Sources of Lifetime Inequality

Mark Huggett, Gustavo Ventura and Amir Yaron*

May 7, 2006

Abstract

Is lifetime inequality mainly due to differences across people established early in life or to differences in luck experienced over the lifetime? To answer this question, we develop a framework that integrates risky human capital investments into an otherwise standard life-cycle, permanent-income model. We find that differences in *initial conditions* account for more of the variation in lifetime utility and in lifetime wealth than do differences in shocks received over the life cycle. Among initial conditions, differences in *learning ability* account for substantially more of the variation in expected lifetime utility or expected lifetime wealth than do differences in initial human capital and initial financial wealth.

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

*Affiliation: Georgetown University, Pennsylvania State University, and The Wharton School of Business and NBER respectively. 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 the 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 importance of the myriad policies directed at modifying initial conditions or providing insurance for initial conditions against those that provide insurance for 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 fields as disparate as macroeconomics, labor economics and public economics. In these fields policies are often analyzed by abstracting from an important role for luck over the lifetime.[1]

This paper quantifies the importance of different sources of lifetime inequality. The analysis focuses on the inequality in lifetime utility and lifetime wealth. The former measure is theoretically well motivated since in standard economic theory agents maximize expected lifetime utility. The focus on lifetime wealth is important since in standard life-cycle theory an agent's lifetime budget constraint is determined by lifetime wealth.

We start from the premise that a useful framework for analyzing lifetime inequality needs to be consistent with some critical distributional facts over the life cycle. To this end, we document for US males how mean earnings and how a measure of earnings dispersion evolve for a typical cohort as the cohort ages. We find that mean earnings are hump shaped over the working lifetime and that earnings dispersion increases over most of the working lifetime. These features of US data have been widely documented in other studies and have been viewed as key facts that a theory of inequality should explain.[2] In

[1]In labor economics, much of the analysis of human capital policies abstracts from such luck (see the review by Carneiro and Heckman (2003)). In public economics, the Mirrlees (1971) model is central. It allows for initial differences but abstracts from ex-post luck.

[2]Mincer (1974) documented related patterns in US cross-section data. Many studies use repeated cross-section data at the individual or the household level to document these patterns for a cohort as the cohort ages (e.g. Deaton and Paxson (1994), Storesletten, Telmer and Yaron (2004), Heathcote,

particular, the rise in earnings dispersion has been viewed as reflecting an important role for idiosyncratic shocks as well as a role for initial differences in learning ability.

The model that we explore integrates risky human capital into an otherwise standard incomplete-markets model. Within the model, individuals differ in terms of three *initial conditions:* initial human capital, learning ability and financial wealth. As individuals 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 uninsurable, idiosyncratic shocks each period.

Our model produces the hump-shaped mean earnings profile by a standard human capital argument. 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 as human capital depreciates and little or no time is put into producing new human capital.

Two forces account for the increase in earnings dispersion. One force is that agents differ in learning ability. Agents with higher learning ability have steeper sloped mean earnings profiles than low ability agents, other things equal.[3] The other force is that agents differ in idiosyncratic, human capital shocks received over the life cycle. To identify the contribution of each force, 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 entirely determined by shocks, independently of all other model parameters. Thus, we estimate the shock process using these moments. Second, given an estimate of the shock process, we choose the initial distribution of financial wealth, human capital and learning ability across agents to best match the earnings facts described above, given all other model parameters.[4] We then use these estimates of initial conditions and

Storesletten and Violante (2005a) and Huggett, Ventura and Yaron (2006)).

[3]This mechanism is supported by the literature, reviewed by Card (1999), on the shape of the mean age-earnings profiles by years of education. It is also supported by the work of Baker (1997) and Guevenen (2005). They estimate a statistical model of earnings and find that permanent differences in individual earnings growth rates are negatively correlated with initial earnings levels and produce large differences in the present value of earnings. Our model endogenously produces these relationships.

[4]Since a measure of financial wealth is observable, we choose the tri-variate initial distribution to match the unconditional distribution of financial wealth among young males.

shocks in our decomposition of lifetime inequality into its sources.

Our preliminary findings indicate that the majority of the variation in lifetime utility and in lifetime wealth is due to differences in initial conditions. Among the initial conditions, we find that differences in learning ability are the most important source of variation in expected lifetime utility or in expected lifetime wealth. Differences in initial human capital serve to raise or lower the lifetime expected earnings profile whereas differences in learning ability serve to change the shape of this profile. Higher learning ability lowers initial earnings but raises future expected earnings, other things equal. We also find that differences in learning ability are key in that much of the rise in earnings dispersion over the lifetime within the model is due to ability differences, given our estimate of the magnitude of human capital risk. Specifically, eliminating learning ability differences leads, other things equal, to an almost flat profile of earnings dispersion with age.

Related Literature 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 is exogenous.[5] They conclude that within their model, highly persistent earnings shocks arriving each period are needed to produce the rise in consumption dispersion with age documented by Deaton and Paxson (1994). In fact, the degree of persistence in shocks to earnings that is required to reproduce the consumption dispersion observations is roughly the persistence that the authors estimate. In their model about half of the variation in lifetime earnings or in lifetime utility is due to differences in shocks received over the lifetime.[6]

We note three difficulties related to this view. First, there is a much smaller rise in US consumption dispersion with age over 1980-1997 than over the period 1980-90 –see

[5]Models with these properties have been used in the macroeconomic literature on economic inequality. Some of this literature is surveyed in Quadrini and Rios-Rull (1997). Recent papers in the literature include Castañeda, Diaz-Jimenez and Rios-Rull (2003), Conesa and Krueger (1999), DeNardi (2004), French (2004), Heathcote, Storesletten and Violante (2005b), Imrohoroglu, Imrohoroglu and Joines (1995), Huggett (1996), Huggett and Ventura (1999), Krueger and Perri (2005), Quadrini (2000) and Storesletten et. al. (2004), among many others.

[6]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.

Deaton and Paxson (1994) and Heathcote et. al. (2005a). Thus, very persistent shocks may no longer be essential within their model. Second, their model may over estimate the importance of idiosyncratic earnings risk after entering the labor market because all the rise in earnings dispersion with age is attributed to the accumulated effect of shocks. Baker (1997) and Guvenen (2005) provide evidence for economically important differences in permanent earnings growth rates across people. These permanent growth differences emerge naturally in the human capital framework we explore. Third, and most importantly, the model considered by these authors may not be 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 has nothing to say about how policy may affect inequality in lifetime labor earnings or may affect welfare through earnings. Models with exogenous wages (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 these wage differences via human capital theory.

The paper is organized as follows. Section 2 presents the model we analyze. Section 3 documents US earnings distribution facts and infers properties of shocks from wage rate data. Section 4 explains how model parameters are set. Section 5 and 6 presents results for earnings distribution dynamics produced by the model and for the decomposition of the variation in lifetime inequality into sources.

2 The Model

We study the decision problem of an agent in an otherwise standard incomplete-markets framework, but modified so that earnings are endogenously determined by human capital decisions.[7] An agent's preferences over consumption allocations are determined by a calculation of expected utility as indicated below. Consumption $c_j(z^j)$ in any period j is

[7]The model is a generalization of the Ben-Porath (1967) model 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. Work by Williams (1978), Krebs (2004) and Pries (2004) also analyzes human capital investment under idiosyncratic risk. Our work differs in how the model is related to data and in its focus on lifetime inequality, among other differences.

risky as it depends on the j-period history of shocks z^j. $P(z^j)$ is the probability of history 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 period shock realizations.

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

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.

$$V(h_1, k_1, a) \equiv \max E[\sum_{j=1}^{J} \beta^{j-1} u(c_j)]$$

subject to

(1) $\sum_{j=1}^{J} c_j(z^j)/(1+r)^{j-1} \leq k_1(1+r) + \sum_{j=1}^{J} e_j(z^j)/(1+r)^{j-1}$

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

(3) $h_{j+1}(z^{j+1}) = z_{j+1} F(h_j(z^j), l_j(z^j), a)$

(4) $L_j(z^j) + l_j(z^j) = 1$

In this decision problem the agent faces a present value budget constraint. The present value budget constraint condition (1) states that the present value of consumption over the lifetime is no more than lifetime wealth which equals the present value of earnings plus the value of initial wealth. This equation holds for each J-period history of shocks. Earnings $e_j(z^j)$ before a retirement age J_R equal the product of a human capital rental rate R_j, an agent's human capital $h_j(z^j)$ and the fraction of available time $L_j(z^j)$ put into market work. Earnings are zero at and after the retirement age. An agent's future human capital is determined (see condition (3)) 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 $h_j(z^j)$, time devoted to human capital production $l_j(z^j)$ 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 over the life cycle, 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 over the lifetime - agents have different learning abilities and agents have different shock realizations over the lifetime. The next section characterizes empirically the rise in US earnings dispersion for cohorts of males. The third key feature is that 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 cycle because optimal human capital investment goes to zero. The next section develops the logic of this point and estimates parameters of a parametric shock process from the relevant moments of wage rate data.

3 Data and Empirical Analysis

In this section we use US data to address two issues. First, we characterize how mean earnings and a measure of earnings dispersion evolve with age for a cohort of individuals. Second, we estimate a stochastic process for human capital shocks by using wage rate data.

3.1 Age Profiles

The age profiles are based on earnings data from the PSID 1969-2003 family files. We consider males between the ages of 23 and 60. We utilize earnings of males who are the head of the household. We apply several additional selection restrictions. We restrict agents to work no less than 520 hours and no more than 5820 per year. We also restrict the agent to earn at least 4000 dollars (in 1968 prices). 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

[8] 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.

