in a parametric class. In the benchmark model, initial assets are zero and initial human capital and learning ability are jointly log-normally distributed so that $\log(x) \sim N(\mu_x, \Sigma)$, when $x = (h_1, a)$. We explore later in the paper a tri-variate distribution, where the initial asset distribution matches features of net wealth holdings for young households in the PSID.

When we examine the low shock case (i.e. $\sigma = 0.088$), we find that $\alpha = 0.675$ is the parameter value which best matches the earnings facts. For the high shock case, we find that $\alpha = 0.625$ best matches the data. We note that the parameter α has been estimated in the human capital literature. These estimates, surveyed by Browning et. al. (1999, Table 2.3- 2.4), lie in the range 0.5 to just over 0.9. These previous estimates, however, are based upon models that abstract from idiosyncratic risk.

We have examined the fit of the model at prespecified values of the parameter α , while choosing the parameters of the initial distribution to best match the U.S. earnings facts. The distance between model and data statistics displays a U-shaped pattern in the parameter α , where the bottom of the U is the value in Table 2. This distance increases sharply for values of α exceeding 0.80. Over this range there is a strong tension between fitting the dispersion profile and the skewness profile. If one were to ignore the skewness profile in choosing initial conditions, then the model would substantially overstate the rise in skewness over the lifetime we find in U.S. data.

5 Earnings in the Model

In this section, we report on the ability of the model to reproduce the earnings facts documented in section 3.¹⁶ The analysis focuses on the benchmark model without initial wealth differences.

¹⁶Methods used to compute solutions to the model are described in the Appendix.

5.1 Dynamics of the Earnings Distribution

The age profiles of mean earnings, earnings dispersion and skewness produced by the benchmark model are displayed in Figure 3(a)-(c). The model generates the hump-shaped earnings profile for a cohort by a standard human capital accumulation argument. Early in the working life cycle, individuals devote more time to human capital production than at later ages. These time allocation decisions lead to a net accumulation of human capital in the early part of the working life cycle. Thus, mean earnings increase with age as human capital and mean time worked increase with age.

Towards the end of the working life-cycle, mean human capital levels fall. This happens as the mean multiplicative shock to human capital is smaller than one (i.e. $E[z] = e^{\mu+\sigma^2/2} < 1$). This corresponds to the notion that on average human capital depreciates. The implication is that average earnings in Figure 3 fall late in life because growth in the rental rate of human capital is not enough to offset the mean fall in human capital.

Figure 4 shows the age profile of the mean fraction of time allocated to human capital production in the model. Approximately 25 percent of available time is directed at human capital production early in life. After age 55 less than 5 percent of time is directed at human capital production. Recall that for the purpose of identifying human capital shocks, we use the assumption that time devoted to human capital production and, hence, human capital production is negligible towards the end of the working lifetime.¹⁷

Two forces account for the rise in earnings dispersion in Figure 3. First, since individual human capital is repeatedly hit by shocks, these shocks are a source of increasing dispersion in human capital and earnings as a cohort ages. Second, differences in learning ability across agents produce mean earnings profiles with different slopes. This follows since within an age group, agents with high learning ability choose to produce more human capital and devote more time to human capital production than their low ability counterparts. Huggett et. al. (2006, Proposition 1) establish that this holds in the absence of human capital risk. This mechanism implies that earnings of high ability individuals are relatively low early in life and relatively high late in life compared to agents with lower learning ability, holding initial human capital equal.

¹⁷Mincer (1997) reviews evidence related to time allocation decisions. He cites evidence that time directed at skill accumulation decreases with age for U.S. males.

5.2 Earnings Dispersion: Risk versus Ability Differences

We now try to understand the quantitative importance of risk and ability differences for producing the increase in earnings dispersion displayed in Figure 3. We do so by alternatively eliminating ability differences or eliminating shocks. The analysis focuses on the high shock case where $\sigma = 0.108$.

5.2.1 Eliminating Ability Differences

We eliminate ability differences by changing the initial distribution so that all agents have the same learning ability, which we set equal to mean ability. In the process of changing learning ability, we do not alter any agent's initial human capital.

Figure 5(a) shows that eliminating ability differences leads to the striking result that the rise in earnings dispersion over the lifetime is almost completely eliminated. This result is due to two opposing forces. First, human capital risk leads ex-ante identical agents to differ ex-post in human capital and earnings. Second, the model has a force which leads to decreasing dispersion in human capital and earnings with age which has received little attention in work which interprets patterns of earnings dispersion over the lifetime. Without risk and without ability differences, all agents within an age group produce the same amount of new human capital regardless of the current level of human capital – see Huggett et. al. (2006, Proposition 1). This holds for any value of the elasticity parameter α of the human capital production function. This implies that both the distribution of human capital and earnings are Lorenz ordered by age. Thus, measures of earnings or human capital dispersion that respect the Lorenz order decrease for a cohort as the cohort ages.

