Sources of Lifetime Inequality

Mark Huggett, Gustavo Ventura and Amir Yaron*

February 9, 2007

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 with risky human capital. 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 human capital raises an agent's lifetime expected earnings profile, whereas an increase in learning ability rotates this profile counter-clockwise.

JEL Classification: D3, D91, E21
KEYWORDS: Lifetime Inequality, Human Capital, Idiosyncratic Risk.

*Affiliation: Georgetown University, University of Iowa, and The Wharton School–University of Pennsylvania and NBER respectively. We thank Richard Rogerson, Victor Rios Rull, Tom Sargent for useful discussions, and 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, University of Oslo, Penn, UCLA and The Wharton School. 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.
Corresponding author: Amir Yaron.
Address: The Wharton School; University of Pennsylvania, Philadelphia PA 19104- 6367.
E-mail: `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, 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 against those directed at shocks over the lifetime. 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 disparate fields.

We start from the premise that a useful framework for analyzing lifetime inequality needs to account for some key distributional facts. To this end, we document how mean earnings and how measures of earnings dispersion evolve for U.S. males. We find that mean earnings are hump shaped and that earnings dispersion increases over most of the working lifetime.[1]

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.

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.

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

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 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 life cycle.

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 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 and in realized lifetime wealth. Lifetime wealth is the realized present value of earnings plus the value of initial financial wealth. We find that between 62 to 77 percent of the variation in lifetime utility and between 54 to 72 percent of the variation in lifetime wealth is due to variation in initial conditions. The higher estimate of the role for initial conditions 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.

[2]This 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.

[3]Since 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.

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. More specifically, we find that a one standard deviation increase in initial wealth increases expected lifetime wealth by 3-5 percent. In contrast, a one standard deviation increase in learning ability or initial human capital increases expected lifetime wealth by 11-15 percent and 32-36 percent, respectively. An increase in initial human capital raises an agent's lifetime expected earnings profile, whereas an increase in learning ability rotates this profile counter-clockwise. It is important to know whether these changes in expected lifetime wealth reflect how the agent's in the model value these changes in initial conditions. Broadly, we find that this holds. Specifically, we ask what is the permanent percentage change in consumption which is equivalent for an agent to these changes in initial conditions. We find that the equivalent changes in consumption are roughly in line with the impact on 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). Their model is an incomplete-markets model in which labor earnings over the lifetime is exogenous.[4] They 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 lifetime utility is due to differences in initial conditions.[5]

We note three difficulties related to this alternative incomplete-markets view. First, the importance of idiosyncratic earnings risk may be overestimated. 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 differences in permanent earnings growth rates. Second, although the

[4] Similar models have been used in the literature on economic inequality. This literature is surveyed by Quadrini and Rios-Rull (1997). Recent papers in the literature include Castañeda, Diaz-Jimenez and Rios-Rull (2003), Heathcote, Storesletten and Violante (2005b), Huggett (1996), Krueger and Perri (2005), among many others.

[5] In 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.

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 a number of recent studies. A key reason why our model produces less of a rise in consumption dispersion than the exogenous-earnings model is that part of the rise in earnings dispersion is due to initial conditions. Thus, part of the rise is anticipated by agents and therefore reflected in consumption early in life. Third, the 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 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. (2005b)) 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 earnings in 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.[6] 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 $Z^j \equiv \{z^j = (z_1, ..., z_j) : z_i \in Z, i = 1, ..., j\}$, where Z is a finite set of possible shock realizations. $P(z^j)$ denotes the probability of history z^j.

[6]The 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.

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

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.

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

subject to

(1) $c_j + k_{j+1} = k_j(1+r) + e_j, \forall j$ and $k_{J+1} \geq 0$
(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.[7] 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

[7] More 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.

from all other model parameters. This holds, as an approximation, towards the end of the working life. 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 a measure of earnings dispersion evolve with age for a cohort. Second, we estimate a process for human capital shocks by using wage rate data.

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.[8] 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

[8] Real 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.

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.

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

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

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.[9] 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.[10]

[9] See Weiss and Lillard (1978) and Deaton and Paxson (1994) among others.

[10] A 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.

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.

A similar analysis can be carried out in order to extract the age profile of measures of earnings dispersion. We consider two standard measures of dispersion: the variance of log earnings and the Gini coefficient of earnings. Let $v_{j,t}$ and $g_{j,t}$ respectively be the cross-sectional variance of log earnings and the Gini coefficient of those agents who are of age j in year t. Then the age profiles are derived from the following regressions.[11]

$$v_{j,t} = \alpha_s^v + \beta_j^v + \gamma_t^v + \epsilon_{j,t}^v$$

$$g_{j,t} = \alpha_s^g + \beta_j^g + \gamma_t^g + \epsilon_{j,t}^g$$

Figure 2 and 3 provide the age effects based on cohort and time effects for each measure of earnings dispersion. Again, the cohort effects are derived by setting $\gamma_t = 0, \ \forall t$, while the time effects are constructed by setting $\alpha_s = 0, \ \forall s$ for both the log variance and Gini regressions. In Figure 2 we see that the cohort effect view implies a rise in the variance of log earnings of about 0.4 from age 23 to 60 while the time effects imply a smaller rise of only about 0.2. The same qualitative pattern can be seen for the Gini coefficient in Figure 3.

We will ask the model to best match the time effects view of the evolution of the earnings distribution. Heathcote et al (2005a) present an argument for this view. Their argument is based partly on the fact that within-cohort and within-age group changes in earnings and wage dispersion vary over time but are of similar magnitude over much of the period they examine.

[11] We use 5-year age bins centered at age j to compute $v_{j,t}$ and $g_{j,t}$. Figure 2 and 3 are normalized so that age profiles run through the mean values of $v_{j,t}$ and $g_{j,t}$ across years at age $j = 38$.