23-60, we have at least 100 observations in each age-year bin with which to calculate age and year-specific earnings statistics. 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 we analyze. 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. The earnings data can be viewed as being generated by several factors, that we name as cohort, time, and age effects. Ultimately, we are interested in the age effect, as they provide the clean measure of earnings and inequality per age which the model should match. However, as described in more detail below, this measure depends crucially on the identifying assumptions regarding cohort and time effects. To introduce notation we denote a cohort as $s = t - j$ that is agents who were born in year $t - j$.[9] 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.

$$e_{j,t} = \alpha_s \beta_j \gamma_t \epsilon_{j,t} \tag{1}$$

$$\tilde{e}_{j,t} = \tilde{\alpha}_s + \tilde{\beta}_j + \tilde{\gamma}_t + \tilde{\epsilon}_{j,t} \tag{2}$$

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 identi-

[9]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.

fication problem is well known in the econometrics literature.[10] In effect any trend in the data can be arbitrarily reinterpreted as a year (time) trend or alternatively as trends in ages and cohorts.

Given this problem, we provide two alternative measures of the age effects. These correspond to case 1 in which we assume cohort effects and thus restrict $\gamma_t = 0\ \forall t$ and case 2 in which we assume only time effects and therefore set $\alpha_s = 0\ \forall s$. For the cohort effect case we use ordinary least squares to estimate the coefficients $\tilde{\alpha}_s$ and $\tilde{\beta}_j$. Each 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 $T = 2003 - 1969$. For the time effects case each 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.

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.

[Insert Figure 1]

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. Specifically, 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.[11] Then the age profiles for the variance of log earnings are derived from the following regressions.

$$v_{j,t} = \alpha_s^v + \beta_j^v + \gamma_t^v + \epsilon_{j,t}^v \tag{3}$$

[10]See, for example, Weiss and Lillard (1978), Hanoch and Honig (1985) and Deaton and Paxson (1994) among others.

[11]More specifically, we use 5 year age bins centered at age j to compute these statistics, which is the same construction as for the analysis of mean earnings.

$$g_{j,t} = \alpha_s^g + \beta_j^g + \gamma_t^g + \epsilon_{j,t}^g. \tag{4}$$

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. Figure 2 plots the variance of log earnings while Figure 3 plots the age effects for the Gini measure. In each case, the solid line corresponds to cohort effects whereas the dashed line corresponds to time effects. In Figure 2 we see that the cohort effect view implies a rise of about 0.4 from age 23 to 60 while the time effects imply a much smoother rise in variance of log earnings of only about 0.2. The same qualitative pattern can be seen in Figure 3.

[Insert Figures 2 and 3]

3.2 Human Capital Shocks

The model implies that an agent's wage rate, measured as market compensation per unit of work time, is precisely equal to 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.

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 a 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. This is the first equation below. This equation then implies that the three cross-sectional moment conditions below hold. These moments are based on the assumption that human capital shocks $\hat{z}_t$ are independent across time and people and that $\hat{z}_t \sim N(\mu, \sigma^2)$ and that measurement errors ϵ_t are independent and identically distributed across time and people with variance $Var(\epsilon_t) = \sigma_\epsilon^2$. We also assume that measurement errors and human capital shocks are jointly independent. Therefore the s period wage growth follows the first equation below. The remaining equations put restrictions on the key variance parameter coming from the moments of wage rate data. Note that the s-period growth rate in wages can be written as,

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

which leads to the following moments,

$$E[y_{t,s}] = \hat{R}_{t+s} - \hat{R}_t + s\mu \quad (5)$$

$$Var(y_{t,s}) = s\sigma^2 + 2\sigma_\epsilon^2 \quad (6)$$

$$Cov(y_{t,s}, y_{t,r}) = r\sigma^2 + \sigma_\epsilon^2 \text{ for } r < s \quad (7)$$

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 learning is a reasonable approximation. The focus on older workers solves 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 use wages defined as total labor earnings divided by total hours for male head of household who are between the ages of 55 to 65 (or 50 to 60). We impose additionally the same selection criteria as those presented in Section 3.1 for earnings. For each year of the PSID we follow these agents for at least three years (i.e., the maximal s, $\bar{s} = 2$), this allows us to utilize moments in equation (5)-(6). When equation (7) is also utilized, we do so by using all covariance permutations.[12]. In estimation we use cross sectional moments aggregated across panel years. That is, for each year we generate the sample analog to the moments in equation (5)-(7), that is $\frac{1}{N_t}\sum_{i=1}^{N_t} y_{t,s}^i$ and $\frac{1}{N_t}\sum_{i=1}^{N_t}(y_{t,s}^i - \mu)^2$ and $\frac{1}{N_t}\sum_{i=1}^{N_t}(y_{t,s}^i - \mu)(y_{t,r}^i - \mu)$. 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: Estimation Human Capital Shocks

Min-Age	Max-Age	Sample	N	$\bar{s}$	Cov	σ	S.E.(σ)
55	65	1969-2004	252	3	1	0.074	(0.039)
55	65	1969-2004	252	3	0	0.075	(0.043)
55	65	1969-2004	252	2	1	0.079	(0.041)
55	65	1969-1994	236	2	0	0.071	(0.042)
50	60	1969-2004	472	2	1	0.063	(0.040)
50	60	1969-2004	472	2	0	0.062	(0.042)
50	60	1969-1994	441	2	0	0.054	(0.043)

The entries provide the estimates for σ for various samples which vary across time period and the cut off criteria for ages. The first and second column provide the minimum and maximum respective age in the sample, the column labelled sample refers to which PSID years are included. The column N refers to the median of number of observation across panel years. The column denoted $\bar{s}$ refers to the maximal s value used. The column Cov indicates whether covariance moments (equation (7)) are included in estimation, in which case all possible covariance permutations for $r < \bar{s}$ are employed. The column labelled $S.E.$ refer to standard errors.

Table 1 provides the estimation results for various data selections which differ in the age criteria used and the sample period. The table demonstrates a tradeoff between resorting to older ages (more likely to capture no investment in human capital – a feature underlying our analysis above) and the number of observations. Overall, the results are quite consistent across different selections. The estimated magnitude of shocks is

[12]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.

somewhat smaller relative to their counterparts when data on earnings is used.

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. The parameters for the functional forms that we consider have been estimated in the consumption and saving literature and in the human capital literature. 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 from the previous section. In choosing this initial distribution, we take all other model parameters as given.

Model parameter values are summarized in Table 2. The top panel in Table 2 highlights parameters that we set based upon other studies. We set the model period to be a year and let 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 $r = 0.042$ and equals 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 is around the middle of the estimates of this parameter from 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. Estimates for the elasticity parameter α governing the production of new human capital is set to lie in the range of estimates $(0.5, 1.0)$ from the literature which is surveyed by Browning et. al. (1999, Table 2.3- 2.4).

The setting of parameter values in the lower panel of Table 2 is described next. Here we set parameters to values we estimated in section 3 or so that the model best matches the US earnings distribution facts we documented in section 3. The growth rate g of the rental rate of human capital is set equal to zero. In future versions of this paper we will set g equal to the annual geometric mean growth rate of mean male earnings in the US.[13] Human capital shocks are assumed to be independent across periods and to be lognormally distributed. Specifically, we assume that $\log(z) \sim N(\mu, \sigma^2)$. The variance σ^2 used in the benchmark analysis is set equal to the value $\sigma^2 = 0.0062$ estimated in section 3. This corresponds to the largest value for σ in Table 1. Thus, a one standard deviation shock is approximately an 8 percent change in an agent's human capital stock in any period. We also explore the sensitivity of the benchmark results to higher and lower settings of σ^2. The mean μ of the log human capital shock is set 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 in section 3. 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^2, so that this holds.

We restrict the initial distribution of human capital, assets and learning ability to lie in a parametric class. Specifically, we assume that these initial conditions are jointly lognormally distributed so that $\log(x) \sim N(\mu_x, \Sigma)$. 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. In the current version of the paper, we set initial financial wealth to zero for all agents. In future versions, the mean and variance parameter governing log wealth holding will be restricted to match the (scaled) empirical values for young males in the PSID. The next section highlights properties of the initial distribution that best matches these data.

[13] Discuss steady state growth justification.

Table 2: 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	$R_j = (1+g)^{j-1}$ $g = 0$
Human Capital Shocks	z	$\log(z) \sim N(\mu, \sigma^2)$ $\sigma^2 = 0.0062$ $\mu = -0.0284$
Distribution of Initial Conditions	ψ	$x \equiv (h_1, k_1, a) \sim \psi \equiv LN(\mu_x, \Sigma)$ see following section

5 Earnings Facts

In this section, we report on the ability of the model to reproduce earnings facts we documented in section 3. We first show the model implications for the evolution of the earnings distribution for a cohort. We then assess the importance of learning ability and human capital risk in generating these facts. From now on, we concentrate on the case in which all individuals start their life cycle with zero wealth.

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 below. Figure 4 shows that the model is able to produce the qualitative properties of mean earnings and earnings dispersion we document in US data.

[Insert Figure 4 a-b Here]

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. We note that 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 absense 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 for producing the results in Figure 4 by alternatively eliminating ability differences or eliminating shocks, holding all other parameters at their previous values. To 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 in this way, we do not alter any agent's initial human capital or initial assets. Figure 5 shows the results for the case of $\sigma^2 = 0.01$. We find 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 which largely cancel out except towards the end of the working lifetime. First, human capital risk leads ex-ante identical agents to have ex-post differences in human capital and earnings. Second, the model has a force which leads to decreasing dispersion in human capital and earnings for a cohort as the cohort ages. 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. Huggett et. al. (2006, Proposition 1) establish that this holds for the human capital production function considered in this paper. Thus, under these circumstances, human capital dispersion decreases for a cohort as the cohort ages.[14]

[Insert Figure 5 Here]

To highlight the role of human capital risk, we eliminate idiosyncratic risk altogether by setting $\sigma^2 = 0$, while adjusting the mean log shock μ to keep the mean shock level constant in both cases for the a *given* initial distribution, which we set equal to the one consistent with $\sigma^2 = 0.01$. Figure 6 a-b illustrate that removing idiosyncratic risk leads to a counter-clockwise rotation of the mean earnings profile. This also leads to a U-shaped earnings dispersion profile.