Figure 5(a) shows that earnings dispersion increases at the end of the working lifetime. This occurs as human capital production at the end of life goes to zero because the time allocated to production (see Figure 4) goes to zero. This means that the opposing force leading to convergence is gradually eliminated with age.

5.2.2 Eliminating Idiosyncratic Risk

To highlight the role of human capital risk, we eliminate idiosyncratic risk altogether by setting $\sigma = 0$. We adjust the mean log shock μ to keep the mean shock level constant but

maintain all other initial conditions. This analysis is used in the next section to understand the effects we observe when we change the shocks and at the same time allow the model to refit the initial conditions. Removing idiosyncratic risk leads to a counter-clockwise rotation of the mean earnings profile and leads to a U-shaped earnings dispersion profile. Figure 5(b) shows the effect on earnings dispersion of eliminating risk.

When idiosyncratic risk is eliminated, human capital accumulation becomes more attractive for risk-averse agents. Thus, all else equal, agents spend a greater fraction of time accumulating human capital early in life. The result is a counter-clockwise movement in the mean earnings profile.¹⁸ In terms of dispersion in labor earnings, human capital shocks are more important for agents with relatively high learning ability. These agents are the ones who would allocate an even larger fraction of time into human capital accumulation for lower values of the variance of idiosyncratic shocks. When human capital risk is eliminated, these agents allocate less time to work early in life and more time to human capital accumulation.

5.3 Properties of the Initial Distribution

Table 3 summarizes properties of the distribution of initial conditions. These properties are given both for the highest and lowest shock estimates and for the two estimated values of the elasticity parameter α . When the shock variance increases, the initial distributions that best reproduce the earnings facts require higher levels of mean learning ability and lower levels of ability dispersion and human capital dispersion. A consequence of the lower levels of ability and human capital dispersion is a reduction in the relative importance of initial conditions for lifetime inequality. We will see this shortly in the next section of the paper.

What accounts for these changes in the initial distributions? Recall from our previous analysis that eliminating shocks for a given initial distribution and elasticity parameter leads to a counter-clockwise rotation of the mean earnings profile. This occurs because the time input into human capital accumulation over the life cycle increases as human capital risk decreases. This is consistent with the result of Levhari and Weiss (1974) whereby, in a two-period model, risk-averse agents reduce human capital investments with human capital risk compared to the no risk case.

¹⁸This is effectively the central result of Levhari and Weiss (1974) extended to a multi-period setting. They showed in a two-period model that time input into human capital production is smaller with human capital risk than without when agents are risk averse.

Following this intuition, to produce the earnings facts as risk increases, holding the elasticity parameter fixed, the distribution of initial conditions needs to be adjusted. A higher mean learning ability level leads to a counter-clockwise rotation of the mean earnings profile to counteract the clockwise rotation of the mean earnings profile produced by adding risk to the model with fixed initial distribution. The intuition for why ability dispersion falls as human capital risk increases is that human capital risk is itself a source of increased earnings dispersion. Thus, greater human capital risk leaves less room for ability differences in accounting for the rise in earnings dispersion with age.

6 Lifetime Inequality

6.1 Initial Conditions Versus Shocks

We decompose the variance in lifetime inequality into variation due to initial conditions versus variation due to shocks. This is done for lifetime utility and lifetime earnings in the benchmark model. Later on we also decompose the variance in lifetime wealth. Lifetime wealth equals the realized present value of earnings (i.e. lifetime earnings) plus initial wealth.¹⁹ Such a decomposition makes use of the fact that a random variable can be written as the sum of its conditional mean plus the variation from its conditional mean. As these two components are orthogonal, the total variance equals the sum of the variance in the conditional mean plus the variance around the conditional mean.

Table 4 presents lifetime inequality within the benchmark model. Lifetime inequality is analyzed as of the start of the working life cycle, which we set to a real-life age of 20. We find that from 62 to 73 percent of the variation in lifetime utility and from 60 to 71 percent of the variation in lifetime earnings is due to initial conditions. For each statistic, initial conditions have a greater role when the the magnitude of shocks is set to our lowest point estimate.

Figure 6 describes lifetime inequality as the elasticity parameter α of the human capital production function is varied over the interval [.5, .9]. This interval includes the values, $\alpha = .625$ and $\alpha = .675$, that best match the earnings profiles in the high and low shock cases.

¹⁹Lifetime utility and lifetime wealth along a lifetime shock history z^J are defined as follows: $U(z^J; h_1, k_1, a) = \sum_{j=1}^J \beta^{j-1} u(c_j(z^J; h_1, k_1, a))$ and $W(z^J; h_1, k_1, a) = k_1(1+r) + \sum_{j=1}^J e_j(z^J; h_1, k_1, a)/(1+r)^{j-1}$.

Figure 6 shows that the fraction of the variance in lifetime utility and lifetime earnings that is due to initial conditions falls as the elasticity parameter increases.

We now determine how the decomposition of lifetime inequality changes when we account for variation in initial wealth found in U.S. data. To examine this issue, we use PSID net-wealth data for households with a male head age 20 to 25.²⁰ We express net wealth as a ratio to mean male earnings in the age group 20 -25 in each year. We then pool these ratios across years.