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.[12]

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

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

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

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 $(\hat{z}_t, \epsilon_t)$ 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.

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

[12] Heckman et al (1998) use a similar line of reasoning to estimate differences in rental rates across skill groups within a model which abstracts from idiosyncratic risk.

$$E[y_{t,s}] = \hat{R}_{t+s} - \hat{R}_t + s\mu$$
$$Var(y_{t,s}) = s\sigma^2 + 2\sigma_\epsilon^2$$
$$Cov(y_{t,s}, y_{t,r}) = r\sigma^2 + \sigma_\epsilon^2 \text{ for } r < s$$

To make use of these moment restrictions, one needs to be able to measure the variable $y_{t,s}$ and to have agents 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.[13] In estimation we use 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 ($\bar{s} = 3$).[14] 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

[13] The PSID data is not available for the years 1997, 1999, 2001, and 2003. In the years preceding those years we impose that the agent is available for three consecutive years and use a two year growth rate.

[14] We 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.

of the persistent wage shocks estimated by Heathcote et al (2005b) 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 well-known fact that the increase in cross-sectional earnings inequality in the U.S. occurred after 1981. Over the early part of the sample period we find that the estimated shocks for the 50-60 and the 55-65 age groups are smaller than those estimated 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, interest rates and the human capital production function based upon previous studies. Third, we set the parameters governing the distribution of initial conditions so that the model best matches the age profiles of mean earnings and earnings dispersion estimated in the previous section. In choosing this initial distribution, 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). The law of motion for human capital embodies the Ben-Porath (1967) human capital production function. The human capital literature has estimated the elasticity parameter α governing the production of human capital to lie in the range 0.5 to just over 0.9. These estimates are surveyed by Browning et. al. (1999, Table 2.3- 2.4). We set α to lie in the middle of this range.

We set the value g so that the rental rate of human capital and average male earnings in US data grow at the same rate. We calculate that in the PSID over the period 1968-yy 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, 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 $55 - 65$ 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 restrict the distribution G of initial human capital, assets and learning ability to lie in a parametric class. In the benchmark model, we set initial assets to zero and specify that initial human capital and learning ability are jointly lognormally distributed so that $\log(x) \sim N(\mu_x, \Sigma)$, when $x = (h_1, a)$. We then choose (μ_x, Σ) to best match the dynamics of the US earnings distribution documented in section 3, given all other model parameters. The next section highlights properties of the resulting initial distributions. Later in the paper we explore a tri-variate distribution of initial conditions, where the initial asset distribution matches features of net wealth holdings for young households in the PSID.

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.[15] The analysis focuses on the benchmark model without initial wealth differences.

5.1 Dynamics of the Earnings Distribution

The age profiles of mean earnings and earnings dispersion produced by the benchmark model are displayed in Figure 4. The model generates the hump-shaped earnings profile for a cohort by a standard human capital accumulation argument. Early in the life cycle, the bulk of individuals accumulate human capital in net terms and progressively devote increasing fractions of their time to market work. Both effects act to increase mean earnings as agents age. Towards the end of the working life-cycle, mean human capital accumulation for a cohort levels off and eventually falls. Human capital falls when the mean multiplicative shock to human 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 fall late in life when growth in the rental rate of human capital is not enough to offset the mean fall in human capital.

Two forces account for the rise in earnings dispersion. 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 for a common level of current human capital, agents with high learning ability choose to produce more human capital 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.

We now try to understand the quantitative importance of these two forces by alternatively eliminating ability differences or eliminating shocks, other things equal. To

[15]Methods used to compute solutions to the model are described in the Appendix.

highlight the role of learning ability differences, we change 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 shows that eliminating ability differences leads to the striking result that the rise in earnings dispersion over the lifetime is almost completely eliminated.

The pattern of dispersion that remains after removing ability differences 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. 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 constant elasticity human capital production functions. An implication is that 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. We note that earnings dispersion in Figure 5 at the end of the working lifetime increases. This occurs because human capital production at the end of life goes to zero. Thus, the force leading to convergence is eliminated.

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 a U-shaped earnings dispersion profile. Figure 6 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.[16] In terms of dispersion in labor earnings, human

[16]This 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.

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. Consequently, earnings dispersion is higher at the start of the working life-cycle. Earnings dispersion falls for the first 10 years of the working lifetime. At this point the earnings of higher learning ability agents are overtaking the earnings of their lower ability counterparts.

5.2 Properties of the Initial Distribution

Table 3 summarizes a number of properties of the distribution of initial conditions which produce the results highlighted in the previous section. When the standard deviation of shocks 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 from the model for a given initial distribution leads to a counter-clockwise shift in 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.

Following this intuition, to produce the earnings facts as risk increases 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. A higher mean ability level has this effect as it leads to an increase in the time put into human capital production early in life. 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, lifetime earnings and lifetime wealth. Lifetime wealth equals the realized present value of earnings (i.e. lifetime earnings) plus initial wealth.[17] Such a decomposition makes use of the fact that a random variable can be written as the sum of its condition mean plus the variation from its conditional mean. As these two components are orthogonal, the total variance in the random variable equals the sum of the variance in the conditional mean plus the variance around the conditional mean.

Table 4 decomposes lifetime inequality into the sources highlighted by the model. Each measure of lifetime inequality is analyzed as of the start of the working life cycle, which is taken to be a real-life age of 20. Based on the estimates of the shocks from Table 1, we focus on two cases for the magnitude of idiosyncratic shocks: a high shock case $\sigma = 0.108$ as well as a low shock case $\sigma = 0.088$. Table 4 shows that for each case the majority of the variation in lifetime utility and lifetime earnings in the benchmark model is due to variation in initial conditions. Specifically, $63 - 73$ percent of the variation in lifetime utility and $60 - 71$ percent of the variation in lifetime earnings is due to initial conditions.