[Insert Figure 6 a-b Here]

[14]The marginal cost of spending an extra unit of time producing new human capital is proportional to an agent's human capital at any age. The marginal benefit is related to the extra units of new human capital produced. Agents with higher human capital produce more extra units of human capital per marginal unit of time than agents with lower human capital but this marginal productivity is less than proportional to an agents human capital. This is because α is less than one in the law of motion for human capital (i.e. $F(h, l, a) = h + a(hl)^{\alpha}$). Thus, learning ability held fixed, agents with higher human capital spend less time in human capital production compared to agents with lower amounts of human capital.

When idiosyncratic risk is eliminated, human capital accumulation becomes more attractive for risk-averse individuals. Thus, all else equal, individuals' response dictates lower mean earnings early in life from the greater time spent accumulating human capital and a higher growth rate in mean earnings early in life. The result is a counter-clockwise movement in the mean earnings profile.[15] 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. 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. Subsequently, earnings dispersion increases because of the differences in slopes of age-earnings profiles across ability levels.

5.2 Properties of the Initial Distribution

What are the properties of the initial distributions that best reproduce the earnings facts? To better highlight the role of human capital risk, Table 2 shows these properties for three values of the variance of the idiosyncratic shocks. The first column displays results for σ^2=0.0026. The second column displays results for $\sigma^2 = 0.0062$, which is the benchmark model. The last column displays results for a higher case of dispersion, $\sigma^2 = 0.01$. In each case, we adjust μ so that the decline in mean earnings matches the observed value at the end of the life cycle.

Table 2 shows that the distribution of initial conditions moves in a systematic way with the shock variance. We focus the discussion on changes in the initial distribution related to learning ability. The reason for this is that in the next section we find that variation in learning ability is the most important source of variation in either expected lifetime utility or expected lifetime wealth among the initial conditions present in the

[15]This is effectively the central result of Lehvari 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.

model. Table 2 shows that as σ^2 increases the initial distributions that best reproduce the earnings facts require higher levels of mean initial learning ability and lower levels of ability dispersion.

Table 2: Properties of Initial Distributions

Statistic	$\sigma^2 = 0.0026$	$\sigma^2 = 0.0062$	$\sigma^2 = 0.01$
Mean Learning Ability (a)	0.314	0.334	0.356
Coefficient of Variation (a)	0.261	0.231	0.201
Mean Initial Human Capital (h_1)	119.2	118.2	114.6
Coefficient of Variation (h_1)	0.514	0.474	0.406
Correlation (a, h_1)	0.771	0.741	0.706

What accounts for these changes in the initial distributions? To understand these changes, the reader should 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 Lehvari 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. Thus, higher mean learning ability counteracts 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 straightforward. This is because human capital risk is itself a source of increased earnings dispersion over the life cycle. Thus, greater human capital risk leaves less room for ability differences in accounting for the rise in earnings dispersion with age over the life cycle.

6 Lifetime Inequality

6.1 Analysis of the Benchmark Model

We decompose the variance in lifetime inequality into variation due to initial conditions versus variation due to shocks. This is done both for lifetime utility and for lifetime wealth. Such a decomposition makes use of the fact that any 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 by construction, the total variance in the random variable equals the sum of the variance in the conditional mean plus the mean of the variance around the conditional mean.[16]

Table 3 decomposes lifetime inequality into the sources highlighted by the model. We focus on the three cases for the variance of idiosyncratic shocks: the benchmark case $\sigma^2 = .0062$ as well as a higher and a lower variance case. The top panel of Table 3 expresses the fraction of the variance in lifetime utility and lifetime wealth due to variation across agents in initial conditions. The bottom panel of Table 3 expresses the fraction of the variance in expected lifetime utility and expected lifetime wealth due to variation across agents in learning ability. Thus, in the bottom panel of Table 3 the remaining source of variation is due to variation in human capital and initial wealth, conditional on learning ability. We note that the variance decomposition of lifetime utility is unchanged if the expected utility function is represented by an affine transformation of the original utility function.

Table 3 shows that for the benchmark variance estimated in section 3 that the majority of the variation in lifetime utility or lifetime wealth is due to variation in initial conditions.[17] Specifically, 82 percent of the variation in lifetime utility and 78 percent of the variation in lifetime wealth is due to initial conditions. We also find that even when

[16] Formally, let $f(x,y) = E[f(x,y)|x] + (f(x,y) - E[f(x,y)|x]) = E[f(x,y)|x] + \epsilon(x,y)$. Orthogonality then implies $var(f(x,y)) = var_x(E[f(x,y)|x]) + E_x[var_y(\epsilon(x,y))]$. In our application of this result, x captures initial conditions (i.e. human capital, learning ability and financial wealth) and y captures shock histories.

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

the shock variance is set to $\sigma^2 = .01$ - a higher value than we estimate in US data- initial conditions still account for the majority of the variation in the two measures of lifetime inequality. The bottom panel of Table 3 shows that, among initial conditions, differences in learning ability are more important than differences in initial human capital in accounting for variation in lifetime inequality. We find that differences in learning ability accounts for 68 percent of the variation in expected utility and for 74 percent of the variation in expected wealth.

Table 3: Sources of Lifetime Inequality

Statistic	$\sigma^2 = 0.0026$	$\sigma^2 = 0.0062$	$\sigma^2 = 0.01$
Fraction of Variance in Lifetime Utility Due to Initial Conditions	0.927	0.821	0.706
Fraction of Variance in Lifetime Wealth Due to Initial Conditions	0.904	0.782	0.626
Fraction of Variance in Expected Lifetime Utility Due to Learning Ability	0.690	0.675	0.664
Fraction of Variance in Expected Lifetime Wealth Due to Learning Ability	0.733	0.741	0.706

Another useful way to highlight the relative importance of different initial conditions is to ask the agents in the model how much they need to be compensated to start life with a one standard deviation change in any initial condition and be equally well off, other things equal. We express this compensation in terms of the percentage change in consumption that would be required to leave an agent with given initial conditions with the same expected utility as an agent with a one standard deviation change in the relevant initial condition. We carry out this analysis by focusing on an agent with initial conditions set at the mean log values of human capital, wealth and learning ability $\mu_x = (\mu_h, \mu_k, \mu_a)$. [Results: To Be Provided!]

6.2 How Important is Modelling Social Insurance?

[Results of this section are not yet available!]

One might conjecture that the relative importance of shocks over the life cycle versus initial conditions for lifetime inequality might be sensitive to the structure of the social insurance system. To discuss this issue, we denote lifetime wealth (lifetime earnings plus initial wealth) as $W(z^J; h_1, k_1, a)$ and lifetime net taxes as $T(z^J; h_1, k_1, a)$ so that lifetime net wealth is $W(z^J; h_1, k_1, a) - T(z^J; h_1, k_1, a)$. If the social insurance system has the feature that most of the variation in taxes $T(z^J; h_1, k_1, a)$ is due to initial conditions (h_1, k_1, a) rather than lifetime shock histories z^J and that lifetime taxes and lifetime wealth are negatively correlated, then one may conjecture that the social insurance system serves to reduce the fraction of variation in lifetime utility $U(z^J; h_1, k_1, a)$ due to initial conditions and to increase the fraction of variation due shocks over the lifetime.

To examine this issue we add a stylized social insurance system to the model. The social insurance system consists of an income tax system and a stylized social security system. For this new model we then choose initial conditions to best match the earnings distribution facts highlighted earlier, holding all the other model parameters at the previous values. Table 4 reassess the sources of lifetime inequality originally presented in Table 3.

Table 4: Sources of Lifetime Inequality: Model with Social Insurance

Statistic	$\sigma^2 =$	$\sigma^2 =$
Fraction of Variance in Lifetime Utility Due to Initial Conditions		
Fraction of Variance in Lifetime Wealth Due to Initial Conditions		
Fraction of Variance in Lifetime After-Tax Wealth Due to Initial Conditions		
Fraction of Variance in Expected Lifetime Utility Due to Learning Ability		
Fraction of Variance in Expected Lifetime Wealth Due to Learning Ability		
Fraction of Variance in Expected Lifetime After-Tax Wealth Due to Learning Ability		

References

Abowd, J. and D. Card (1989), On the Covariance Structure of Earnings and Hours Changes, *Econometrica,* 57, 411-45.

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.

Bowlus, A. and J.M. Robin (2003), Twenty Years of Rising Inequality in Lifetime Labor Income Values, University of Western Ontario, manuscript.

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).

Carneiro, P. and J. Heckman (2003), Human Capital Policy, National Bureau of Economic Research Paper 9495.

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. (2005), Learning your Earning: Are Labor Income Shocks Really That Persistent?, manuscript.

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

Keane, M. and K. Wolpin (1997), The Career Decisions of Young Men, *Journal of Political Economy*, 105(3), 473-522.

Keane, M. and K. Wolpin (2001), The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment, *International Economic Review*, November.

Krebs, T (2004), Human Capital Risk and Economic Growth, *Quarterly Journal of Economics,* 118, 709- 744.

Levhari, D. and Y. Weiss (1974), The Effect of Risk on the Investment in Human Capital, *American Economic Review,* 64, 950-63.

Lochner, L. and A. Monge (2002), Human Capital Formation with Endogenous Credit Constraints, NBER working paper 8815.

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

Mirrlees, J. (1971), An Exploration into the Theory of Optimum Income Taxation, *Review of Economic Studies*, 38, 175- 208.

Press, W. et. al. (1992), *Numerical recipes in FORTRAN,* Second Edition, (Cambridge University Press, Cambridge).