We maintain the multi-variate log-normal structure for describing initial conditions. However, we do allow for negative wealth holding. Specifically, we approximate the empirical pooled wealth distribution with a lognormal distribution which is shifted a distance δ . We choose δ so that 95 percent of the distribution has a wealth to mean earnings ratio above $-\delta$. The distribution of the wealth-earnings ratio in the model is given by $e^x - \delta$, where x is distributed $N(\mu_1, \sigma_1^2)$. The parameters (μ_1, σ_1^2) are set equal to the sample mean and sample variance of the log of the sum of the wealth-earnings ratio plus δ for ratios above $-\delta$. The median, mean and standard deviation of the wealth-earnings ratio in the model is then $(0.377, 0.778, 1.340)$.²¹ This implies that there is a substantial amount of initial wealth dispersion within the model. Specifically, a one standard deviation change in initial wealth is 1.34 times mean yearly earnings for young agents.

The distribution of initial wealth, human capital and learning ability is selected to best match the earnings facts documented earlier when α is set to the corresponding value in Table 2. The distribution is a tri-variate lognormal, where the parameters describing the mean and variance of shifted log wealth are those calculated above in U.S. data. Thus, wealth in the model is right skewed and mean wealth is more than double median wealth.

Table 5 analyzes lifetime inequality when initial wealth differences are set to the magnitudes we find in U.S. data. We find that initial conditions account for 66 to 77 percent of the variation in lifetime utility, 61 to 72 percent of the variation in lifetime wealth and 59 to 71 percent of the variation in lifetime earnings.²² Thus, we find that when we account for

²⁰The data is from the PSID wealth supplement for 1984, 1989, 1994, 1999, 2001 and 2003. The sample size is 1176 when pooled across these years.

²¹In the PSID sample we calculate that $(\mu_1, \sigma_1^2, \delta) = (-0.277, 0.849, 0.381)$ and that the median, mean and standard deviation of the wealth-earnings ratio is $(0.313, 0.776, 1.432)$.

²²Our results for lifetime earnings inequality are close to some existing results that are based upon a statistical model of earnings. Such a model does not attempt to explain how earnings arise endogenously from optimal decisions over the life cycle. For example, Geweke and Keane (2000, Table 11) estimate such

initial wealth differences the majority of the variation in lifetime inequality is due to initial conditions.

6.2 How Important are Different Initial Conditions?

The analysis so far has not addressed how important variation in one type of initial condition is compared to variation in other types for how an agent fares in life. We analyze the importance of different initial conditions by asking the agents in the model how much compensation is equivalent to starting life with a one standard deviation change in any initial condition. We express this compensation, which we call an equivalent variation, in terms of the percentage change in consumption in all periods that would be required to leave an agent with the same expected lifetime utility as an agent with a one standard deviation change in the relevant initial condition. The baseline initial condition is set equal to the mean log values of initial human capital and learning ability and equal to the mean of the shifted log initial wealth. The changes in initial conditions are also in standard deviations of log variables.

The importance of changes in initial conditions, stated in terms of equivalent variations, is presented in the upper panel of Table 6. We find that a one standard deviation movement in log human capital is substantially more important than a one standard deviation movement in either log learning ability or log initial wealth. A one standard deviation increase in initial human capital is equivalent to a 31 – 34 percent increase in consumption. In contrast, a one standard deviation increase in learning ability or initial wealth is equivalent to an 6 – 8 percent and 4 – 5 percent increase in consumption, respectively. Thus, we find that an increase in human capital leads to the largest impact, an increase in learning ability has the next largest impact and an increase in initial financial wealth has the smallest impact on equivalent variations.

We also analyze the importance of different initial conditions by determining how changes in initial conditions affect an agent’s budget constraint. More specifically, we determine the percent by which an agent’s expected lifetime wealth changes in response to a one standard deviation change in an initial condition. The lower panel of Table 6 presents the results of this analysis. In interpreting these results, it is useful to keep two points in mind. First, an

a model of male earnings using PSID data and find that initial conditions account for between 66 and 72 percent of the variation in the simulated present value of earnings from age 25 to 65.

increase in human capital acts as a vertical shift of the expected earnings profile, whereas an increase in learning ability rotates this profile counter clockwise. Second, the impact of additional initial financial wealth is both through a direct impact on lifetime resources as well as the indirect impact through earnings.

Broadly speaking, the lower panel of Table 6 finds that the impact on expected lifetime wealth of changes in initial conditions are roughly in line with their impact on equivalent variations. Thus, a one standard deviation change in initial human capital has the greatest impact, the corresponding change in learning ability has the next biggest impact and initial wealth changes have the least important impact on expected lifetime wealth.²³ We note, however, that the impact of a one standard deviation change in learning ability on equivalent variations and lifetime wealth displays a weaker link than for other initial conditions. Specifically, an increase in learning ability raises expected lifetime wealth substantially more than it raises an agent's equivalent variation. Intuitively, this occurs because higher ability leads to higher mean earnings and a higher earnings variance later in life but lowers earnings early in life. Thus, with incomplete insurance markets, a risk-averse agent values such an increase in expected lifetime wealth at less than the equivalent change in current wealth.