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.[18] We express net

[17] Lifetime utility and lifetime wealth along a given 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 $W(z^J; h_1, k_1, a) = k_1(1 + r) + \sum_{j=1}^{J} e_j(z^J; h_1, k_1, a)/(1 + r)^{j-1}$.

[18] The 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.

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 $(0.377, 0.778, 1.340)$.[19] 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. 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 presents results for sources of lifetime inequality when we allow for initial wealth differences. We find that initial conditions account for $62 - 77$ percent of the variation in lifetime utility and $54-72$ percent of the variation in lifetime wealth. Thus, the majority of the variation in all measures of lifetime inequality is due to initial conditions.[20] Table 5 also finds that initial conditions play a greater role in lifetime inequality when the magnitude of shocks is set to the lowest point estimate (i.e. $\sigma = 0.088$). The intuition is that the greater the magnitude of human capital shocks the smaller the role for initial conditions in accounting for patterns of earnings dispersion over the lifetime.

[19] In 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 $(0.313, 0.776, 1.432)$.

[20] We have also investigated lifetime inequality when a social insurance system is added. In the model with social insurance and no initial wealth difference, initial conditions account for a similar fraction of the variation in lifetime utility as compared to the models analyzed in Tables 4 and 5. The social insurance system is described in detail in the Appendix.

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.

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 earnings or expected lifetime wealth change in response to a one standard deviation change in an initial condition. The baseline initial condition is set so that $(\log h, \log a)$ equal the mean log values of initial human capital and learning ability and that the shifted log initial wealth is set to its mean. The changes in initial conditions are also in standard deviations of log variables.

Table 6 presents the results of this analysis. 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.[21] A one standard deviation increase in initial human capital is equivalent to a $32 - 36$ percent increase in expected lifetime wealth. In contrast, a one standard deviation increase in learning ability or initial wealth increases expected lifetime wealth by $11 - 15$ percent and $3 - 5$ percent, respectively. Thus, in comparing the impact on expected lifetime wealth for one standard deviation increases in each initial condition, 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 wealth has the smallest impact on lifetime wealth.

Table 6 also analyzes equivalent variation measures of the importance of different initial conditions. We find that how an agent evaluates changes in initial conditions

[21] A one standard deviation movement in $\log h_1$ and $\log a$ is xx and yy, respectively. Thus, in percentage terms, variation in initial human capital is larger than learning ability variation.

in terms of equivalent percentage changes in consumption is roughly in line with the impact on expected wealth. Thus, we find that one standard deviation changes in initial human capital have the greatest impact, the corresponding change in learning ability have the next biggest impact and initial wealth changes have the least important impact on equivalent variations. We note, however, that the impact of one standard deviation changes in learning ability on these two measures displays the weakest link. Specifically, the changes in learning ability analyzed in Table 6 have a greater impact on expected wealth than on equivalent variations. Intuitively, this occurs because higher ability leads to higher mean earnings and a higher earnings variance later 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.

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 and realized 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 32 to 36 percent increase in consumption each period, whereas a one standard deviation increase in learning ability is equivalent to a 8 to 10 percent increase in consumption. A one standard deviation increase in initial wealth is the least important of these initial conditions. It is equivalent to between 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. Implicitly, the importance of one's family and one's environment early in life is captured in our work by their impact on 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 human capital policies and for many other issues. It has the potential to replace the standard life-cycle model with exogenous earnings or exogenous wages. For this reason, future work should investigate the framework in detail. We mention three directions to pursue. First, evidence bearing on the nature of human capital risk is key. The methodology we employ for estimating this risk could certainly be extended beyond the case of independent shocks. Second, the earnings implications of the framework could be investigated in more detail. Here it would be valuable to allow for valued leisure. Third, it is important to evaluate the consumption implications of the framework. The Appendix takes a step in this last direction. It shows that the rise in consumption dispersion within cohort produced by the model is substantially less than the rise found in U.S. data by Deaton and Paxson (1994) but similar to the rise found in more recent data by a number of researchers.

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.

Guvenen, F. (2006), Learning your Earning: Are Labor Income Shocks Really Very Persistent?, Institute for Empirical Macroeconomics, Discussion Paper 145.

Heathcote, J., Storresletten, K. and G. Violante (2005a), Two Views on Inequality over the Life Cycle, *Journal of the European Economic Association: Papers and Proceedings,* 765- 75.

Heathcote, J., Storresletten, K. and G. Violante (2005b), 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 (2005), 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.

Lucas, R. E. (2003), Macroeconomic Priorities, *American Economic Review,* 93, 1-14.

Mincer, J. (1974), *Schooling, Experience and Earnings,* Columbia University Press, New York.

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.

Quadrini, V. and V. Rios-Rull (1997), Understanding the U.S. Wealth Distribution, *Quarterly Review of the Federal Reserve Bank of Minneapolis,* 22-36.

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

A.1 Consumption

A number of studies analyze the variance of log household-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 household consumption in U.S. data from two such studies. Deaton and Paxson (1994) analyze U.S. Consumer Expenditure Survey (CEX) data from 1980 - 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 studies find that consumption dispersion rises with age for a cohort and that the rise is smaller than the rise in Deaton and Paxson (1994).

The sharp rise in consumption dispersion documented by Deaton and Paxson (1994) has been viewed as evidence for the presence of highly persistent, uninsurable earnings shocks (e.g. Lucas (2003)). Storesletten et al (2004) showed that the persistent earnings process that they estimate is consistent with the growth in earnings dispersion for a cohort and, within a life-cycle model with social insurance, produces an increase in consumption dispersion in line with the Deaton and Paxson estimates.