Pries, M. (2001), Uninsured Idiosyncratic Risk and Human Capital Accumulation, manuscript.

Storesletten, K., Telmer, C. and A. Yaron (2004), Consumption and Risk Sharing Over the Life Cycle, *Journal of Monetary Economics,* 51, 609- 33.

Williams, J. (1978), Risk, Human Capital, and the Investor's Portfolio, *Journal of Business,* 51, 65-89.

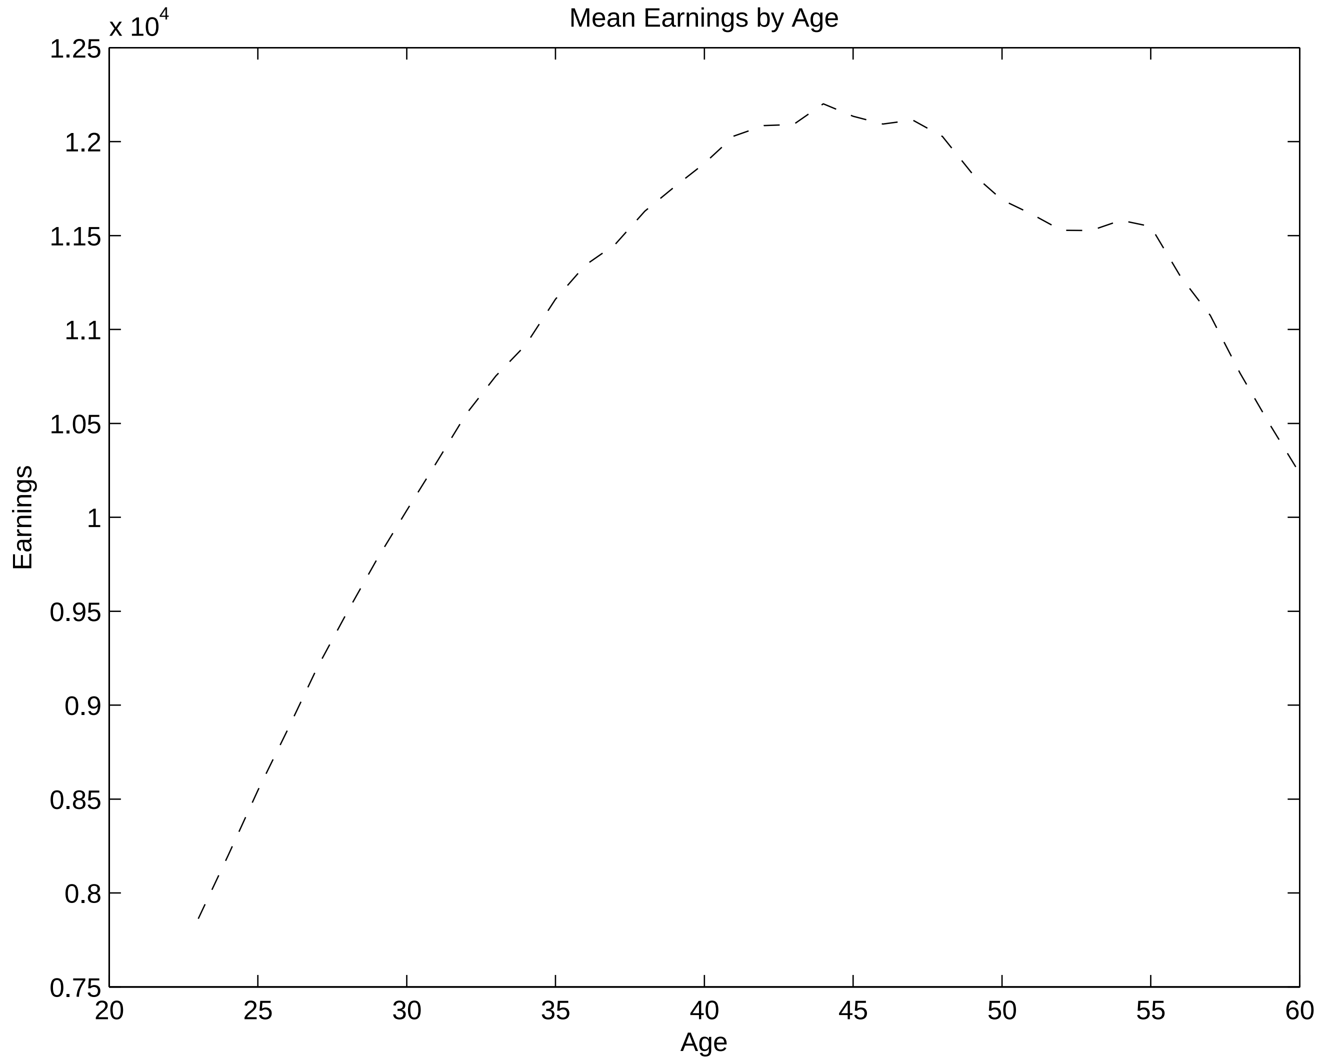

Mean Earnings by Age
x 10^4
1.25
1.2
1.15
1.1
1.05
1
0.95
0.9
0.85
0.8
0.75
Earnings
20
25
30
35
40
45
50
55
60
Age

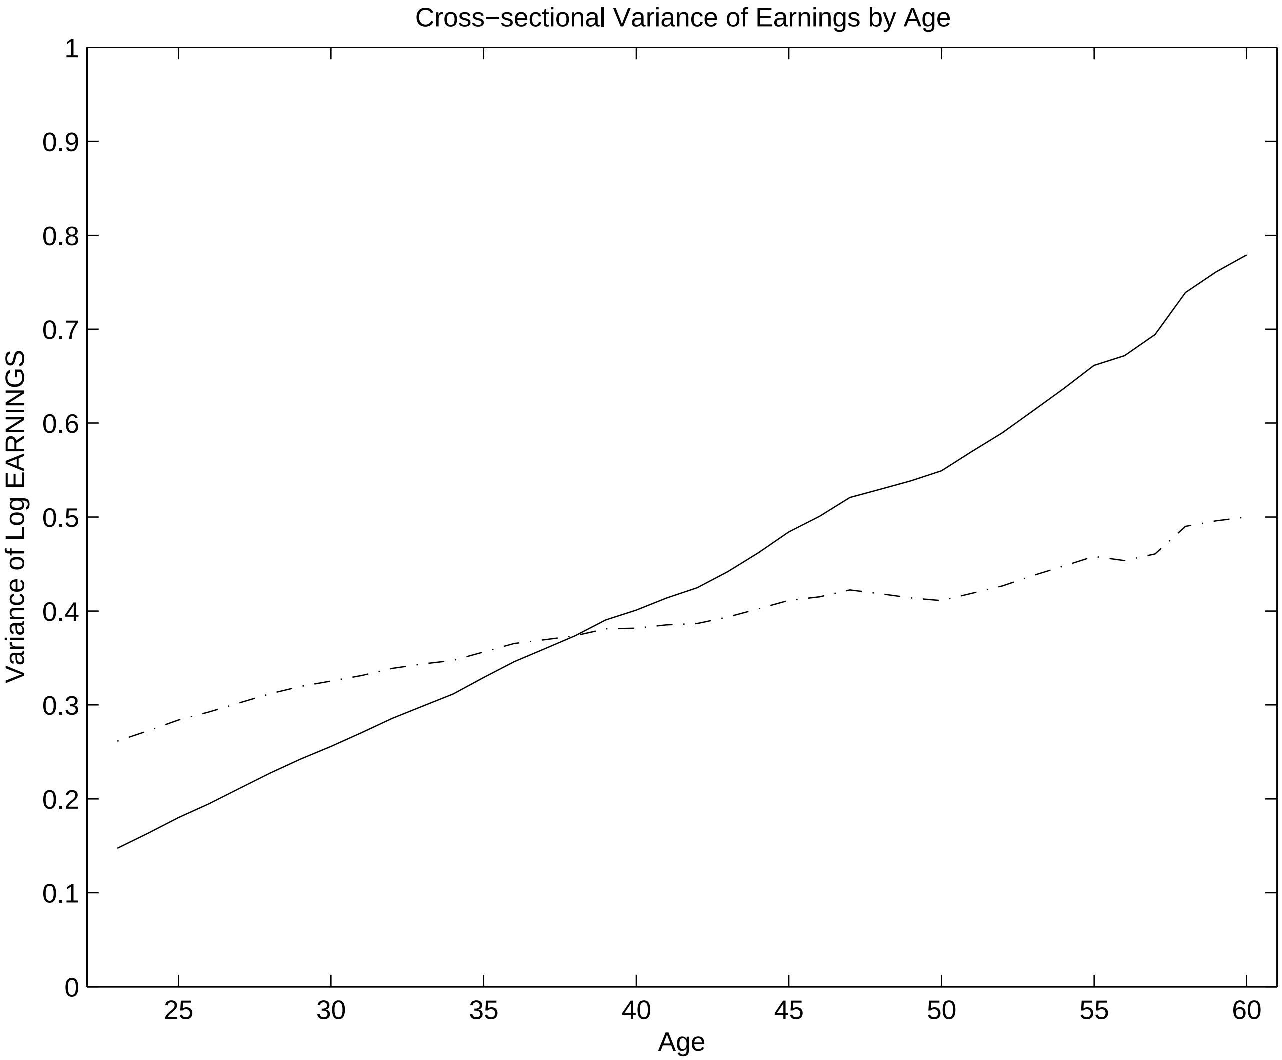
Cross−sectional Variance of Earnings by Age
Variance of Log EARNINGS
Age
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
25
30
35
40
45
50
55
60