6.3 Consumption and Social Insurance Implications

One may argue that a useful model for analyzing lifetime inequality within an incomplete-markets framework should also be broadly consistent in terms of its implications for consumption inequality. We therefore compare the model's implications for the rise in consumption dispersion over the lifetime with the patterns found in U.S. consumption data.

A number of studies analyze the variance of log adult-equivalent consumption in U.S. data. These studies regress the variance of log adult-equivalent consumption for households in different age groups on age and time dummies or alternatively on age and cohort dummies. The coefficients on age dummies are then used to highlight how consumption dispersion varies for a cohort with age.

Figure 7 plots the variance of log adult-equivalent consumption in U.S. data from two

²³We also calculate expected lifetime wealth elasticities for an increase in each initial condition. At the benchmark initial condition, the expected lifetime wealth elasticities are .692 to .723 for human capital, .356 to .414 for learning ability and .026 to .031 for initial wealth. For each initial condition, the larger (smaller) elasticity corresponds to the high (low) shock case.

such studies. Deaton and Paxson (1994) analyze U.S. Consumer Expenditure Survey (CEX) data from 1980 to 1990. Heathcote et. al. (2005a), Slesnick and Ulker (2005) and Primiceri and van Rens (2006) reexamine this issue using CEX data over a longer time period. All three of these later studies find that the rise in dispersion with age is substantially smaller than the rise in Deaton and Paxson (1994).

The exogenous earnings model analyzed in Storesletten et. al. (2004) produces the rise in consumption dispersion documented in Deaton and Paxson (1994). This is the case when their exogenous earnings model has a social insurance system. We analyze the consumption implications of our benchmark risky human capital model with a social insurance system.

Specifically, we augment our benchmark model to include a social security system and an income tax system. The model social security system features a proportional earnings tax of 10.6 percent, which is the old-age and survivors insurance benefit tax rate in the US social security system. The social security system has a common retirement benefit paid to all agents after the retirement age set equal 45 percent of mean earnings in the last period of the working lifetime. The income tax in the model captures the pattern of effective average federal tax rates in the US documented in Congressional Budget Office (2004, Table 3A and 4A) for the tax year 2001. Effective average federal tax rates in the US rise from approximately 0 percent for low income households to approximately 20 percent for very high income households.²⁴ We set the distribution of initial conditions so that the model with social insurance and no initial wealth differences best matches the U.S. earnings facts, setting all other model parameters to the values in Table 2.

Figure 7 shows that the rise in consumption dispersion in our model is less than the rise in Deaton and Paxson (1994). The rise in the model from age 25 to 60 is approximately 7 log points when shocks are set to our highest point estimate and is approximately 6 log points when shocks are set to the lowest point estimate. The rise over the life cycle found by Primiceri and van Rens (2006), using CEX data from 1980 to 2000, is very similar to the model results for both the high and low shock case.²⁵

To better understand this result, we emphasize that our risky human capital model decouples the rise in consumption dispersion from the rise in earnings dispersion. At one

²⁴The details for implementing this income tax function within our model follows closely Huggett and Parra (2006).

²⁵The benchmark model without a social insurance system produces a rise of approximately 16 and 12 log points from age 25 to 60 in the high and low shock cases, respectively.

extreme, the model can produce the rise in earnings dispersion with no rise in consumption dispersion when risk is absent as then initial conditions account for all the rise in earnings dispersion. What is important for determining the rise in consumption dispersion within the model is the magnitude of residual risk that remains after taxation and any insurance opportunities rather than the rise in earnings dispersion.

Finally, we note that our analysis of lifetime inequality is not very sensitive to the inclusion of a social insurance system. In the model with social insurance, initial conditions account for almost the same fraction of the variation in lifetime utility and lifetime earnings as compared to the benchmark model. Furthermore, in the model with social insurance the relative importance of variation in initial human capital compared to learning ability still holds. Specifically, a one standard deviation increase in human capital is, as of age 20, several times more valuable to an agent compared to a one standard deviation increase in learning ability.

7 Conclusion

This paper analyzes the proximate sources of lifetime inequality. We find that differences in initial conditions as of a real-life age of 20 account for more of the variation in realized lifetime utility, lifetime earnings and lifetime wealth than do shocks over the lifetime. Among initial conditions, a one standard deviation change in human capital is substantially more important as of age 20 than either a one standard deviation change in learning ability or initial wealth for how an agent fares in life. A one standard deviation increase in human capital is equivalent to between a 31 to 34 percent increase in consumption each period, whereas a one standard deviation increase in learning ability is equivalent to a 6 to 8 percent increase in consumption. A one standard deviation increase in initial wealth is the least important and is equivalent to a 4 to 5 percent increase in consumption.