Figure 7 also highlights the consumption patterns implied by our model with social insurance.[22] The rise in consumption dispersion in the model is less than the rise in Deaton and Paxson (1994). The rise in the model is approximately 8 log points when shocks are set to our highest point estimate and is even smaller when shocks are set to the lowest point estimate. The rise found by Primiceri and van Rens (2006), using CEX data for 1980-2000, is between our results for the high and low shock case. We note that social insurance is quantitatively important for this result. Specifically, the models for the high and low shock case produce a rise of approximately 18 and 13 log points from age 25 to 60, respectively, without social insurance.

We emphasize that the risky human capital model decouples the rise in consumption dispersion from the rise in earnings dispersion. In particular, the model can produce the rise in U.S. earnings dispersion with no rise in consumption dispersion when risk is absent. In this

[22] The model social insurance system consists of a social security system and 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. We follow the steps in Huggett and Parra (2006) for how we implement this tax function within our model.

case, initial conditions fully account for 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.

A.2 Computation

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

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

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 a parametric distribution $G(h_1, k_1, a)$ of initial conditions. The distribution under consideration is described in section 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.[23]

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

[23]Tables 4 and 5 are computed from $120,000$ draws of initial conditions and lifetime histories.

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

where (m_{1j}, d_{1j}) denote mean earnings in model and data at different ages j and (m_{2j}, d_{2j}) denote earnings Gini coefficients in the model and 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	$\beta = 1.0/(1 + r)$
Preferences	$u(c)$	$u(c) = \frac{c^{(1-\rho)}}{(1-\rho)}$ $\rho = 2$
Law of Motion for Human Capital	$F(h, l, a)$	$F(h, l, a) = h + a(hl)^{\alpha}$ $\alpha = 0.7$
Rental Rate	R_j	$R_j = (1 + g)^{j-1}$ $g = 0.0019$
Human Capital Shocks	z	$\log(z) \sim N(\mu, \sigma^2)$ $\sigma = 0.088, 0.108$ $\mu = -0.029, -0.031$
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 $\sigma = 0.108$
Mean Learning Ability (a)	0.332	0.348
Coefficient of Variation (a)	0.215	0.206
Mean Initial Human Capital (h_1)	117.6	118.9
Coefficient of Variation (h_1)	0.426	0.403
Correlation (a, h_1)	0.784	0.753

Note: Entries show the moments of the distribution of initial conditions that best reproduces the mean earnings and earnings dispersion profiles.

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	.725	.629
Fraction of Variance in Lifetime Earnings Due to Initial Conditions	.708	.596

Note: Entries show the fraction of the variance accounted for by initial conditions (initial human capital and learning ability), under the distribution of initial conditions that best reproduces the mean earnings and earnings dispersion profiles.

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

Statistic	Low Shock $\sigma = 0.088$	High Shock $\sigma = 0.108$
Fraction of Variance in Lifetime Utility Due to Initial Conditions	.766	.616
Fraction of Variance in Lifetime Earnings Due to Initial Conditions	.708	.524
Fraction of Variance in Lifetime Wealth Due to Initial Conditions	.716	.537

Note: Entries show the fraction of the variance accounted for by initial conditions (initial human capital, initial wealth and learning ability). Wealth differences are measured directly from PSID data as explained in the text. The underlying distribution of initial conditions is the one that best reproduces the mean earnings and earnings dispersion profiles.

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	35.8	31.5
	− 1 st. deviation	-27.5	-24.3
Learning Ability	+ 1 st. deviation	7.8	10.1
	− 1 st. deviation	-5.7	-4.8
Initial Wealth	+ 1 st. deviation	3.8	5.4
	− 1 st. deviation	-3.2	-2.2

Expected Lifetime Wealth			
Variable	Change in Variable	Low Shock $\sigma = 0.088$	High Shock $\sigma = 0.108$
Human Capital	+ 1 st. deviation	35.7	32.2
	− 1 st. deviation	-25.1	-22.3
Learning Ability	+ 1 st. deviation	10.5	14.5
	− 1 st. deviation	-6.8	-6.1
Initial Wealth	+ 1 st. deviation	3.3	4.8
	− 1 st. deviation	-2.4	-1.6

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.