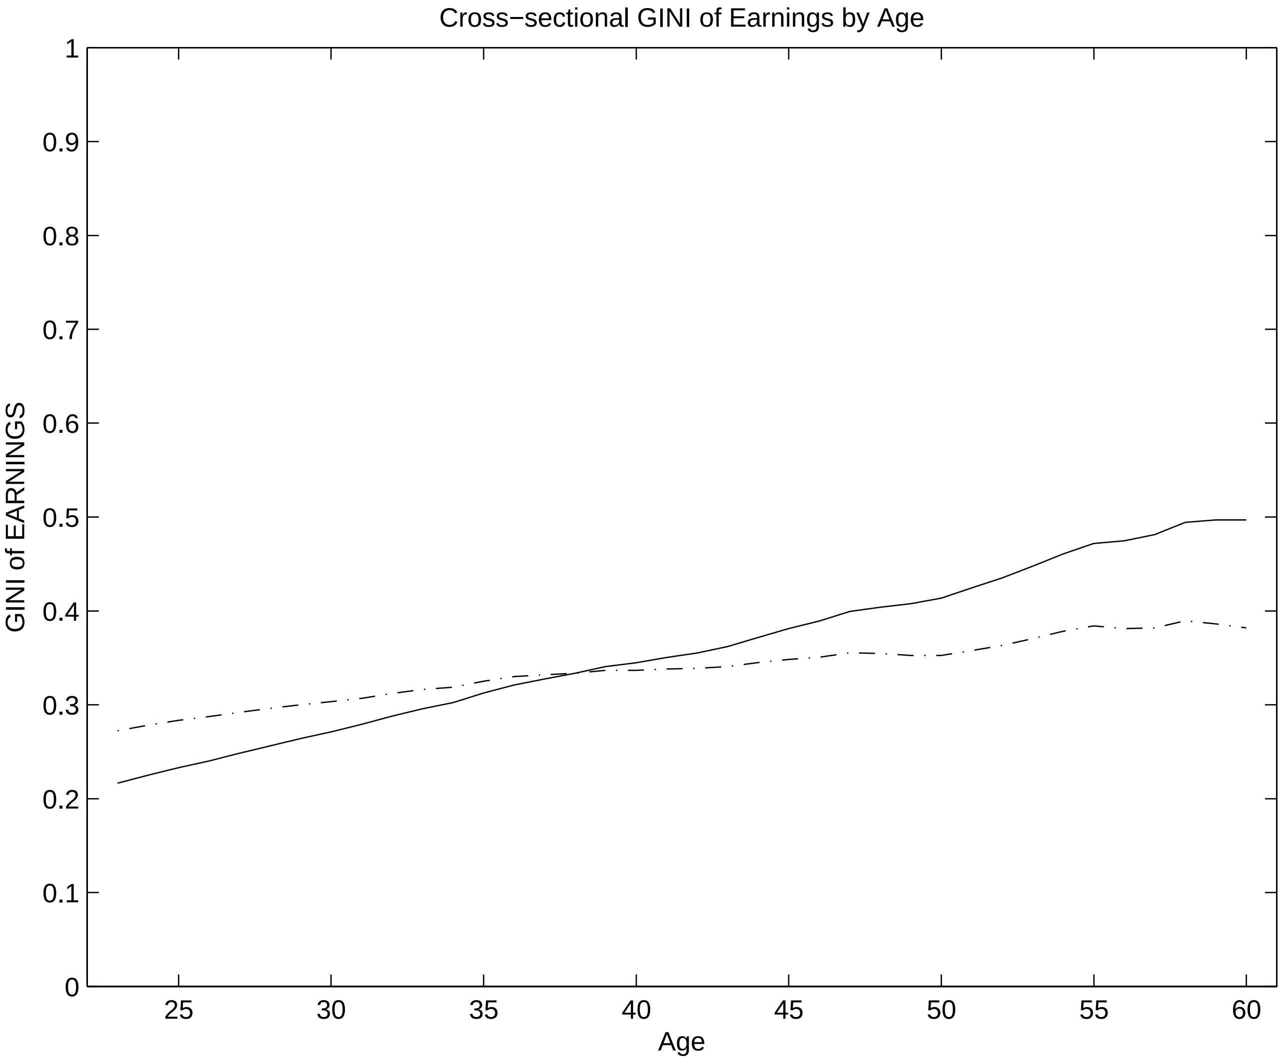

Cross−sectional GINI of Earnings by Age
GINI of EARNINGS
Age
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
25
30
35
40
45
50
55
60