Initial human capital and learning ability are positively correlated in the initial distribution which best matches the earnings distribution facts. This may suggest to some that the importance of learning ability differences relative to human capital differences would be greater if one were to evaluate lifetime inequality at a younger age. Some intuition for this position would be that learning ability is crystallized before age 20 and that learning ability differences are an important source of human capital differences as of age 20. We think

that such a line of reasoning is valuable to pursue. However, pushing back the age at which lifetime inequality is evaluated will raise the issue of the importance of one's family more directly than is pursued here. The importance of one's family and one's environment up to age 20 is not modeled in our work but is implicitly captured through their impact on initial conditions: human capital, learning ability and initial wealth.

Our analysis of lifetime inequality is based upon a parsimonious model. Thus, it is easy to think of initial differences or shocks that are not captured by the model. For example, shocks to mortality, health and preferences or shocks leading to the formation and dissolution of households are not captured by the model. It is not obvious to us that adding more sources of shocks will necessarily imply a more important role for shocks. The reason is that initial differences as of a young age may play a role in future health and preference states as well as a role in who forms households with whom.

In our view the risky human capital framework we have analyzed is likely to be important for the analysis of a number of economic policies and for many other issues. It has the potential to replace the standard incomplete-markets model with exogenous earnings or exogenous wages for both positive and normative analysis. For example, on the policy side, the framework is ideal to study tax policy. Analyzing the replacement of progressive taxation by flat-rate taxes in this setting would be of special interest. This follows from the role of progressive taxes in distorting human capital decisions as well as in reducing labor market risk, and the corresponding unexplored implications for human capital accumulation. All these reasons suggest that future work should investigate this framework in more detail.

References

- Baker, M. (1997), Growth-rate Heterogeneity and the Covariance Structure of Life Cycle Earnings, *Journal of Labor Economics*, 15, 338-75.
- Ben-Porath, Y. (1967), The Production of Human Capital and the Life Cycle of Earnings, *Journal of Political Economy*, 75, 352-65.
- Browning, M., Hansen, L. and J. Heckman (1999), Micro Data and General Equilibrium Models, in *Handbook of Macroeconomics*, ed. J.B. Taylor and M. Woodford, (Elsevier Science B.V, Amsterdam).
- Card, D. (1999), The Causal Effect of Education on Earnings, In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, Volume 3, (Elsevier, Amsterdam).
- Castañeda, A. Diaz-Jimenez, J. and V. Rios-Rull (2003), Accounting for Earnings and Wealth Inequality, *Journal of Political Economy*, 111, 818- 57.
- Congressional Budget Office (2004), Effective Federal Tax Rates 1979- 2001, available at <http://www.cbo.gov/>.
- Deaton, A. and C. Paxson (1994), Intertemporal Choice and Inequality, *Journal of Political Economy*, 102, 437-67.
- Eaton, J. and H. Rosen (1980), Taxation, Human Capital and Uncertainty, *American Economic Review*, 70, 705- 15.
- Geweke, J. and M. Keane (2000), An Empirical Analysis of Earnings Dynamics Among Men in the PSID: 1968- 1989, *Journal of Econometrics*, 96, 293- 356.
- Guvenen, F. (2006), Learning your Earning: Are Labor Income Shocks Really Very Persistent?, Institute for Empirical Macroeconomics, Discussion Paper 145.
- Heathcote, J., Storesletten, K. and G. Violante (2005), Two Views on Inequality over the Life Cycle, *Journal of the European Economic Association: Papers and Proceedings*, 765- 75.
- Heathcote, J., Storesletten, K. and G. Violante (2006), The Cross-Sectional Implications of Rising Wage Inequality in the United States, manuscript.
- Heckman, J., Lochner, L. and C. Taber (1998), Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents, *Review of Economic Dynamics*, 1, 1-58.

- Huggett, M. (1996), Wealth Distribution in Life-Cycle Economies, *Journal of Monetary Economics*, 38, 469- 94.
- Huggett, M., Ventura, G. and A. Yaron (2006), Human Capital and Earnings Distribution Dynamics, *Journal of Monetary Economics*, 53, 265- 90.
- Huggett, M. and J.C. Parra (2006), How Well Does the US Social Insurance System Provide Social Insurance?, manuscript.
- Keane, M. and K. Wolpin (1997), The Career Decisions of Young Men, *Journal of Political Economy*, 105(3), 473-522.
- Krebs, T. (2004), Human Capital Risk and Economic Growth, *Quarterly Journal of Economics*, 118, 709- 744.
- Krueger, D. and F. Perri (2006), Does Income Inequality Lead to Consumption Inequality? Evidence and Theory, *Review of Economic Studies*, 73, 163- 93.
- Levhari, D. and Y. Weiss (1974), The Effect of Risk on the Investment in Human Capital, *American Economic Review*, 64, 950-63.
- Lillard, L. and Y. Weiss (1979), Components of Variation in Panel Earnings Data: American Scientists 1960-70, *Econometrica*, 47, 437- 454.
- Mincer, J. (1974), *Schooling, Experience and Earnings*, Columbia University Press, New York.
- Mincer, J. (1997), The Production of Human Capital and the Life Cycle of Earnings: Variations on a Theme, *Journal of Labor Economics*, 15 (1), Part 2.
- Press, W. et. al. (1992), *Numerical recipes in FORTRAN*, Second Edition, (Cambridge University Press, Cambridge).
- Primiceri, G. and T. van Rens (2006), Predictable Life-Cycle Shocks, Income Risk and Consumption Inequality, manuscript.
- Siegel, J. (2002), *Stocks for the Long Run*, Third Edition, (McGraw-Hill, New York).
- Slesnick, D. and A. Ulker (2005), Inequality and the Life Cycle: Age, Cohort Effects and Consumption, manuscript.
- Storesletten, K., Telmer, C. and A. Yaron (2004), Consumption and Risk Sharing Over the Life Cycle, *Journal of Monetary Economics*, 51, 609- 33.