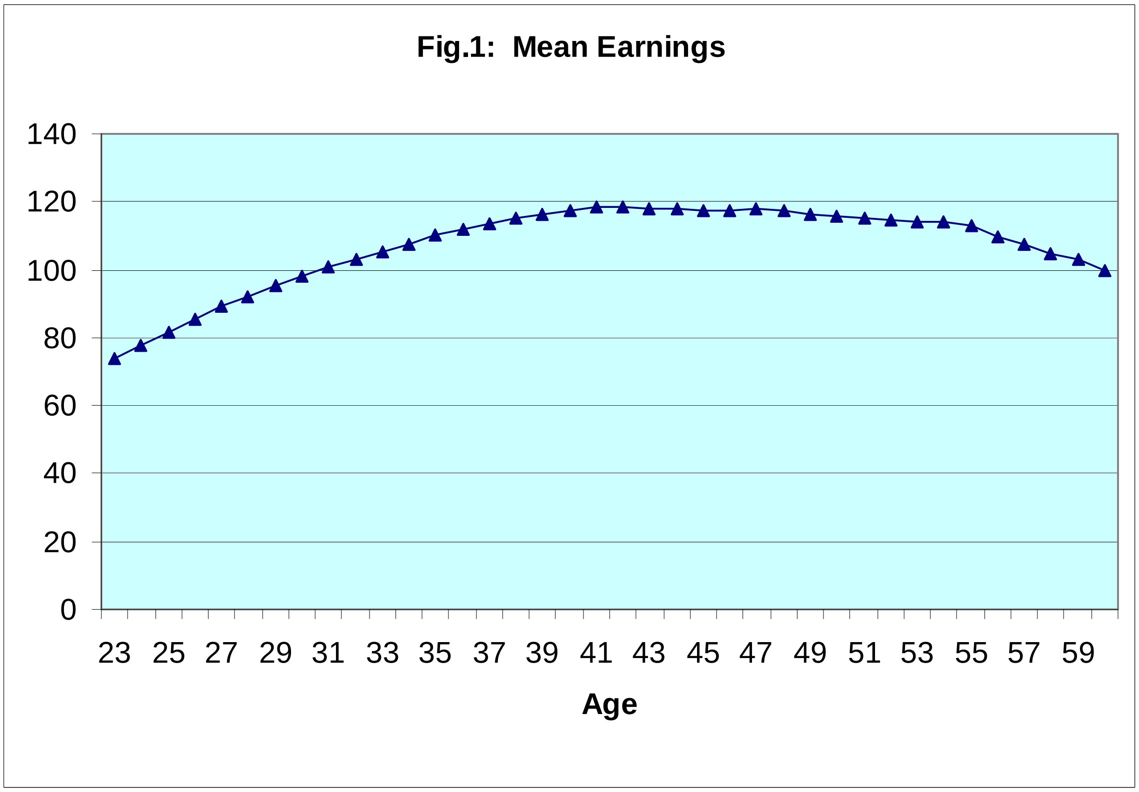
Fig.1: Mean Earnings
140
120
100
80
60
40
20
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age

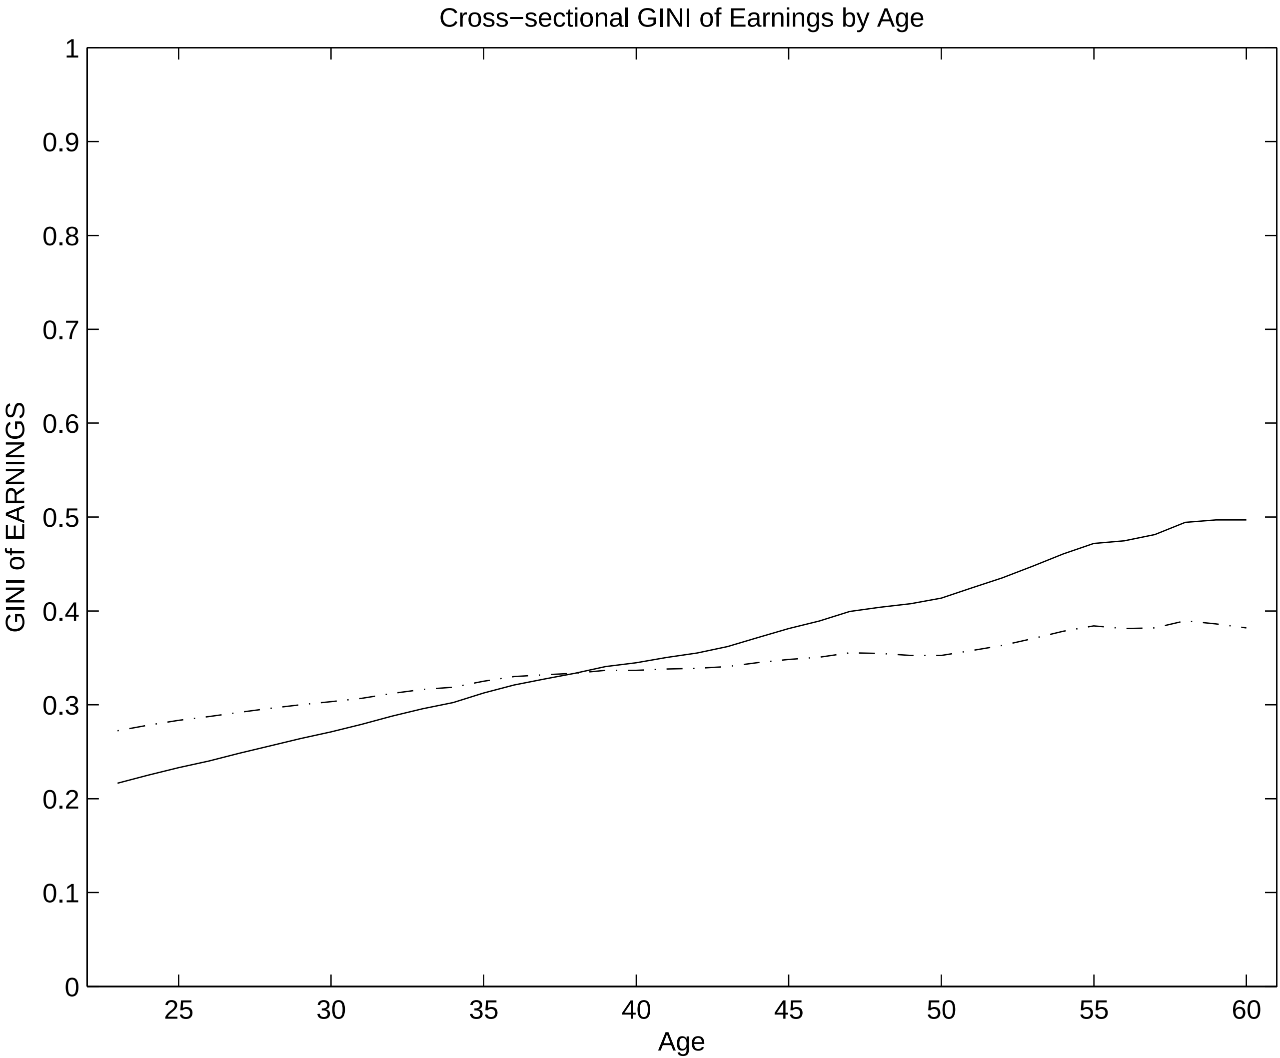
Cross−sectional GINI of Earnings by Age
GINI of EARNINGS
Age
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
25
30
35
40
45
50
55
60