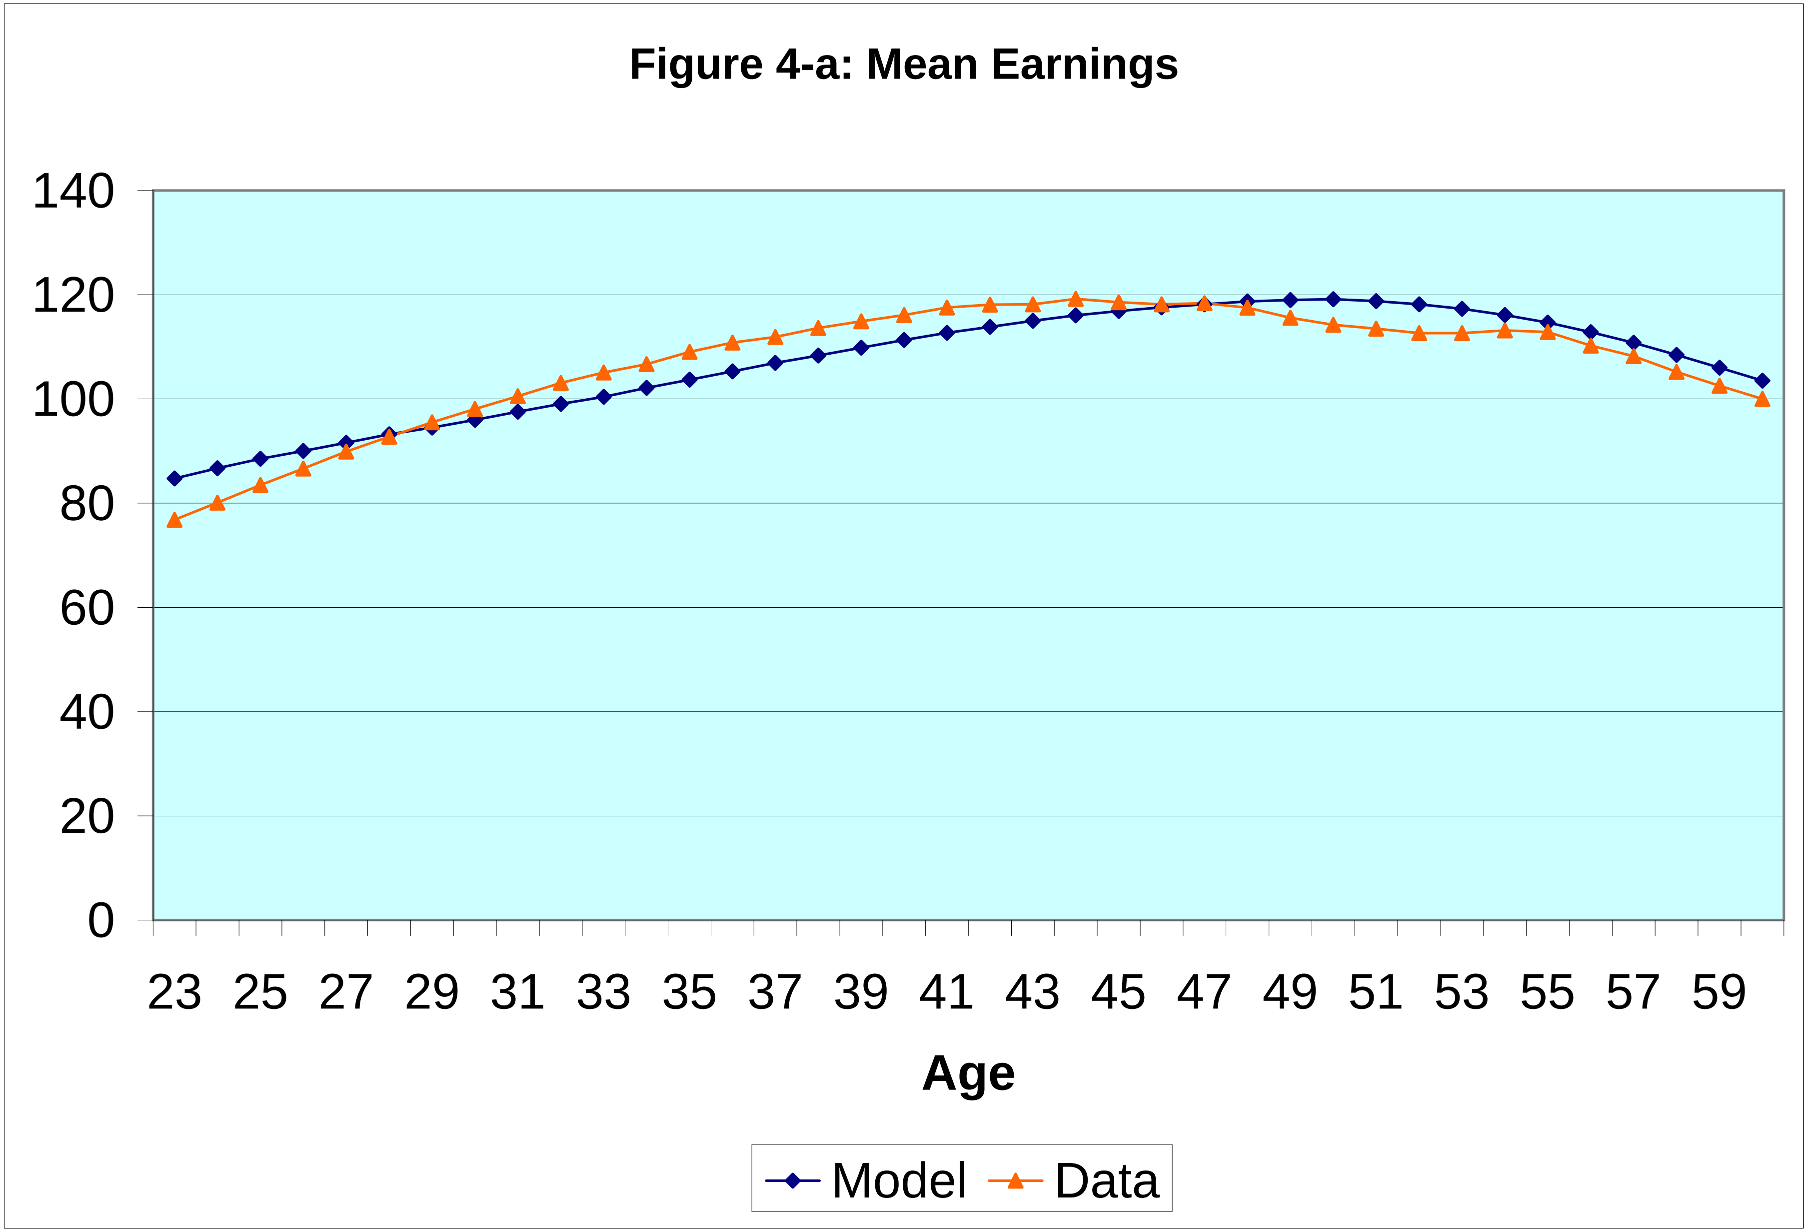

Figure 4-a: Mean Earnings
140
120
100
80
60
40
20
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
Model
Data

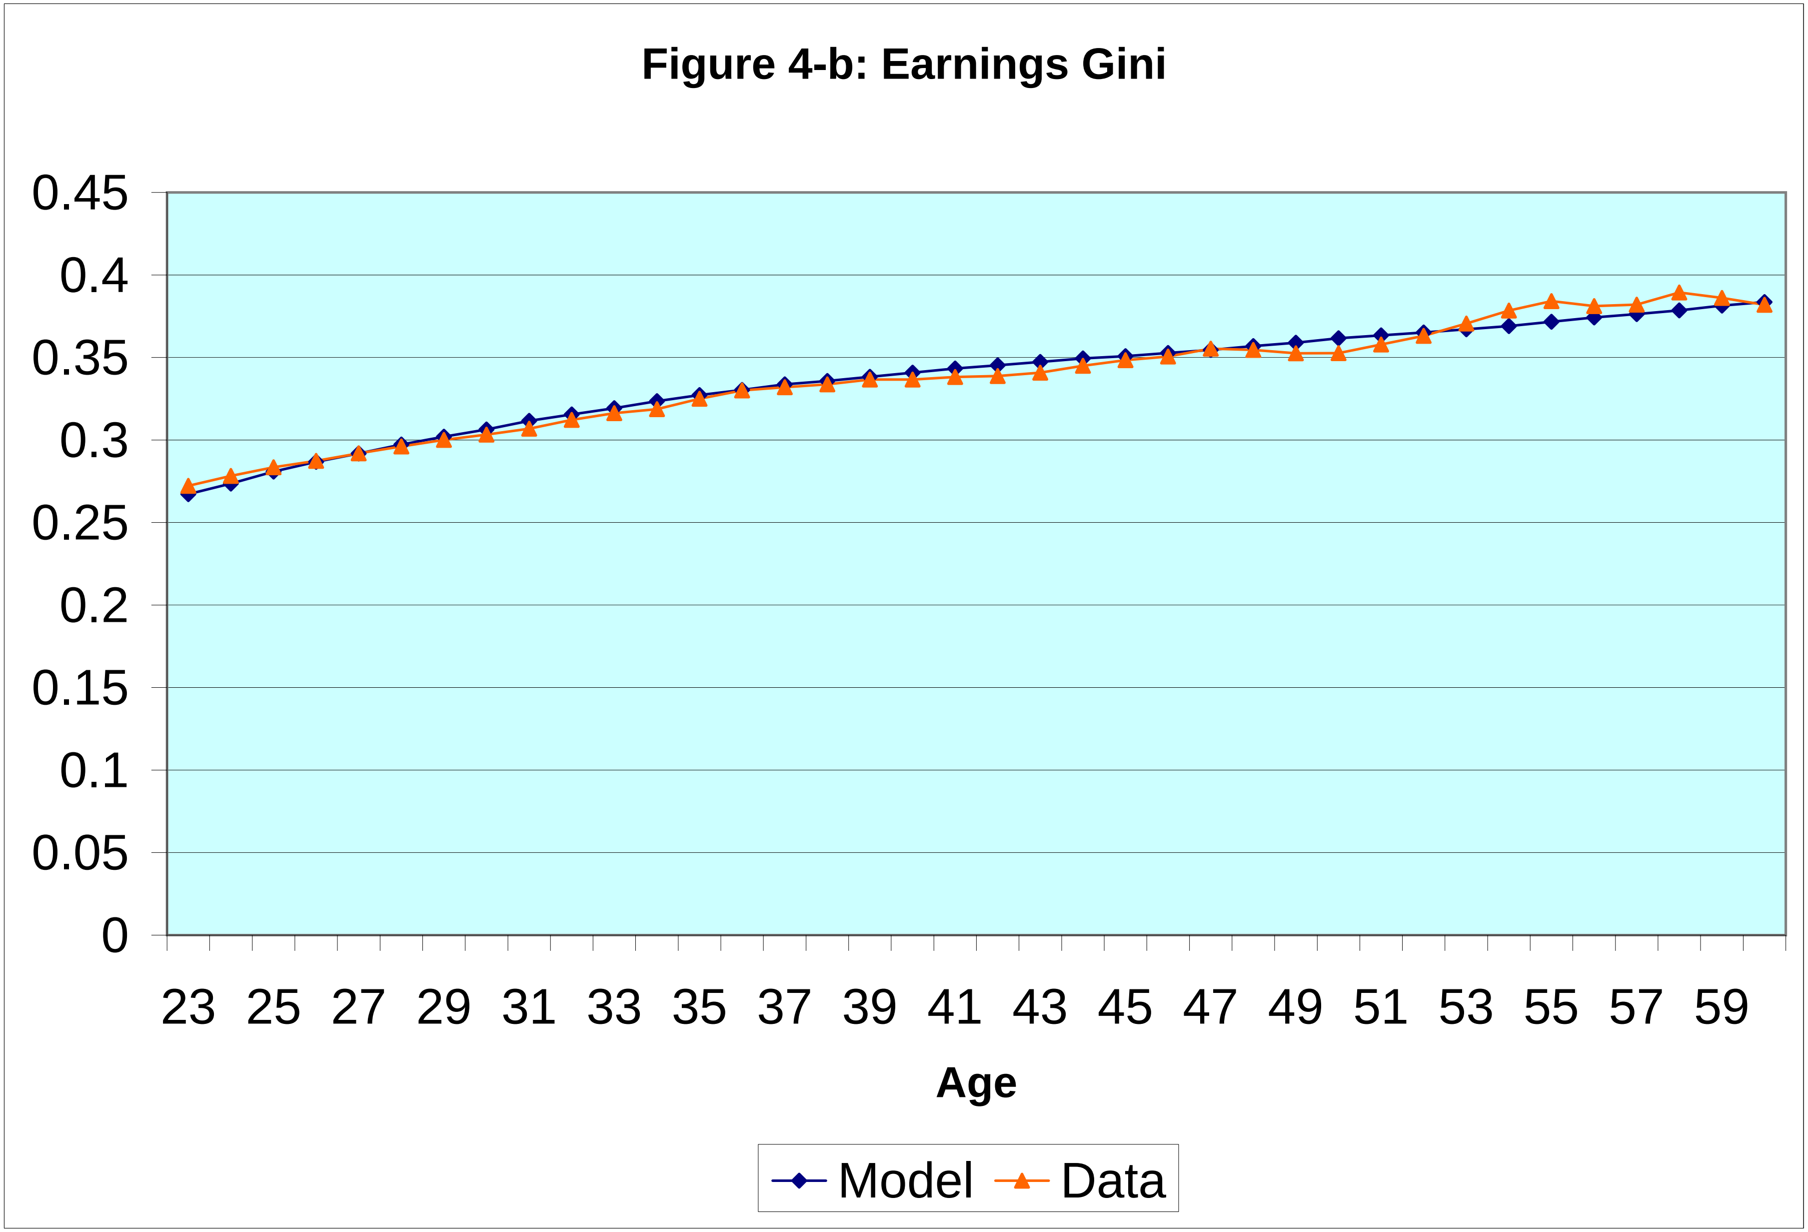

Figure 4-b: Earnings Gini
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
Model
Data