Tauchen, G. (1986), Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions, *Economics Letters*, 20, 177-81.

Weiss, Y. and L. Lillard (1978), Experience, Vintage and Time Effects in the Growth of Earnings: American Scientists 1960-1970, *Journal of Political Economy*, 86, 427- 47.

A Appendix

We analyze a dynamic programming formulation of an agent's decision problem. The dynamic programming problem is given below, where the state is $x = (h, k, a)$. The model implies that the period borrowing limits should depend upon age, human capital, learning ability and the distribution of shocks. We impose ability-specific limits $\underline{k}(a)$ and relax these limits until they are not binding. We also directly penalize choices leading to negative consumption later in life. This is a device for effectively imposing the endogenous limits implied by the model.

$$V_j(x) = \max_{(c, k', L)} u(c) + \beta E[V_{j+1}(h', k', a)]$$

subject to

$$c + k' \leq R_j h L + k(1 + r), \quad h' = z' F(h, l, a), \quad l + L = 1, \quad k' \geq \underline{k}(a)$$

We compute solutions to this problem by backwards recursion. We use a rectangular grid on the state variables (h, k) which is learning-ability specific. For each gridpoint and age j , we numerically solve the maximization problem on the right-hand-side of the Bellman's equation. Evaluating the objective involves a bi-linear interpolation of V_{j+1} across gridpoints.

To compute expectations, we follow Tauchen (1986) and discretize the shock into 5 equally-spaced values on the log scale. Values range from minus 2 to plus 2 standard deviations from the mean log-shock. Proceeding in this way gives a computed value function $V_j(x)$ and decision rules $(c_j(x), k_j(x), L_j(x))$ at gridpoints.

Given decision rules at each age, we simulate lifetime histories from the parametric distribution $G(h_1, k_1, a)$ of initial conditions described in section 4-6. To simulate histories, we put a grid on (h_1, k_1, a) . We draw a gridpoint (h_1, k_1, a) with a probability proportional to the density of the distribution at (h_1, k_1, a) . For any draw of an initial condition, we also draw a lifetime history of shocks from the relevant distribution. We calculate realizations of all endogenous variables using the computed decision rules, initial conditions and shock histories. Earnings statistics are computed from 40,000 draws of initial conditions and lifetime histories.²⁶

We determine the parameters of the distribution G by minimizing the squared distance of log model moments from log data moments. The objective of the minimization problem is

²⁶The decomposition of lifetime inequality is computed from 120,000 draws of initial conditions and lifetime histories.

$$\sum_{j=1}^{J_R-1} [(\log(m_{1j}/d_{1j}))^2 + (\log(m_{2j}/d_{2j}))^2 + (\log(m_{3j}/d_{3j}))^2],$$

where (m_{1j}, m_{2j}, m_{3j}) denote mean earnings, earnings Gini and the mean to median earnings ratio at age j in the model and where (d_{1j}, d_{2j}, d_{3j}) are the corresponding statistics from U.S. data. The simplex minimization routine AMOEBA, from Press et. al. (1992), is used to solve this minimization problem.

Table 1: Estimation of Human Capital Shocks

Min-Age	Max-Age	Period	N	σ	S.E. (σ)	σ_ϵ	S.E. (σ_ϵ)	\bar{s}
55	65	1969-2004	125	0.108	(0.029)	0.153	(0.013)	2
50	60	1969-2004	223	0.110	(0.023)	0.157	(0.011)	2
23	60	1969-2004	1521	0.158	(0.010)	0.177	(0.006)	2
55	65	1969-2004	106	0.103	(0.023)	0.149	(0.012)	3
50	60	1969-2004	200	0.104	(0.019)	0.151	(0.010)	3
23	60	1969-2004	1406	0.140	(0.009)	0.178	(0.005)	3
55	65	1969-1981	119	0.088	(0.040)	0.150	(0.015)	2
50	60	1969-1981	225	0.105	(0.027)	0.146	(0.013)	2
23	60	1969-1981	1322	0.152	(0.013)	0.166	(0.008)	2

Note: The entries provide the estimates for σ and σ_ϵ for various samples. The first and second column provide the minimum and maximum respective age in the sample. The third column refers to which PSID years are included. The column labeled N refers to the median number of observation across panel years. Columns labeled $S.E.$ refer to standard errors. The column denoted \bar{s} refers to the maximum s value used in computing log wage differences. In estimation all variance and covariance restrictions are always imposed.