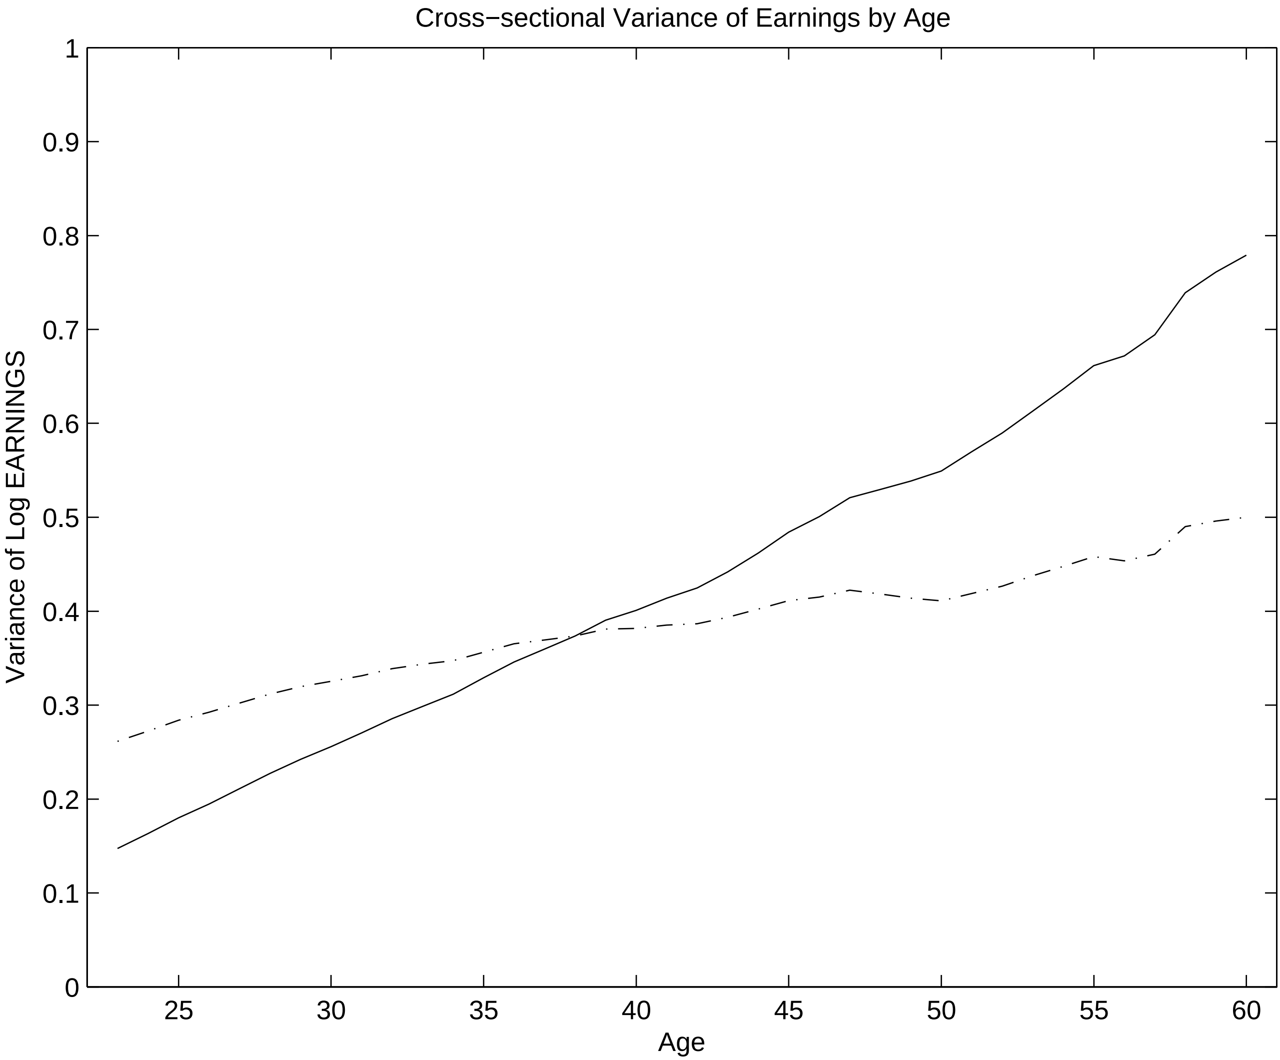
Cross−sectional Variance of Earnings by Age
Variance of Log EARNINGS
Age
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
25
30
35
40
45
50
55
60