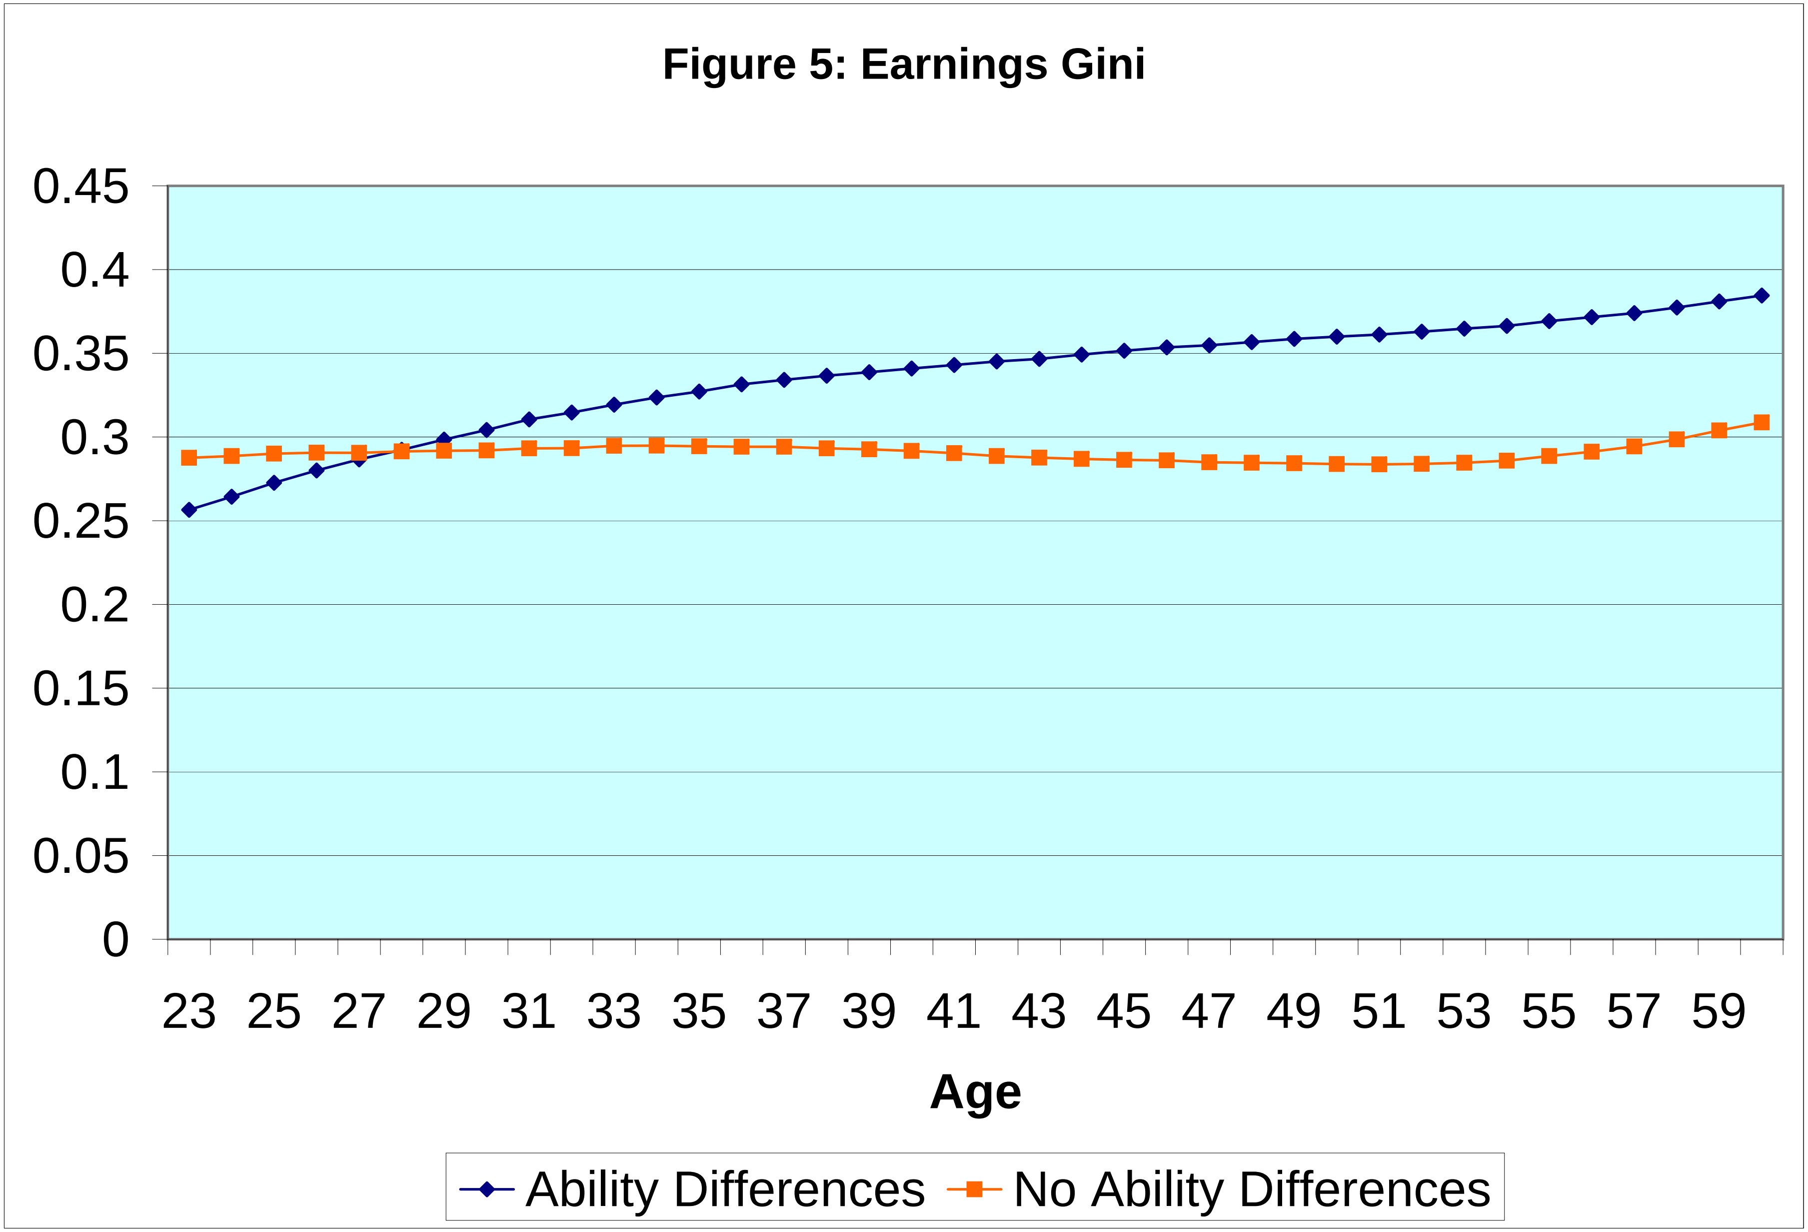

Figure 5: Earnings Gini
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
Ability Differences
No Ability Differences