Table 2: Benchmark Parameter Values

Definition	Symbol	Value
Model Periods	J	$J = 56$
Retirement Period	J_R	$J_R = 42$
Interest Rate	r	$r = 0.042$
Discount Factor	β	$\beta = 1.0/(1 + r)$
Period Utility Function	$u(c)$	$u(c) = \frac{c^{(1-\rho)}}{(1-\rho)}$ $\rho = 2$
Rental Rate	R_j	$R_j = (1 + g)^{j-1}$ $g = 0.0019$
Human Capital Shocks	z	$\log(z) \sim N(\mu, \sigma^2)$ $\sigma = 0.088, 0.108$ $\mu = -0.029, -0.031$
Law of Motion for Human Capital	$F(h, l, a)$	$F(h, l, a) = h + a(hl)^\alpha$ $\alpha = 0.675$ when $\sigma = 0.088$ $\alpha = 0.625$ when $\sigma = 0.108$
Distribution of Initial Conditions	G	$x \equiv (h_1, k_1, a) \sim G$ discussed in text

Table 3: Properties of Initial Distributions: Benchmark Model

Statistic	Low Shock ($\sigma = 0.088$)		High Shock ($\sigma = 0.108$)	
	$\alpha = 0.625$	$\alpha = 0.675$	$\alpha = 0.625$	$\alpha = 0.675$
Mean Learning Ability (a)	0.465	0.373	0.485	0.385
Coefficient of Variation (a)	0.223	0.213	0.205	0.204
Mean Initial Human Capital (h_1)	114.4	116.2	113.3	116.0
Coefficient of Variation (h_1)	0.439	0.432	0.391	0.387
Correlation (a, h_1)	0.802	0.720	0.783	0.789

Note: Entries show the moments of the distribution of initial conditions that best match the profiles of mean earnings, earnings dispersion and skewness, at the specified values of the shock σ and the elasticity parameter α . The elasticity value $\alpha = 0.675$ is the best estimate when $\sigma = 0.088$, whereas $\alpha = 0.625$ is the best estimate when $\sigma = 0.108$.

Table 4: Sources of Lifetime Inequality: Benchmark Model

Statistic	Low Shock $\sigma = 0.088$	High Shock $\sigma = 0.108$
Fraction of Variance in Lifetime Utility		
Due to Initial Conditions	.731	.621
Fraction of Variance in Lifetime Earnings		
Due to Initial Conditions	.708	.596

Note: Entries show the fraction of the variance of lifetime utility and lifetime earnings accounted for by initial conditions (initial human capital and learning ability).

Table 5: Sources of Lifetime Inequality: Model with Initial Wealth Differences

Statistic	Low Shock $\sigma = 0.088$	High Shock $\sigma = 0.108$
Fraction of Variance in Lifetime Utility		
Due to Initial Conditions	.770	.663
Fraction of Variance in Lifetime Earnings		
Due to Initial Conditions	.712	.590
Fraction of Variance in Lifetime Wealth		
Due to Initial Conditions	.722	.612

Note: Entries show the fraction of the variance of lifetime utility, lifetime earnings and lifetime wealth accounted for by initial conditions (initial human capital, learning ability and initial wealth). Wealth differences are measured directly from PSID data as explained in the text.

Table 6: Importance of Changes in Initial Conditions: Model with Initial Wealth

Equivalent Variations			
Variable	Change in Variable	Low Shock $\sigma = 0.088$	High Shock $\sigma = 0.108$
Human Capital	+ 1 st. deviation	34.1	30.7
	- 1 st. deviation	-25.9	-25.4
Learning Ability	+ 1 st. deviation	7.9	6.4
	- 1 st. deviation	-5.0	-7.6
Initial Wealth	+ 1 st. deviation	4.7	3.7
	- 1 st. deviation	-2.1	-3.6

Expected Lifetime Wealth			
Variable	Change in Variable	Low Shock $\sigma = 0.088$	High Shock $\sigma = 0.108$
Human Capital	+ 1 st. deviation	34.3	29.7
	- 1 st. deviation	-23.9	-22.4
Learning Ability	+ 1 st. deviation	9.9	9.0
	- 1 st. deviation	-5.9	-8.4
Initial Wealth	+ 1 st. deviation	4.1	2.5
	- 1 st. deviation	-1.4	-2.8

Note: The top panel states equivalent variations, whereas the bottom panel states the percentage change in the expected lifetime wealth associated with changes in each initial condition. The baseline initial condition is set equal to the mean log values of initial human capital, learning ability and wealth. Changes in initial conditions are also in log units.