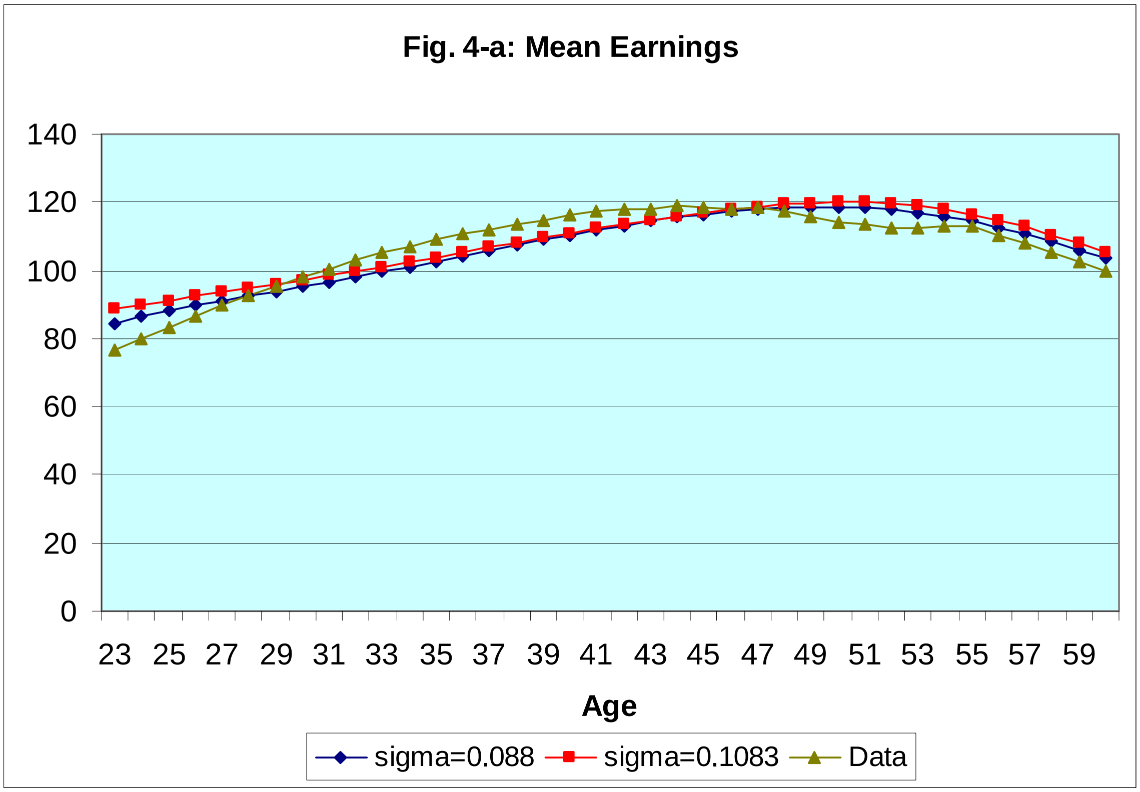

Fig. 4-a: Mean Earnings
140
120
100
80
60
40
20
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
sigma=0.088
sigma=0.1083
Data

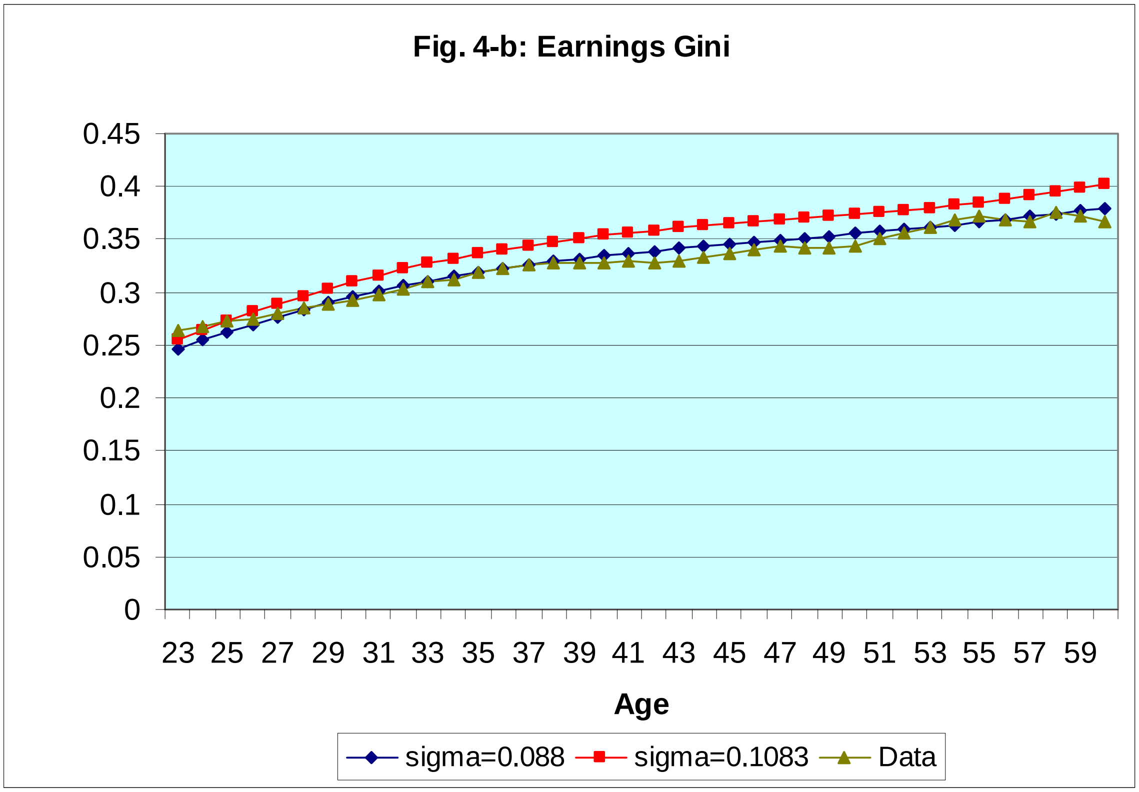
Fig. 4-b: Earnings Gini
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
sigma=0.088
sigma=0.1083
Data