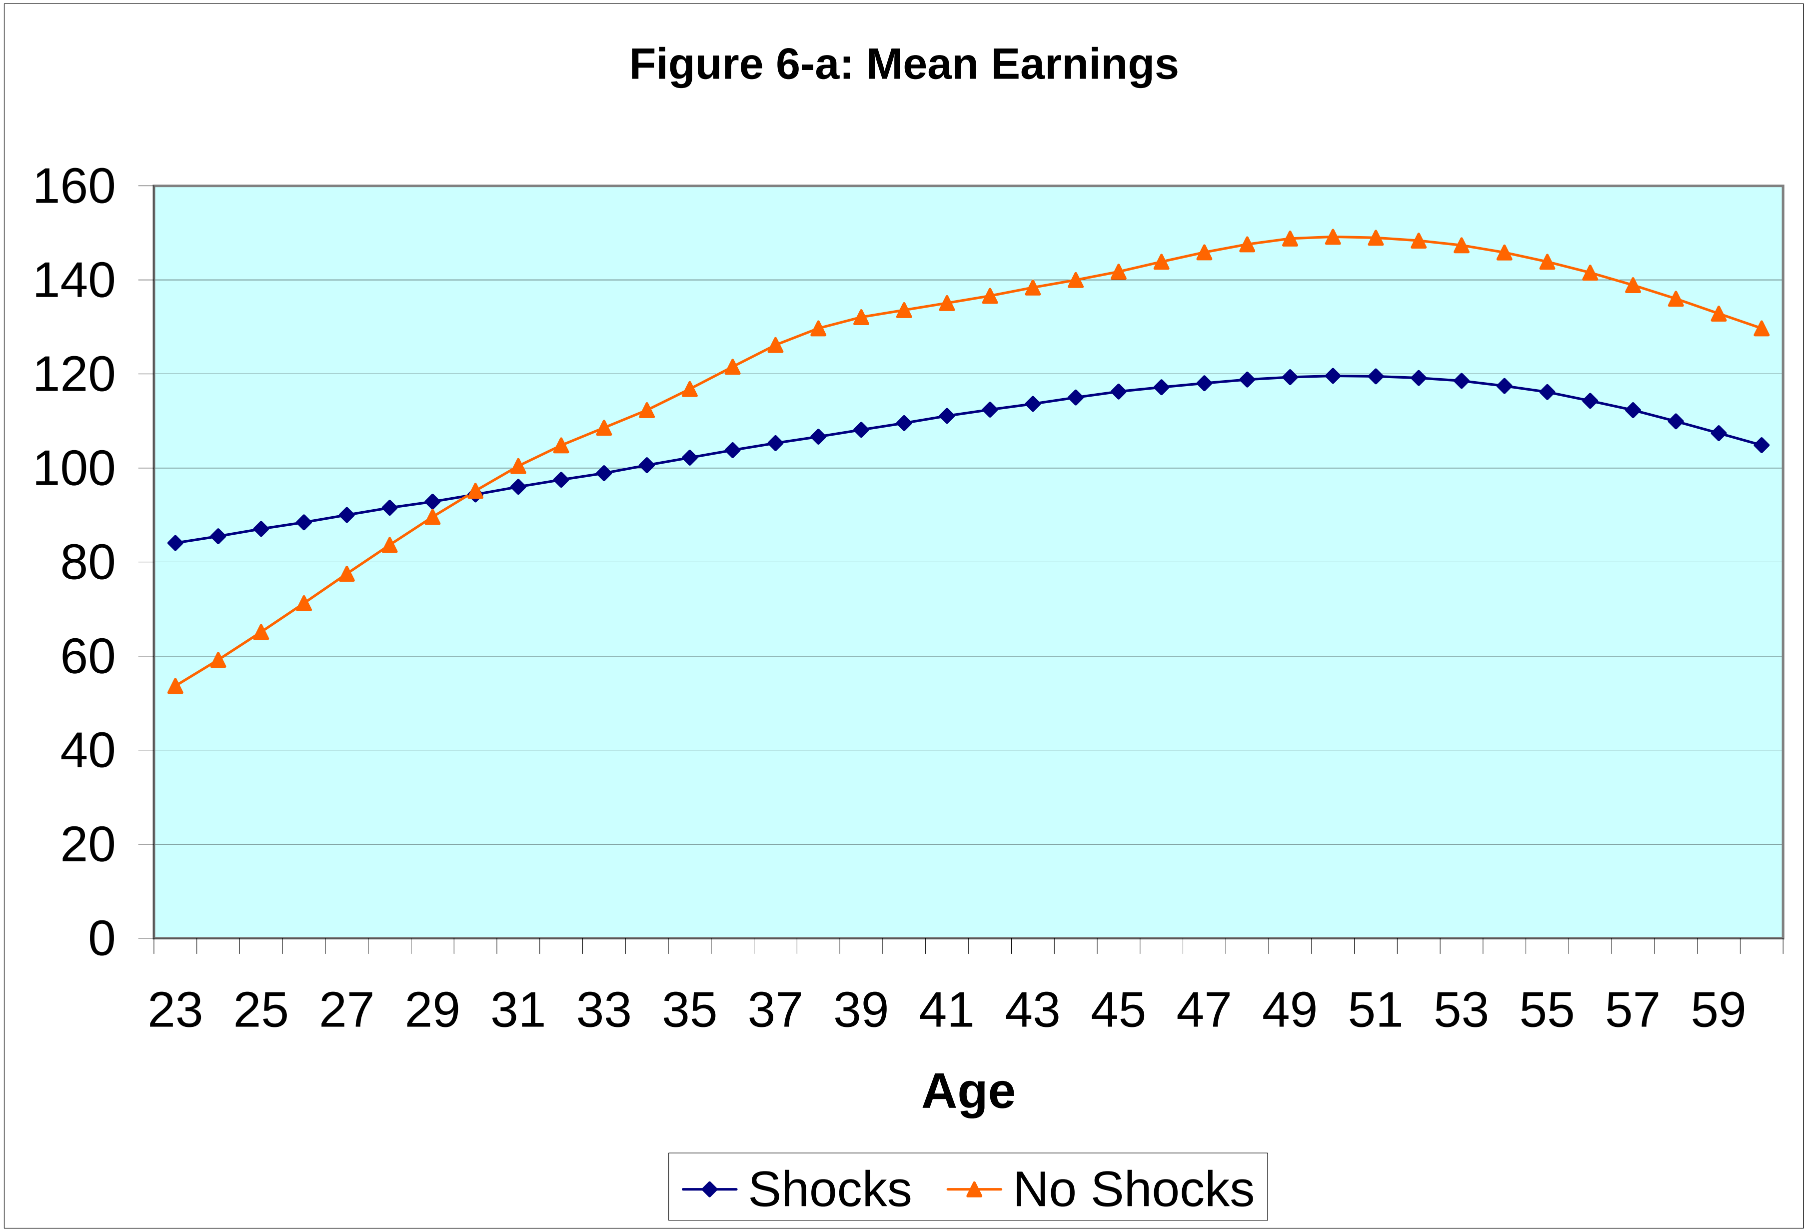

Figure 6-a: Mean Earnings
160
140
120
100
80
60
40
20
0
23 25 27 29 31 33 35 37 39 41 43 45 47 49 51 53 55 57 59
Age
Shocks
No Shocks

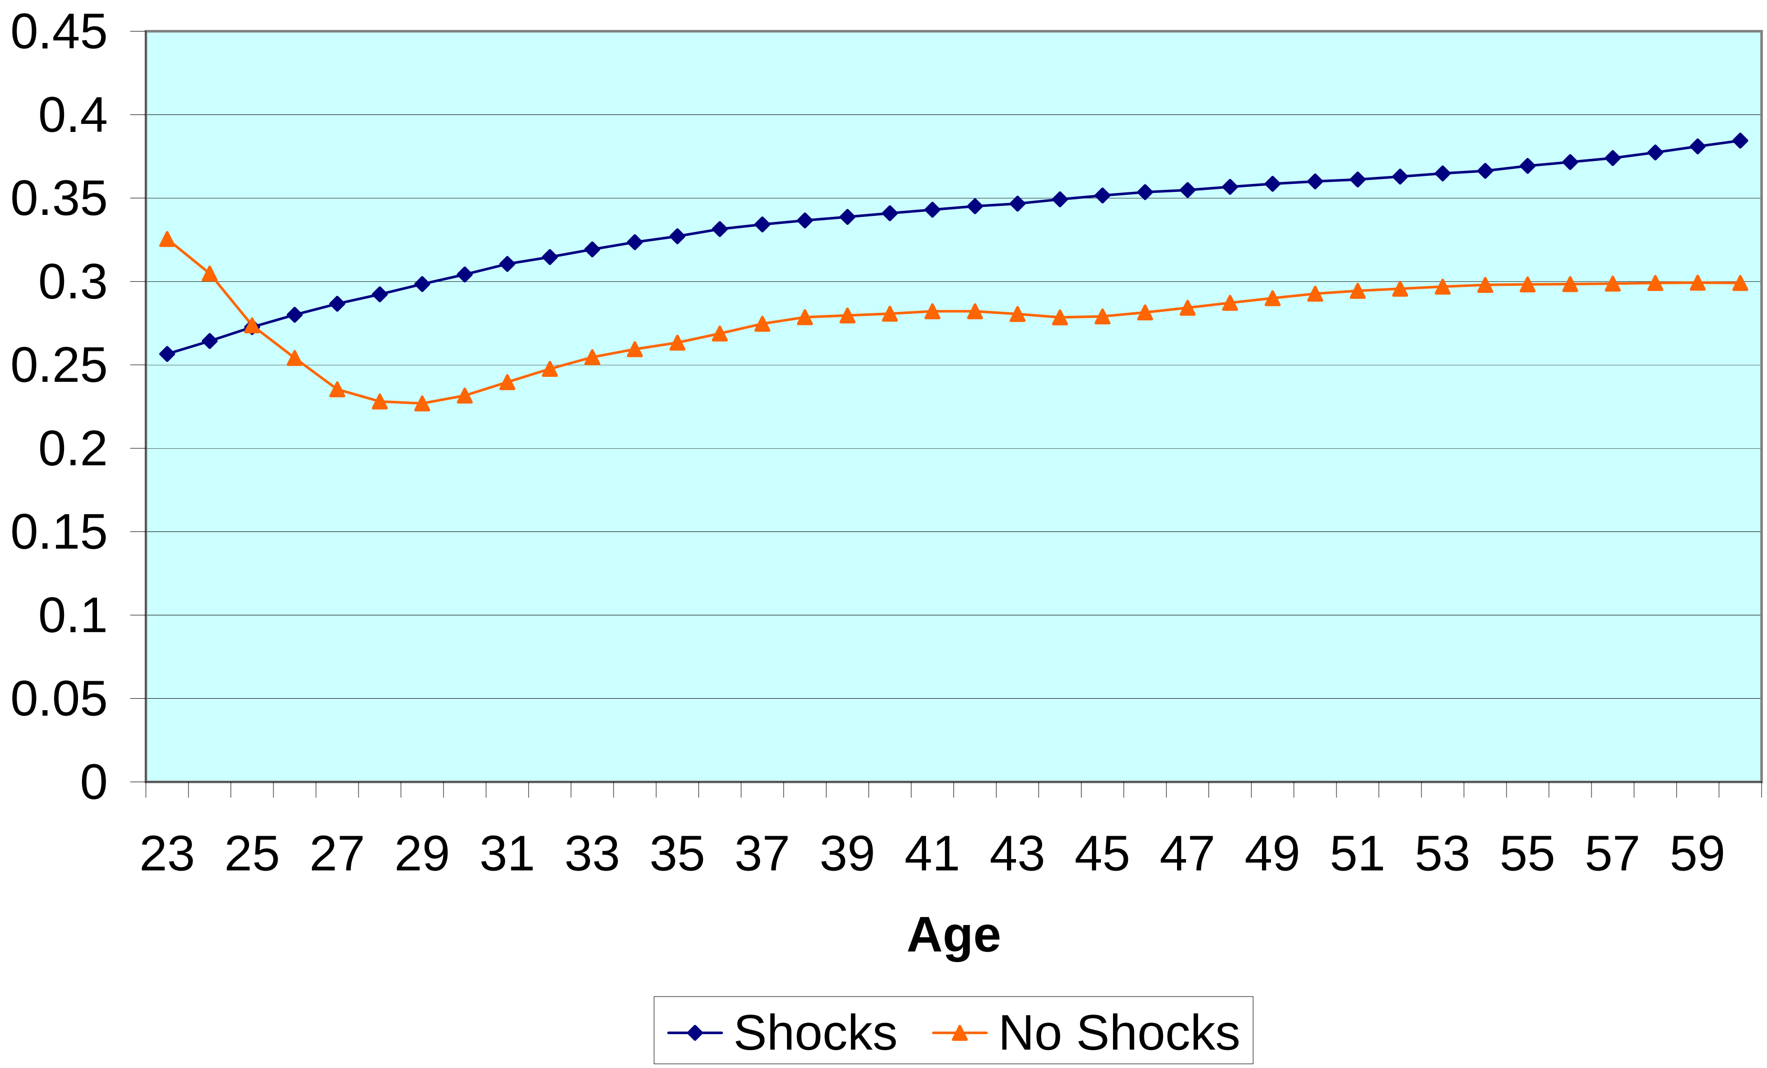

Figure 6-b: Earnings Gini