Figure 1. Mean Earnings by Age

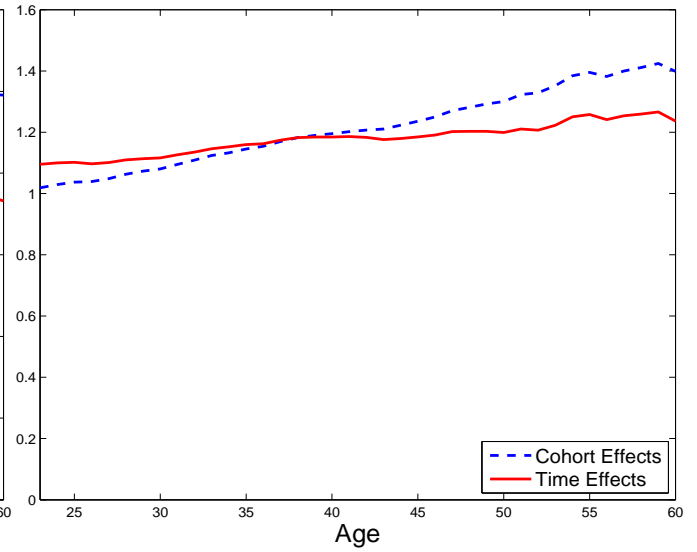
Note: Figure 1 plots the age effects in mean earnings, after controlling for either time or cohort effects based on data from PSID, 1969-2004.



(a) Variance of Log Earnings




(b) Earnings Gini




(c) Earnings Skewness (Mean/Median)

Figure 2. Dispersion and Skewness of Earnings by Age


Note: Figure 2 plots the age effects in earnings dispersion and in a measure of skewness after controlling for either time or cohort effects based on data from the PSID, 1969-2004.



(a) Mean Earnings



(b) Earnings Gini



(c) Earnings Skewness (Mean/Median)

Figure 3. Earnings in Model and Data

Note: Figure 3 displays the model implications for the age profiles of mean earnings, the earnings Gini coefficient and a measure of earnings skewness. The model implications are displayed for low shock case ($\sigma = 0.088$) as well as for the high shock case ($\sigma = 0.108$). The figure also reports the U.S. data counterpart based on the time effects case.






Figure 4. Mean Time in Human Capital Accumulation

Note: Figure 4 plots the age profile for the fraction of time spent in human capital accumulation, for both specification of shocks.




(a) Earnings Gini: Eliminating Learning Ability Differences



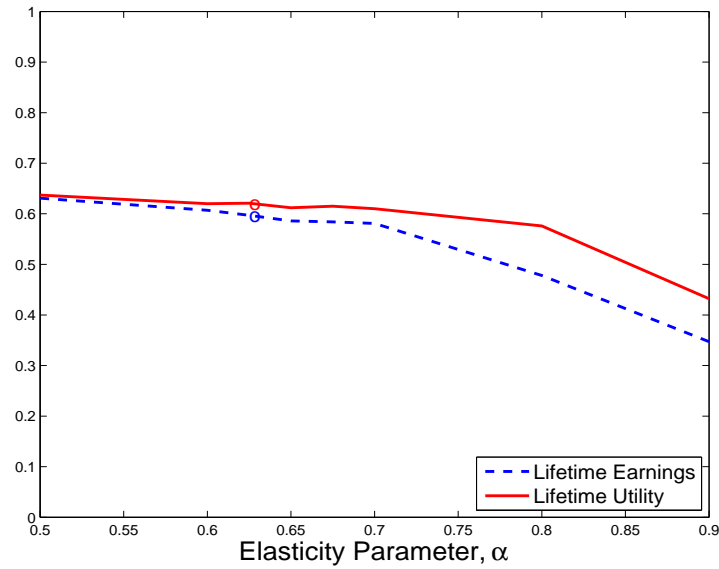
(b) Earnings Gini: Eliminating Shocks

Figure 5. Earnings Dispersion: Learning Ability versus Shocks

Note: Figure 5 is based on the benchmark model with $\sigma = 0.108$. In Panel (a), all differences in learning ability are eliminated. In Panel (b) shocks are eliminated, but differences in learning ability and initial human capital are the same as in the benchmark case.



(a) Lifetime Inequality: $\sigma = 0.088$



(b) Lifetime Inequality: $\sigma = 0.108$

Figure 6. Lifetime Inequality and the Elasticity Parameter α

Note: Figure 6 displays the fraction of the variance in lifetime earnings and utility due to initial conditions as the elasticity parameter α varies. In each panel, the values corresponding to the best estimate of α are highlighted.




Figure 7. Rise in Consumption Dispersion: Model and Data.

Note: Figure 7 plots the rise in the variance of log-consumption for ages 25-60 both in U.S. data and the model. The data is based on Deaton and Paxson (1994) and Primiceri and van Rens (2006). The model results are for the version of the model with a social insurance system (SI)– see text for details.