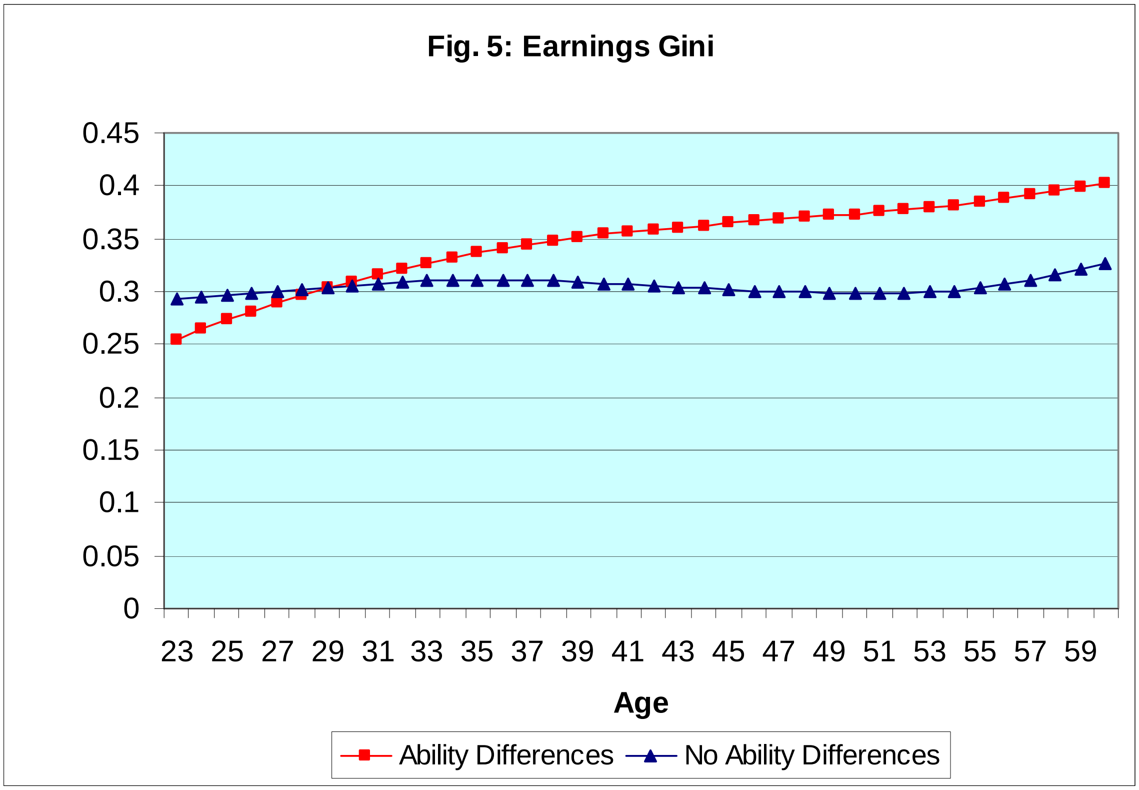
Fig. 5: Earnings Gini
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
Ability Differences
No Ability Differences

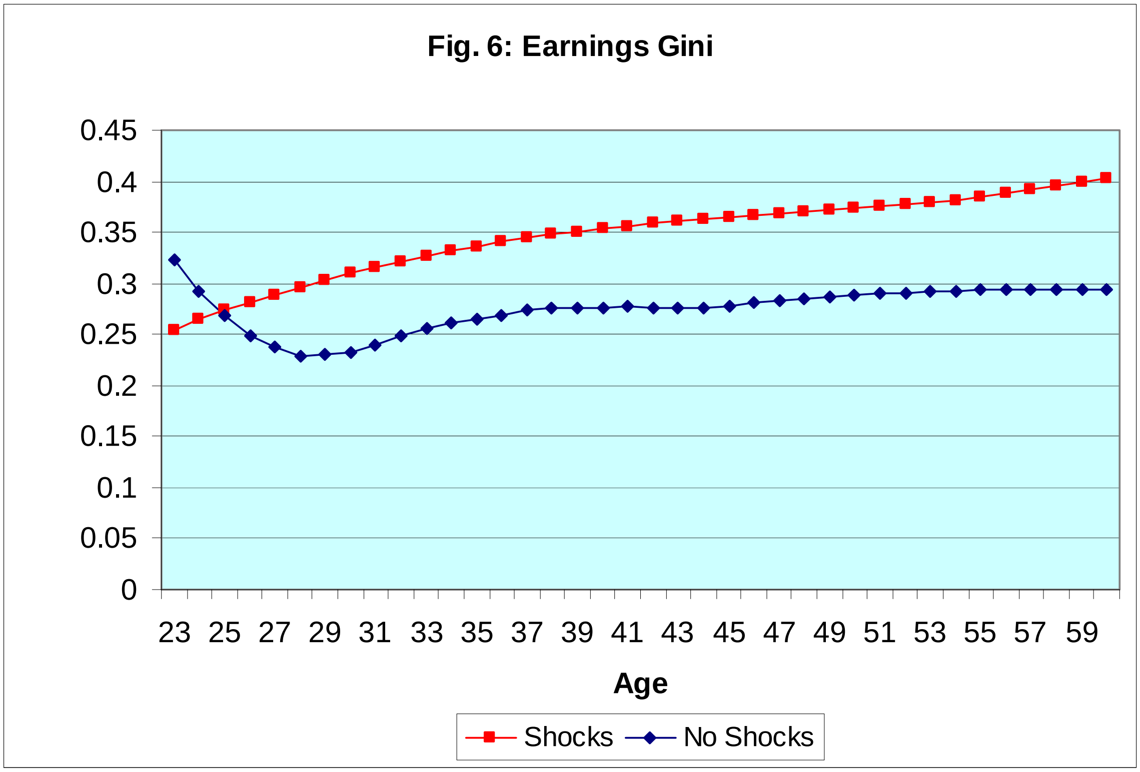
Fig. 6: Earnings Gini
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
Shocks
No Shocks

Fig. 7: Consumption Dispersion